



IN PARTNERSHIP WITH:
CNRS

Université Rennes 1

**École normale supérieure de
Rennes**

Activity Report 2016

Project-Team CAIRN

Energy Efficient Computing Architectures

IN COLLABORATION WITH: Institut de recherche en informatique et systèmes aléatoires (IRISA)

RESEARCH CENTER
Rennes - Bretagne-Atlantique

THEME
**Architecture, Languages and Compila-
tion**

Table of contents

1. Members	1
2. Overall Objectives	3
3. Research Program	4
3.1. Panorama	4
3.2. Reconfigurable Architecture Design	5
3.3. Compilation and Synthesis for Reconfigurable Platforms	5
4. Application Domains	6
5. Highlights of the Year	7
6. New Software and Platforms	7
6.1. Panorama	7
6.2. Gecos	8
6.3. ID-Fix	9
6.4. Zyggye	9
7. New Results	10
7.1. Reconfigurable Architecture Design	10
7.1.1. Dynamic Reconfiguration Support in FPGA	10
7.1.2. Hardware Accelerated Simulation of Heterogeneous Platforms	11
7.1.3. Optical Interconnections for 3D Multiprocessor Architectures	11
7.1.4. Communication-Based Power Modelling for Heterogeneous Multiprocessor Architectures	12
7.1.5. Arithmetic Operators for Cryptography and Fault-Tolerance	12
7.1.6. Adaptive Overclocking, Error Correction, and Voltage Over-Scaling for Error-Resilient Applications	13
7.2. Compilation and Synthesis for Reconfigurable Platform	14
7.2.1. Adaptive dynamic compilation for low power embedded systems	14
7.2.2. Leveraging Power Spectral Density for Scalable System-Level Accuracy Evaluation	14
7.2.3. Approximate Computing	14
7.2.4. Real-Time Scheduling of Reconfigurable Battery-Powered Multi-Core Platforms	15
7.2.5. Optimization of loop kernels using software and memory information	15
7.2.6. Adaptive Software Control to Increase Resource Utilization in Mixed-Critical Systems	16
8. Partnerships and Cooperations	16
8.1. Regional Initiatives	16
8.2. National Initiatives	16
8.2.1. ANR Blanc - PAVOIS (2012–2016)	16
8.2.2. ANR Ingénierie Numérique et Sécurité - ARDyT (2011-2016)	17
8.2.3. Labex CominLabs - BoWI (2012-2016)	17
8.2.4. Labex CominLabs - 3DCORE (2014-2018)	17
8.2.5. Labex CominLabs - RELIASIC (2014-2018)	17
8.2.6. Labex CominLabs & Lebesgue - H-A-H (2014-2017)	18
8.3. European Initiatives	18
8.3.1. H2020 ARGO	18
8.3.2. ANR International ARTEFaCT	18
8.4. International Initiatives	19
8.4.1. Inria Associate Teams	19
8.4.2. Inria International Partners	19
8.4.2.1. Declared Inria International Partners	19
8.4.2.1.1. LRS	19
8.4.2.1.2. HARAMCOP	19
8.4.2.1.3. SPINACH	19

8.4.2.2. Informal International Partners	19
8.5. International Research Visitors	20
8.5.1. Visits of International Scientists	20
8.5.2. Visits to International Teams	20
9. Dissemination	20
9.1. Promoting Scientific Activities	20
9.1.1. Scientific Events Selection	20
9.1.1.1. General Chair, Scientific Chair	20
9.1.1.2. Chair of Conference Program Committees	21
9.1.1.3. Member of the Conference Program Committees	21
9.1.2. Journal	21
9.1.3. Invited Talks	21
9.1.4. Leadership within the Scientific Community	21
9.1.5. Scientific Expertise	21
9.2. Teaching - Supervision - Juries	22
9.2.1. Teaching	22
9.2.2. Teaching Responsibilities	22
9.2.3. Supervision	23
10. Bibliography	24

Project-Team CAIRN

Creation of the Project-Team: 2009 January 01

CAIRN is located on two campuses: Rennes (Beaulieu) and Lannion (ENSSAT).

Keywords:

Computer Science and Digital Science:

- 1.1. - Architectures
 - 1.1.1. - Multicore
 - 1.1.2. - Hardware accelerators (GPGPU, FPGA, etc.)
 - 1.1.8. - Security of architectures
 - 1.1.9. - Fault tolerant systems
 - 1.1.10. - Reconfigurable architectures
 - 1.1.12. - Non-conventional architectures
- 1.2.5. - Internet of things
- 1.2.6. - Sensor networks
- 2.2. - Compilation
 - 2.2.1. - Static analysis
 - 2.2.4. - Parallel architectures
 - 2.2.5. - GPGPU, FPGA, etc.
 - 2.2.6. - Adaptive compilation
- 4.4. - Security of equipment and software
- 7.12. - Computer arithmetic

Other Research Topics and Application Domains:

- 4.5. - Energy consumption
 - 4.5.1. - Green computing
 - 4.5.2. - Embedded sensors consumption
- 6.2.2. - Radio technology
- 6.2.4. - Optic technology
- 6.6. - Embedded systems
- 8.1. - Smart building/home
 - 8.1.1. - Energy for smart buildings
 - 8.1.2. - Sensor networks for smart buildings

1. Members

Research Scientists

- François Charot [Researcher, Inria, Rennes]
- Olivier Sentieys [Team Leader, Senior Researcher, Inria, HDR]
- Arnaud Tisserand [Senior Researcher, CNRS, Lannion, until Nov. 2016, HDR]
- Tomofumi Yuki [Inria, Researcher, since Nov 2016]

Faculty Members

- Emmanuel Casseau [Professor, Univ. Rennes I, ENSSAT, Lannion, HDR]
- Daniel Chillet [Professor, Univ. Rennes I, ENSSAT, Lannion, HDR]

Steven Derrien [Professor, Univ. Rennes I, ISTIC, Rennes, HDR]
Cédric Killian [Associate Professor, Univ. Rennes I, IUT, Lannion]
Angeliki Kritikakou [Associate Professor, Univ. Rennes I, ISTIC, Rennes]
Patrice Quinton [Professor, ENS Rennes, Rennes, HDR]
Christophe Wolinski [Professor, Univ. Rennes I, Director of ESIR, Rennes, HDR]

Engineers

Philippe Quémerais [Research Engineer (20%), Univ. Rennes I, ENSSAT, Lannion, until Aug. 2016]
Arnaud Carer [Research Engineer (half time), Univ. Rennes I, Lannion]
Pierre Guilloux [Univ. Rennes I, Lannion]
Christophe Huriaux [Inria, Lannion]
Ali Hassan El Moussawi [Inria, Rennes]
Thomas Lefeuvre [Univ. Rennes I, Rennes]
Nicolas Simon [Univ. Rennes I, Rennes, until Oct. 2016]
Nicolas Estibals [Inria, Rennes, until Dec. 2016]
Raphael Bardoux [Inria, Lannion, until Jan. 2016]

PhD Students

Franck Bucheron [DGA, Rennes, from Oct. 2011 (half time)]
Gaël Deest [Univ. Rennes I, MENRT grant, Rennes, from Oct. 2013]
Rengarajan Ragavan [Univ. Rennes I, granted by FP7 FlexTiles, Lannion, from Oct. 2013]
Mai-Thanh Tran [Univ. Rennes I, granted by Brittany Region/CG22, Lannion, from Oct. 2013]
Xuan Chien Le [Inria, granted by Brittany Region/LTC, Lannion, from Oct. 2013]
Benjamin Barrois [Univ. Rennes I, MENRT grant, Lannion, from Oct. 2014]
Gabriel Gallin [CNRS, granted by CominLabs, from Oct. 2014]
Jiating Luo [Univ. Rennes I, granted by China Gov., from Nov. 2014]
Van Dung Pham [Inria, granted by CominLabs, Lannion, from Dec. 2014]
Baptiste Roux [Inria, granted by DGA and Inria, Rennes, from Oct. 2014]
Audrey Lucas [CNRS, granted by DGA-PEC, Lannion, from Jan. 2016]
Rafail Psiakis [Univ. Rennes I, MENRT grant, from Oct. 2015]
Simon Rokicki [Univ. Rennes I, granted by ENS Rennes, from Oct. 2015]
Aymen Gammoudi [Univ. Rennes I, Lannion, from Sep. 2015]
Mael Gueguen [Univ. Rennes I, MENRT grant, Rennes, from Nov. 2016]
Genevieve Ndour [CEA Leti, Grenoble, from May 2016]
Joel Ortiz Sosa [Inria, Lannion, from Oct. 2016]
Kleanthis Papachatzopoulos [Inria, Rennes, from Oct. 2016]
Tara Petric [Inria, Rennes, from Nov. 2016]
Nicolas Roux [Inria, Lannion, from Oct. 2016]
Florent Berthier [CEA Leti, Grenoble, until Oct. 2016]
Ali Hassan El-Moussawi [Univ. Rennes I, granted by FP7 Alma, Rennes, until Dec. 2016]
Jérémie Métairie [CNRS, granted by ANR Pavois, Lannion, until Apr. 2016]

Post-Doctoral Fellows

Ashraf El-Antably [Inria, Lannion, from May 2016]
Imen Fassi [Univ. Rennes I, Rennes]
Imran Wali [Univ. Rennes I, Lannion, from Sep. 2016]
Atef Dorai [Univ. Rennes I, ATER, ENSSAT, Lannion, from Sep. 2016]
Karim Bigou [Univ. Rennes I, ATER, IUT, Lannion, until Aug. 2016]
Benoit Lopeze [Univ. Rennes I, ATER, ISTIC, Rennes, until Aug. 2016]

Administrative Assistants

Nadia Derouault [Inria, Rennes]
Angélique Le Pennec [Univ. Rennes I, ENSSAT, Lannion, until Aug. 2016]
Emilie Carquin [Univ. Rennes I, ENSSAT, Lannion, from June 2016]

2. Overall Objectives

2.1. Overall Objectives

Abstract — The CAIRN project-team researches new architectures, algorithms and design methods for flexible, secure, fault-tolerant, and energy-efficient domain-specific system-on-chip (SOC). As performance and energy-efficiency requirements of SOCs, especially in the context of multi-core architectures, are continuously increasing, it becomes difficult for computing architectures to rely only on programmable processors solutions. To address this issue, we promote/advocate the use of reconfigurable hardware, i.e., hardware structures whose organization may change before or even during execution. Such reconfigurable chips offer high performance at a low energy cost, while preserving a high level of flexibility. The group studies these systems from three angles: (i) The invention and design of new reconfigurable architectures with an emphasis on flexible arithmetic operator design, dynamic reconfiguration management and low-power consumption. (ii) The development of their corresponding design flows (compilation and synthesis tools) to enable their automatic design from high-level specifications. (iii) The interaction between algorithms and architectures especially for our main application domains (wireless communications, wireless sensor networks and digital security).

Keywords —**Architectures:** Embedded Systems, System-on-Chip, Reconfigurable Architectures, Hardware Accelerators, Low-Power, Computer Arithmetic, Secure Hardware, Fault Tolerance. **Compilation and synthesis:** High-Level Synthesis, CAD Methods, Numerical Accuracy Analysis, Fixed-Point Arithmetic, Polyhedral Model, Constraint Programming, Source-to-Source Transformations, Domain-Specific Optimizing Compilers, Automatic Parallelization. **Applications:** Wireless (Body) Sensor Networks, High-Rate Optical Communications, Wireless Communications, Applied Cryptography.

The scientific goal of the CAIRN group is to research new hardware architectures for domain-specific SOCs, along with their associated design and compilation flows. We particularly focus on on-chip integration of specialized and reconfigurable accelerators. Reconfigurable architectures, whose hardware structure may be adjusted before or even during execution, originate from the possibilities opened up by Field Programmable Gate Arrays (FPGA) [64] and then by Coarse-Grain Reconfigurable Arrays (CGRA) [67], [81] [1]. Recent evolutions in technology and modern hardware systems confirm that reconfigurable systems are increasingly used in recent and future applications (see e.g. Intel/Altera or Xilinx/Zynq solutions). This architectural model has received a lot of attention in academia over the last two decades [71], and is now considered for industrial use in many application domains. One first reason is that the rapidly changing standards or applications require frequent device modifications. In many cases, software updates are not sufficient to keep devices on the market, while hardware redesigns remain too expensive. Second, the need to adapt the system to changing environments (e.g., wireless channel, harvested energy) is another incentive to use runtime dynamic reconfiguration. Moreover, with technologies at 28 nm and below, manufacturing problems strongly impact electrical parameters of transistors, and transient errors caused by particles or radiations also often appear during execution: error detection and correction mechanisms or autonomic self-control can benefit from reconfiguration capabilities.

As chip density increased, power or energy efficiency has become “the Grail” of all chip architects. With the end of Dennard scaling [76], multicore architectures are hitting the *utilisation wall* and the percentage of transistors in a chip that can switch at full frequency drops at a fast pace [69]. However, this unused portion of a chip also opens up new opportunities for computer architecture innovations. Building specialized processors or hardware accelerators can come with orders-of-magnitude gains in energy efficiency. Since from the beginning of CAIRN in 2009, we advocate the interest of heterogeneous multicores, in which general-purpose processors (GPPs) are integrated with specialized accelerators, especially when built on reconfigurable hardware, which provides the best trade-off between power, performance, cost and flexibility. During the period, it therefore turns out that the time has come for these heterogeneous manycore architectures.

Standard multicore architectures enable flexible software on fixed hardware, whereas reconfigurable architectures make possible **flexible software on flexible hardware**.

However, designing reconfigurable systems poses several challenges: the definition of the architecture structure itself, along with its dynamic reconfiguration capabilities, and its corresponding compilation or synthesis tools. The scientific goal of CAIRN is therefore to leverage the background and past experience of its members to tackle these challenges. We propose to approach energy efficient reconfigurable architectures from three angles: (i) the invention and the design of new reconfigurable architectures or hardware accelerators, (ii) the development of their corresponding compilers and design methods, and (iii) the exploration of the interaction between applications and architectures.

3. Research Program

3.1. Panorama

The development of complex applications is traditionally split in three stages: a theoretical study of the algorithms, an analysis of the target architecture and the implementation. When facing new emerging applications such as high-performance, low-power and low-cost mobile communication systems or smart sensor-based systems, it is mandatory to strengthen the design flow by a joint study of both algorithmic and architectural issues.

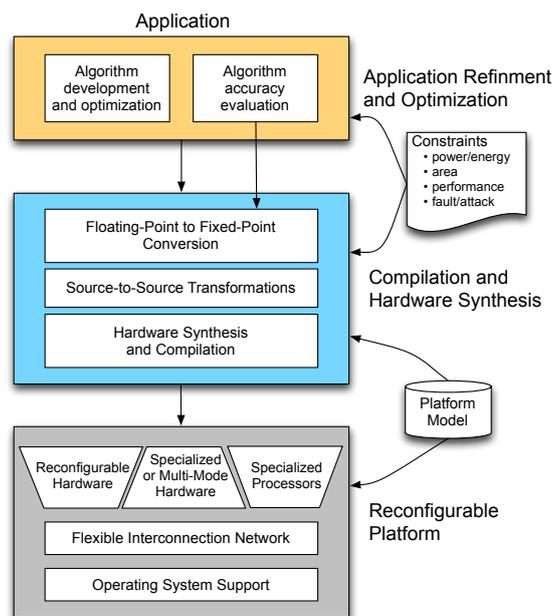


Figure 1. CAIRN's general design flow and related research themes

Figure 1 shows the global design flow we propose to develop. This flow is organized in levels which refer to our three research themes: application optimization (new algorithms, fixed-point arithmetic, advanced representations of numbers), architecture optimization (reconfigurable and specialized hardware, application-specific processors, arithmetic operators and functions), and stepwise refinement and code generation (code transformations, hardware synthesis, compilation).

In the rest of this part, we briefly describe the challenges concerning **new reconfigurable platforms** in Section 3.2 and the issues on **compiler and synthesis tools** related to these platforms in Section 3.3.

3.2. Reconfigurable Architecture Design

Nowadays, FPGAs are not only suited for application specific algorithms, but also considered as fully-featured computing platforms, thanks to their ability to accelerate massively parallelizable algorithms much faster than their processor counterparts [84]. They also support to be dynamically reconfigured. At runtime, partially reconfigurable regions of the logic fabric can be reconfigured to implement a different task, which allows for a better resource usage and adaptation to the environment. Dynamically reconfigurable hardware can also cope with hardware errors by relocating some of its functionalities to another, sane, part of the logic fabric. It could also provide support for a multi-tasked computation flow where hardware tasks are loaded on-demand at runtime. Nevertheless, current design flows of FPGA vendors are still limited by the use of one partial bitstream for each reconfigurable region and for each design. These regions are defined at design time and it is not possible to use only one bitstream for multiple reconfigurable regions nor multiple chips. The multiplicity of such bitstreams leads to a significant increase in memory. Recent research has been conducted in the domain of task relocation on a reconfigurable fabric. All of the related work was conducted on architectures from commercial vendors (e.g., Xilinx, Altera) which share the same limitations: the inner details of the bitstream are not publicly known, which limits applicability of the techniques. To circumvent this issue, most dynamic reconfiguration techniques are either generating multiple bitstreams for each location [66] or implementing an online filter to relocate the tasks [78]. Both of these techniques still suffer from memory footprint and from the online complexity of task relocation.

Increasing the level and grain of reconfiguration is a solution to counterbalance the FPGA penalties. Coarse-grained reconfigurable architectures (CGRA) provide operator-level configurable functional blocks and word-level datapaths [85], [72], [83]. Compared to FPGA, they benefit from a massive reduction in configuration memory and configuration delay, as well as for routing and placement complexity. This in turns results in an improvement in the computation volume over energy cost ratio, although with a loss of flexibility compared to bit-level operations. Such constraints have been taken into account in the design of DART[10], Adres [81] or polymorphous computing fabrics[12]. These works have led to commercial products such as the PACT/XPP [65] or Montium from Recore systems, without however a real commercial success yet. Emerging platforms like Xilinx/Zynq or Intel/Altera are about to change the game.

In the context of emerging heterogenous multicore architecture, CAIRN advocates for associating general-purpose processors (GPP), flexible network-on-chip and coarse-grain or fine-grain dynamically reconfigurable accelerators. We leverage our skills on microarchitecture, reconfigurable computing, arithmetic, and low-power design, to discover and design such architectures with a focus on: -reduced energy per operation, - improved application performance through acceleration, - hardware flexibility and self-adaptive behavior, - tolerance to faults, computing errors, and process variation, - protections against side channel attacks, - limited silicon area overhead.

3.3. Compilation and Synthesis for Reconfigurable Platforms

In spite of their advantages, reconfigurable architectures, and more generally hardware accelerators, lack efficient and standardized compilation and design tools. As of today, this still makes the technology impractical for large-scale industrial use. Generating and optimizing the mapping from high-level specifications to reconfigurable hardware platforms are therefore key research issues, which have received considerable interest over the last years [70], [86], [82], [80], [79]. In the meantime, the complexity (and heterogeneity) of these platforms has also been increasing quite significantly, with complex heterogeneous multi-cores architectures becoming a *de facto* standard. As a consequence, the focus of designers is now geared toward optimizing overall system-level performance and efficiency [77]. Here again, existing tools are not well suited, as they fail at providing an unified programming view of the programmable and/or reconfigurable components implemented on the platform.

In this context, we have been pursuing our efforts to propose tools whose design principles are based on a tight coupling between the compiler and the target hardware architectures. We build on the expertise of the team members in High Level Synthesis (HLS) [6], ASIP optimizing compilers [13] and automatic parallelization for massively parallel specialized circuits [2]. We first study how to increase the efficiency of standard programmable processors by extending their instruction set to speed-up compute intensive kernels. Our focus is on efficient and exact algorithms for the identification, selection and scheduling of such instructions [7]. We address compilation challenges by borrowing techniques from high-level synthesis, optimizing compilers and automatic parallelization, especially when dealing with nested loop kernels. In addition, and independently of the scientific challenges mentioned above, proposing such flows also poses significant software engineering issues. As a consequence, we also study how leading edge software engineering techniques (Model Driven Engineering) can help the Computer Aided Design (CAD) and optimizing compiler communities prototyping new research ideas [14], [5], [3].

Efficient implementation of multimedia and signal processing applications (in software for DSP cores or as special-purpose hardware) often requires, for reasons related to cost, power consumption or silicon area constraints, the use of fixed-point arithmetic, whereas the algorithms are usually specified in floating-point arithmetic. Unfortunately, fixed-point conversion is very challenging and time-consuming, typically demanding up to 50% of the total design or implementation time. Thus, tools are required to automate this conversion. For hardware or software implementation, the aim is to optimize the fixed-point specification. The implementation cost is minimized under a numerical accuracy or an application performance constraint. For DSP-software implementation, methodologies have been proposed [8] to achieve fixed-point conversion. For hardware implementation, the best results are obtained when the word-length optimization process is coupled with the high-level synthesis [73]. Evaluating the effects of finite precision is one of the major and often the most time consuming step while performing fixed-point refinement. Indeed, in the word-length optimization process, the numerical accuracy is evaluated as soon as a new word-length is tested, thus, several times per iteration of the optimization process. Classical approaches are based on fixed-point simulations [74]. Leading to long evaluation times, they can hardly be used to explore the design space. Therefore, our aim is to propose closed-form expressions of errors due to fixed-point approximations that are used by a fast analytical framework for accuracy evaluation [11].

4. Application Domains

4.1. Panorama

keywords: Wireless (Body) Sensor Networks, High-Rate Optical Communications, Wireless Communications, Applied Cryptography.

Our research is based on realistic applications, in order to both discover the main needs created by these applications and to invent realistic and interesting solutions.

Wireless Communication is our privileged application domain. Our research includes the prototyping of (subsets of) such applications on reconfigurable and programmable platforms. For this application domain, the high computational complexity of the 5G Wireless Communication Systems calls for the design of high-performance and energy-efficient architectures. In **Wireless Sensor Networks** (WSN), where each wireless node is expected to operate without battery replacement for significant periods of time, energy consumption is the most important constraint. Sensor networks are a very dynamic domain of research due, on the one hand, to the opportunity to develop innovative applications that are linked to a specific environment, and on the other hand to the challenge of designing totally autonomous communicating objects.

Other important fields are also considered: hardware cryptographic and security modules, high-rate optical communications, machine learning, and multimedia processing.

5. Highlights of the Year

5.1. Highlights of the Year

Our work on accuracy evaluation and optimisation for fixed point arithmetic was presented during a tutorial "Fixed-point refinement, a guaranteed approach towards energy efficient computing" at HiPEAC Conference in January 2016 [60].

Members of CAIRN got six papers accepted at IEEE/ACM Design Automation and Test in Europe for 2017, one of the major events in design automation.

6. New Software and Platforms

6.1. Panorama

With the ever raising complexity of embedded applications and platforms, the need for efficient and customizable compilation flows is stronger than ever. This need of flexibility is even stronger when it comes to research compiler infrastructures that are necessary to gather quantitative evidence of the performance/energy or cost benefits obtained through the use of reconfigurable platforms. From a compiler point of view, the challenges exposed by these complex reconfigurable platforms are quite significant, since they require the compiler to extract and to expose an important amount of coarse and/or fine grain parallelism, to take complex resource constraints into consideration while providing efficient memory hierarchy and power management.

Because they are geared toward industrial use, production compiler infrastructures do not offer the level of flexibility and productivity that is required for compiler and CAD tool prototyping. To address this issue, we designed an extensible source-to-source compiler infrastructure that takes advantage of leading edge model-driven object-oriented software engineering principles and technologies.

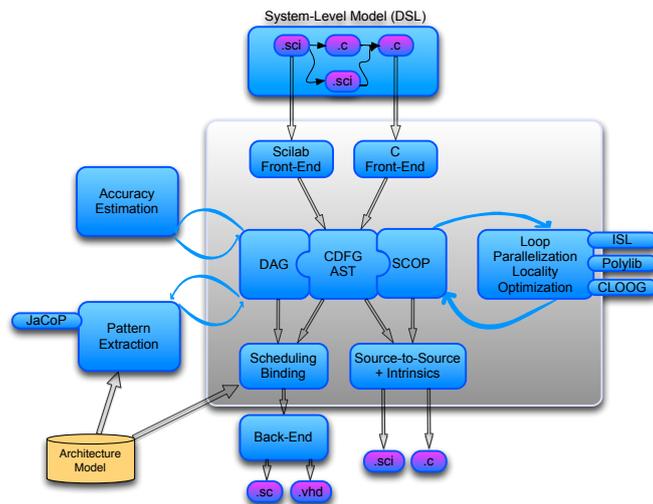


Figure 2. CAIRN's general software development framework.

Figure 2 shows the global framework that is being developed in the group. Our compiler flow mixes several types of intermediate representations. The baseline representation is a simple tree-based model enriched with control flow information. This model is mainly used to support our source-to-source flow, and serves as the backbone for the infrastructure. We use the extensibility of the framework to provide more advanced representations along with their corresponding optimizations and code generation plug-ins. For example, for our pattern selection and accuracy estimation tools, we use a data dependence graph model in all basic blocks instead of the tree model. Similarly, to enable polyhedral based program transformations and analysis, we introduced a specific representation for affine control loops that we use to derive a Polyhedral Reduced Dependence Graph (PRDG). Our current flow assumes that the application is specified as a hierarchy of communicating tasks, where each task is expressed using C or Matlab/Scilab, and where the system-level representation and the target platform model are often defined using Domain Specific Languages (DSL).

Gecos (Generic Compiler Suite) is the main backbone of CAIRN's flow. It is an open source Eclipse-based flexible compiler infrastructure developed for fast prototyping of complex compiler passes. Gecos is a 100% Java based implementation and is based on modern software engineering practices such as Eclipse plugin or model-driven software engineering with EMF (Eclipse Modeling Framework). As of today, our flow offers the following features:

- An automatic floating-point to fixed-point conversion flow (for ASIC/FPGA and embedded processors). **ID.Fix** is an infrastructure for the automatic transformation of software code aiming at the conversion of floating-point data types into a fixed-point representation.
- A polyhedral-based loop transformation and parallelization engine (mostly targeted at HLS).
- A custom instruction extraction flow (for ASIP and dynamically reconfigurable architectures). **Durase** is developed for the compilation and the synthesis targeting reconfigurable platforms and the automatic synthesis of application specific processor extensions. It uses advanced technologies, such as graph matching together with constraint programming methods.
- Several back-ends to enable the generation of VHDL for specialized or reconfigurable IPs, and SystemC for simulation purposes (e.g., fixed-point simulations).

Gecos, ID.Fix or Durase have been demonstrated during "University Booths" in various conference such as IEEE/ACM DAC or DATE.

6.2. Gecos

KEYWORDS: Source-to-source compiler - Model-driven software engineering - Retargetable compilation
SCIENTIFIC DESCRIPTION

The Gecos (Generic Compiler Suite) project is a source-to-source compiler infrastructure developed in the Cairn group since 2004. It was designed to enable fast prototyping of program analysis and transformation for hardware synthesis and retargetable compilation domains.

Gecos is 100% Java based and takes advantage of modern model driven software engineering practices. It uses the Eclipse Modeling Framework (EMF) as an underlying infrastructure and takes benefits of its features to make it easily extensible. Gecos is open-source and is hosted on the Inria gforge at <http://gecos.gforge.inria.fr>.

The Gecos infrastructure is still under very active development, and serves as a backbone infrastructure to projects of the group. Part of the framework is jointly developed with Colorado State University and between 2012 and 2015 it was used in the context of the FP7 ALMA European project. The Gecos infrastructure will also be used by the EMMATRIX start-up, a spin-off from the ALMA project which aims at commercializing the results of the project and in the context of the H2020 ARGO European project.

FUNCTIONAL DESCRIPTION

Gecos provides a program transformation toolbox facilitating parallelisation of applications for heterogeneous multiprocessor embedded platforms. This includes a polyhedral loop transformation toolbox, efficient SIMD code generation for fixed point arithmetic data-types, coarse-grain parallelization engine targeting the data-flow actor model, and a Matlab/Scilab front-end. In addition to targeting programmable processors, Gecos can regenerate optimized code for High Level Synthesis tools.

- Participants: Steven Derrien, Nicolas Simon, Imen Fassi, and Ali Hassan El-Moussawi
- Partner: Université de Rennes 1
- Contact: Steven Derrien
- URL: <http://gecos.gforge.inria.fr/doku/doku.php>

6.3. ID-Fix

KEYWORDS: Energy efficiency - Embedded systems - Analytical accuracy evaluation - Fixed-point arithmetic - Accuracy optimization - Dynamic range evaluation - Code optimisation

SCIENTIFIC DESCRIPTION

The different techniques proposed by the team for fixed-point conversion are implemented on the ID.Fix infrastructure. The application is described with a C code using floating-point data types and different pragmas, used to specify parameters (dynamic, input/output word-length, delay operations) for the fixed-point conversion. This tool determines and optimizes the fixed-point specification and then, generates a C code using different fixed-point data types. The infrastructure is made-up of two main modules corresponding to the fixed-point conversion (ID.Fix-Conv) and the accuracy evaluation (ID.Fix-Eval).

FUNCTIONAL DESCRIPTION

ID.Fix focuses on computational accuracy and can provide an optimised specification using fixed point arithmetic from a C source code with floating point data types. Fixed point arithmetic is very widely used in embedded systems as it provides better performance and is much more energy efficient. ID.Fix used an analytical model of the software code, which means it can explore more solutions and thereby produce much more efficient code than classical simulation-based tools.

- Participants: Olivier Sentieys, Benjamin Barrois and Nicolas Simon
- Partner: Université de Rennes 1
- Contact: Olivier Sentieys
- URL: <http://idfix.gforge.inria.fr/doku.php>

6.4. Zyggye

KEYWORDS: Health - Biomechanics - Wireless body sensor networks - Low power - Gesture recognition - Hardware platform - Software platform - Localization

SCIENTIFIC DESCRIPTION

Zyggye is a hardware and software wireless body sensor network platform. Each sensor node, attached to different parts of the human body, contains inertial sensors (IMU) (accelerometer, gyrometer, compass and barometer), an embedded processor and a low-power radio module to communicate data to a coordinator node connected to a computer, tablet or smartphone. One of the system's key innovations is that it collects data from sensors as well as on distances estimated from the power of the radio signal received to make the 3D location of the nodes more precise and thus prevent IMU sensor drift and power consumption overhead. Zyggye can be used to determine posture or gestures and mainly has applications in sport, healthcare and the multimedia industry.

FUNCTIONAL DESCRIPTION

The Zyggy sensor platform was developed to create an autonomous Wireless Body Sensor Network (WBSN) with the capabilities of monitoring body movements. The Zyggy platform is part of the BoWI project funded by CominLabs. Zyggy is composed of a processor, a radio transceiver and different sensors including an Inertial Measurement Unit (IMU) with 3-axis accelerometer, gyrometer, and magnetometer. Zyggy is used for evaluating data fusion algorithms, low power computing algorithms, wireless protocols, and body channel characterization in the BoWI project.

The Zyggy V2 prototype includes new features: a 32-bit microcontroller to manage a custom MAC layer and process quaternions based on IMU measures, and an UWB radio from DecaWave to measure distances between nodes with Time of Flight (ToF).

- Participants: Arnaud Carer and Olivier Sentieys
- Partners: Lab-STICC - Université de Rennes 1
- Contact: Olivier Sentieys
- URL: <http://www.bowi.cominlabs.ueb.eu/fr/zyggie-wbsn-platform>



Figure 3. CAIRN's Zyggy platform for WBSN

7. New Results

7.1. Reconfigurable Architecture Design

7.1.1. Dynamic Reconfiguration Support in FPGA

Participants: Olivier Sentieys, Christophe Hurliaux.

Almost since the creation of the first SRAM-based FPGAs there has been a desire to explore the benefits of partially reconfiguring a portion of an FPGA at run-time while the remainder of design functionality continues to operate uninterrupted. Currently, the use of partial reconfiguration imposes significant limitations on the FPGA design: reconfiguration regions must be constrained to certain shapes and sizes and, in many cases, bitstreams must be precompiled before application execution depending on the precise region of the placement in the fabric. We developed an FPGA architecture that allows for seamless translation of partially-reconfigurable regions, even if the relative placement of fixed-function blocks within the region is changed.

In [4], we proposed a design flow for generating compressed configuration bitstreams abstracted from their final position on the logic fabric, the Virtual Bit-Streams (VBS). Those configurations can then be decoded and finalized in real-time and at run-time by a dedicated reconfiguration controller to be placed at a given physical location. The VPR (Versatile Place and Route) framework was expanded to include bitstream generation features. The configuration stream format was proposed along with its associated decoding architecture. We analyzed the compression induced by our coding method and proved that compression ratios of at least $2.5\times$ can be achieved on the 20 largest MCNC benchmarks. The introduction of clustering which aggregates multiple routing resources together showed compression ratio up to a factor of $10\times$, at the cost of a more complex decoding step at runtime.

The emergence of 2.5D and 3D packaging technologies enables the integration of FPGA dice into more complex systems. Both heterogeneous manycore designs, which include an FPGA layer, and interposer-based multi-FPGA systems support the inclusion of reconfigurable hardware in 3D-stacked integrated circuits. In these architectures, the communication between FPGA dice or between FPGA and fixed-function layers often takes place through dedicated communication interfaces spread over the FPGA logic fabric, as opposed to an I/O ring around the fabric. In [39], we investigate the effect of organizing FPGA fabric I/O into coarse-grained interface blocks distributed throughout the FPGA fabric. Specifically, we consider the quality of results for the placement and routing phases of the FPGA physical design flow. We evaluate the routing of I/O signals of large applications through dedicated interface blocks at various granularities in the logic fabric, and study its implications on the critical path delay of routed designs. We show that the impact of such I/O routing is limited and can improve chip routability and circuit delay in many cases.

7.1.2. *Hardware Accelerated Simulation of Heterogeneous Platforms*

Participant: François Charot.

When considering designing heterogeneous multi-core platforms, the number of possible design combinations leads to a huge design space, with subtle trade-offs and design interactions. To reason about what design is best for a given target application requires detailed simulation of many different possible solutions. Simulation frameworks exist (such as gem5) and are commonly used to carry out these simulations. Unfortunately, these are purely software-based approaches and they do not allow a real exploration of the design space. Moreover, they do not really support highly heterogeneous multi-core architectures. These limitations motivate the study of the use of hardware to accelerate the simulation, and in particular of FPGA components. In this context, we are currently investigating the possibility of building hardware accelerated simulators using the HASim simulation infrastructure, jointly developed by MIT and Intel. HASim is a FPGA-accelerated simulator that is able to simulate a multicore with a high-detailed pipeline, cache hierarchy and detailed on-chip network on a single FPGA. We work on integrating a model of the RISC-V instruction set architecture in the HASim infrastructure. This work is done with the perspective of studying hardware accelerated simulation of heterogeneous multicore architectures mixing RISC-V cores and hardware accelerators.

7.1.3. *Optical Interconnections for 3D Multiprocessor Architectures*

Participants: Jiating Luo, Ashraf El-Antably, Pham Van Dung, Cédric Killian, Daniel Chillet, Olivier Sentieys.

To address the issue of interconnection bottleneck in multiprocessor on a single chip, we study how an Optical Network-on-Chip (ONoC) can leverage 3D technology by stacking a specific photonics die. The objectives of this study target: i) the definition of a generic architecture including both electrical and optical components, ii) the interface between electrical and optical domains, iii) the definition of strategies (communication protocol) to manage this communication medium, and iv) new techniques to manage and reduce the power consumption of optical communications. The first point is required to ensure that electrical and optical components can be used together to define a global architecture. Indeed, optical components are generally larger than electrical components, so a trade-off must be found between the size of optical and electrical parts. For example, if the need in terms of communications is high, several waveguides and wavelengths must be necessary, and can lead to an optical area larger than the footprint of a single processor. In this case, a solution is to connect (through the optical NoC) clusters of processors rather than each single processor. For the second point, we study how the interface can be designed to take applications needs into account. From the different possible interface designs, we extract a high-level performance model of optical communications from losses induced by all optical components to efficiently manage Laser parameters. Then, the third point concerns the definition of high-level mechanisms which can handle the allocation of the communication medium for each data transfer between tasks. This part consists in defining the protocol of wavelength allocation. Indeed, the optical wavelengths are a shared resource between all the electrical computing clusters and are allocated at run time according to application needs and quality of service. The last point concerns the definition of techniques allowing to reduce the power consumption of on-chip optical communications. The power of each Laser can be dynamically tuned in the optical/electrical interface at run time for a given targeted bit-error-rate. Due to the relatively high power

consumption of such integrated Laser, we study how to define adequate policies able to adapt the laser power to the signal losses.

We are currently designing an Optical-Network-Interface (ONI) to connect one processor, or a cluster of several processors, to the optical communication medium. This interface, constrained by the 10 Gb/s data-rate of the Lasers, integrates Error Correcting Codes and a communication manager. This manager can select, at run-time, the communication mode to use depending on timing or power constraints. Indeed, as the use of ECC is based on redundant bits, it increases the transmission time, but saves power for a given Bit Error Rate (BER). Moreover, our ONI allows for data to be sent using several wavelengths in parallel, hence increasing transmission bandwidth.

However, multiple signals sharing simultaneously a waveguide can lead to inter-channel crosstalk noise. This problem impacts the Signal to Noise Ratio (SNR) of the optical signal, which leads to an increase in the Bit Error Rate (BER) at the receiver side. In [40], [59], we proposed a Wavelength Allocation (WA) method allowing to search for performance and energy trade-offs based on application constraints. We showed that for a 16-core WDM ring-based ONoC architecture using 12 wavelengths, more than 100,000 allocation solutions exist and only 51 are on a Pareto front giving a tradeoff between execution time and energy per bit (derived from the BER). The optimized solutions reached reduce the execution time by 37% or the energy from 7,6fJ/bit to 4,4fJ/bit.

7.1.4. Communication-Based Power Modelling for Heterogeneous Multiprocessor Architectures

Participants: Baptiste Roux, Olivier Sentieys, Steven Derrien.

Programming heterogeneous multiprocessor architectures is a real challenge dealing with a huge design space. Computer-aided design and development tools try to circumvent this issue by simplifying instantiation mechanisms. However, energy consumption is not well supported in most of these tools due to the difficulty to obtain fast and accurate power estimation. To this aim, in [46] we proposed and validated a power model for such platforms. The methodology is based on micro-benchmarking to estimate the model parameters. The energy model mainly relies on the energy overheads induced by communications between processors in a parallel application. Power modelling and micro-benchmarks are validated using a Zynq-based heterogeneous architecture showing the accuracy of the model for several tested synthetic applications.

7.1.5. Arithmetic Operators for Cryptography and Fault-Tolerance

Participants: Arnaud Tisserand, Emmanuel Casseau, Pierre Guilloux, Karim Bigou, Gabriel Gallin, Audrey Lucas, Franck Bucheron, Jérémie Métairie.

Arithmetic Operators for Fast and Secure Cryptography.

Our paper [21], published in IEEE Transactions on Computers, extends our fast RNS modular inversion for finite fields arithmetic published at CHES 2013 conference. It is based on the binary version of the plus-minus Euclidean algorithm. In the context of elliptic curve cryptography (*i.e.* 160–550 bits finite fields), it significantly speeds-up modular inversions. In this extension, we propose an improved version based on both radix 2 and radix 3. This new algorithm leads to 30 % speed-up for a maximal area overhead about 4 % on Virtex 5 FPGAs. This work was done in the ANR PAVOIS project.

Our paper [32], presented at ARITH-23, presents an hybrid representation of large integers, or prime field elements, combining both positional and residue number systems (RNS). Our *hybrid position-residues* (HPR) number system mixes a high-radix positional representation and digits represented in RNS. RNS offers an important source of parallelism for addition, subtraction and multiplication operations. But, due to its non-positional property, it makes comparisons and modular reductions more costly than in a positional number system. HPR offers various trade-offs between internal parallelism and the efficiency of operations requiring position information. Our current application domain is asymmetric cryptography where HPR significantly reduces the cost of some modular operations compared to state-of-the-art RNS solutions. This work was done in the ANR PAVOIS project.

An ASIC circuit has been implemented in the 65nm ST CMOS technology and sent to fabrication in June 2016 (chip delivery is expected for January 2017). The implemented cryptoprocessor was designed for 256-bit prime finite fields elements and generic curves. It embeds: 1 multiplier, 1 adder and 1 inversion units for field-level computations. Various algorithms for scalar multiplication primitives can be programmed in software for curve-level computations. It was designed to evaluate algorithmic and arithmetic protections against side channel attacks (there is no hardware protection embedded in this ASIC version). This work was done in the ANR PAVOIS project.

In the HAH project, funded by CominLabs and Lebesgue Labex, we study hardware implementation of cryptoprocessors for hyperelliptic curves. The poster [61] presents the current state of the project for FPGA implementations.

Arithmetic Operators for Fault-Tolerance.

Various methods have been proposed for fault detection and fault tolerance in digital integrated circuits. In the case of *arithmetic circuits*, the selection of an efficient method depends on several elements: type of operation, type(s) of operand(s), computation algorithms, internal representations of numbers, optimizations at architecture and circuit levels, and acceptable accuracy level (i.e. mathematical error) of the result(s) including both rounding errors and errors due to the faults. High-level mathematical models are not sufficient to capture the effect of faults in arithmetic circuits. Simulation of intensive fault scenarios in all components of the arithmetic circuit (data-path, control, gates with important fan-out such as some partial products generation in large multipliers, etc.) is widely used. But cycle accurate and bit accurate software simulations at gate level are too slow for large circuits and numerous fault scenarios. *FPGA emulation* is a popular method to speed-up fault simulation.

We are developing an hardware-software platform dedicated to fault emulation for ASIC arithmetic circuits. The platform is based on a parallel cluster of Zynq FPGA cards and a Linux server. Various arithmetic circuits and fault models will be demonstrated in the context of digital signal and image processing. Our paper [57], presented at Compas, describes the very first version of our platform. This platform has also been presented in a poster at GDR SoC-SiP [58] and in a Demo Night at DASIP [56]. This work was done in the ANR ARDyT and Reliasic projects.

7.1.6. Adaptive Overclocking, Error Correction, and Voltage Over-Scaling for Error-Resilient Applications

Participants: Rengarajan Ragavan, Benjamin Barrois, Cédric Killian, Olivier Sentieys.

Error detection and correction based on double-sampling is used as common technique to handle timing errors while scaling V_{dd} for energy efficiency. Implementation and advantages of double-sampling technique in FPGAs are simpler and significant compared to the conventional highly pipelined processors due to the higher flexibility of the reconfigurable architectures. It is common practice to insert shadow flipflop in the critical paths of the design, which will fail while scaling down the supply voltage, or to correct timing errors while over clocking the datapaths. Overclocking, and error detection and correction capabilities of these methods are limited due to the fixed speculation window used by these methods. In [44], we presented a Dynamic Speculation Window in double-sampling for timing error detection and correction in FPGAs. The proposed method employs online slack measurement and conventional shadow flipflop approach to adaptively overclock the design and also to detect and correct timing errors due to temperature and other variability effects. We demonstrated this method in the Xilinx VC707 Virtex 7 FPGA for various benchmarks. We achieved maximum of 71% overclocking for unsigned 32-bit multiplier with the area overhead of 1.9% LUTs and 1.7% FFs.

Voltage scaling has been used as a prominent technique to improve energy efficiency in digital systems, scaling down supply voltage effects in quadratic reduction in energy consumption of the system. Reducing supply voltage induces timing errors in the system that are corrected through additional error detection and correction circuits. In [43], we proposed voltage over-scaling based approximate operators for applications that can tolerate errors. We characterized the basic arithmetic operators using different operating triads (combination of supply voltage, body-biasing scheme and clock frequency) to generate models for approximate operators. Error-resilient applications can be mapped with the generated approximate operator models to achieve

optimum trade-off between energy efficiency and error margin. Based on the dynamic speculation technique, best possible operating triad is chosen at runtime based on the user definable error tolerance margin of the application. In our experiments in 28nm FDSOI, we achieved maximum energy efficiency of 89% for basic operators like 8-bit and 16-bit adders at the cost of 20% Bit Error Rate (ratio of faulty bits over total bits) by operating them in near-threshold regime.

7.2. Compilation and Synthesis for Reconfigurable Platform

7.2.1. Adaptive dynamic compilation for low power embedded systems

Participants: Steven Derrien, Simon Rokicki.

Dynamic binary translation (DBT) consists in translating – at runtime – a program written for a given instruction set to another instruction set. Dynamic Translation was initially proposed as a means to enable code portability between different instruction sets and can be implemented in software or hardware. DBT is also used to improve the energy efficiency of high performance processors, as an alternative to out-of-order microarchitectures. In this context, DBT is used to uncover instruction level parallelism (ILP) in the binary program, and then target an energy efficient wide issue VLIW architecture. This approach is used in Transmeta Crusoe [75] and NVidia Denver [68] processors. Since DBT operates at runtime, its execution time is directly perceptible by the user, hence severely constrained. As a matter of fact, this overhead has often been reported to have a huge impact on actual performance, and is considered as being the main weakness of DBT based solutions. This is particularly true when targeting a VLIW processor: the quality of the generated code depends on efficient scheduling; unfortunately scheduling is known to be the most time-consuming component of a JIT compiler or DBT. Improving the responsiveness of such DBT systems is therefore a key research challenge. This is however made very difficult by the lack of open research tools or platform to experiment with such platforms. In this work, we have been addressing these two issues by developing an open hardware/software platform supporting DBT. The platform was designed using HLS tools and validated on a FPGA board. The DBT uses RISC-V as host ISA, and can target varying issue width VLIW architectures. Our platform uses custom hardware accelerators to improve the reactivity of our optimizing DBT flow. Our results show that, compared to a software implementation, our approach offers speed-up by $8\times$ while consuming $18\times$ less energy.

7.2.2. Leveraging Power Spectral Density for Scalable System-Level Accuracy Evaluation

Participants: Benjamin Barrois, Olivier Sentieys.

The choice of fixed-point word-lengths critically impacts the system performance by impacting the quality of computation, its energy, speed and area. Making a good choice of fixed-point word-length generally requires solving an NP-hard problem by exploring a vast search space. Therefore, the entire fixed-point refinement process becomes critically dependent on evaluating the effects of accuracy degradation. In [30], a novel technique for the system-level evaluation of fixed-point systems, which is more scalable and that renders better accuracy, was proposed. This technique makes use of the information hidden in the power-spectral density of quantization noises. It is shown to be very effective in systems consisting of more than one frequency sensitive components. Compared to state-of-the-art hierarchical methods that are agnostic to the quantization noise spectrum, we show that the proposed approach is $5\times$ to $500\times$ more accurate on some representative signal processing kernels.

7.2.3. Approximate Computing

Participants: Benjamin Barrois, Olivier Sentieys.

Many applications are error-resilient, allowing for the introduction of approximations in the calculations, as long as a certain accuracy target is met. Traditionally, fixed-point arithmetic is used to relax accuracy, by optimizing the bit-width. This arithmetic leads to important benefits in terms of delay, power and area. Lately, several hardware approximate operators were invented, seeking the same performance benefits. However, a fair comparison between the usage of this new class of operators and classical fixed-point arithmetic with careful truncation or rounding, has never been performed. In [31], we first compare approximate and fixed-point arithmetic operators in terms of power, area and delay, as well as in terms of induced error, using many state-of-the-art metrics and by emphasizing the issue of data sizing. To perform this analysis, we developed a design exploration framework, APXPERF, which guarantees that all operators are compared using the same operating conditions. Moreover, operators are compared in several classical real-life applications leveraging relevant metrics. In [31], we show that considering a large set of parameters, existing approximate adders and multipliers tend to be dominated by truncated or rounded fixed-point ones. For a given accuracy level and when considering the whole computation data-path, fixed-point operators are several orders of magnitude more accurate while spending less energy to execute the application. A conclusion of this study is that the entropy of careful sizing is always lower than approximate operators, since it requires significantly less bits to be processed in the data-path and stored. Approximated data therefore always contain on average a greater amount of costly erroneous, useless information.

7.2.4. Real-Time Scheduling of Reconfigurable Battery-Powered Multi-Core Platforms

Participants: Daniel Chillet, Aymen Gammoudi.

Reconfigurable real-time embedded systems are constantly increasingly used in applications like autonomous robots or sensor networks. Since they are powered by batteries, these systems have to be energy-aware, to adapt to their environment and to satisfy real-time constraints. For energy harvesting systems, regular recharges of battery can be estimated, and by including this parameter in the operating system, it is then possible to develop a strategy able to ensure the best execution of the application until the next recharge. In this context, operating system services must control the execution of tasks to meet the application constraints. Our objective concerns the proposition of a new real-time scheduling strategy that considers execution constraints such as the deadline of tasks and the energy.

To address this issue, we first focus on mono-processor scheduling [38] and propose to classify the tasks that have similar periods (or WCETs) in packs and to manage the execution parameters of these packs. For each reconfiguration scenario, parameter modifications are performed on packs/tasks to meet the real-time and energy constraints. Compared to previous work, task delaying is significantly improved in [36]. Furthermore, we also develop a strategy for multi-cores systems considering the dependencies between tasks [37] by adding the cost of communication between cores.

7.2.5. Optimization of loop kernels using software and memory information

Participant: Angeliki Kritikakou.

Current compilers cannot generate code that can compete with hand-tuned code in efficiency, even for a simple kernel like matrix–matrix multiplication (MMM). A key step in program optimization is the estimation of optimal values for parameters such as tile sizes and number of levels of tiling. The scheduling parameter values selection is a very difficult and time-consuming task, since parameter values depend on each other; this is why they are found by using searching methods and empirical techniques. To overcome this problem, the scheduling sub-problems must be optimized together, as one problem and not separately. In [24], an MMM methodology is presented where the optimum scheduling parameters are found by decreasing the search space theoretically, while the major scheduling sub-problems are addressed together as one problem and not separately according to the hardware architecture parameters and input size; for different hardware architecture parameters and/or input sizes, a different implementation is produced. This is achieved by fully exploiting the software characteristics (e.g., data reuse) and hardware architecture parameters (e.g., data caches sizes and associativities), giving high-quality solutions and a smaller search space. This methodology refers to a wide range of CPU and GPU architectures.

The size required to store an array is crucial for an embedded system, as it affects the memory size, the energy per memory access and the overall system cost. Existing techniques for finding the minimum number of resources required to store an array are less efficient for codes with large loops and not regularly occurring memory accesses. They have to approximate the accessed parts of the array leading to overestimation of the required resources. Otherwise their exploration time is increased with an increase over the number of the different accessed parts of the array. In [25], we propose a methodology to compute the minimum resources required for storing an array which keeps the exploration time low and provides a near-optimal result for regularly and non-regularly occurring memory accesses and overlapping writes and reads.

7.2.6. Adaptive Software Control to Increase Resource Utilization in Mixed-Critical Systems

Participant: Angeliki Kritikakou.

Automotive embedded systems need to cope with antagonist requirements: on the one hand, the users and market pressure push car manufacturers to integrate more and more services that go far beyond the control of the car itself. On the other hand, recent standardization efforts in the safety domain has led to the development of the ISO 26262 norm that defines means and requirements to ensure the safe operation of automotive embedded systems. In particular, it led to the definition of ASIL (Automotive Safety and Integrity Levels), i.e., it formally defines several criticality levels. Handling the increased complexity of new services makes new architectures, such as multi or many-cores, appealing choices for the car industry. Yet, these architectures provide a very low level of timing predictability due to shared resources, which goes in contradiction with timing guarantees required by ISO 26262. For highest criticality level tasks, Worst-Case Execution Time analysis (WCET) is required to guarantee that timing constraints are respected. The WCET analyzers consider the worst-case scenario: whenever a critical task accesses a shared resource in a multi/many-core platform, a WCET analyzer considers that all cores use the same resource concurrently. To improve the system performance, we proposed in a earlier work an approach where a critical task can be run in parallel with less critical tasks, as long as the real-time constraints are met. When no further interferences can be tolerated, the proposed run-time control in [54] suspends the low critical tasks until the termination of the critical task. In an automotive context, the approach can be translated as a highly critical partition, namely a classic AUTOSAR one, that runs on one dedicated core, with several cores running less critical Adaptive AUTOSAR application(s). We briefly describe in [54] the design of our proven-correct approach. Our strategy is based on a graph grammar to formally model the critical task as a set of control flow graphs on which a safe partial WCET analysis is applied and used at run-time to control the safe execution of the critical task.

8. Partnerships and Cooperations

8.1. Regional Initiatives

8.1.1. Images & Réseaux Competitvity Cluster - Embrace (2014-2016)

Participants: Raphaël Bardoux, Arnaud Carer, Olivier Sentieys.

Embrace (Embedded Radio Accelerator) is a project which involves CAIRN and two Small Medium Enterprises (SMEs): Digidia and PrimeGPS. Embrace aims at developing a software radio platform to enable the digital demodulation of HF signals. Both SMEs will use this platform as the first step to implement new products. These products will be dedicated to two different applications (Global Navigation Satellite System and Navigation Safety) at the heart of the markets of the SMEs. CAIRN goal is the technological transfer of the methods proposed by the team that enable the rapid prototyping of digital radios.

8.2. National Initiatives

8.2.1. ANR Blanc - PAVOIS (2012–2016)

Participants: Arnaud Tisserand, Emmanuel Casseau, Jérémie Métairie, Karim Bigou, Pierre Guilloux.

PAVOIS is a project on Arithmetic Protections Against Physical Attacks for Elliptic Curve based Cryptography that will provide novel implementations of curve based cryptographic algorithms on custom hardware platforms. A specific focus is placed on trade-offs between efficiency and robustness against physical attacks. It involves IRISA-CAIRN (Lannion) and LIRMM (Perpignan and Montpellier). Theoretical aspects include an investigation of how special number representations can be used to speed-up cryptographic algorithms, and protect cryptographic devices from physical attacks. On the practical side, we design innovative cryptographic hardware architectures of a specific processor based on the theoretical advancements described above to implement curve based protocols. For more details see <http://pavois.irisa.fr>.

8.2.2. ANR Ingénierie Numérique et Sécurité - ARDyT (2011-2016)

Participants: Arnaud Tisserand, Pierre Guilloux.

ARDyT is a project on a Reliable and Reconfigurable Dynamic Architecture. It involves IRISA-CAIRN (Lannion), Lab-STICC (Lorient), LIEN (Nancy) and ATMEL. The purpose of the ARDyT project is to provide a complete environment for the design of a fault tolerant and self-adaptable platform. Then, a platform architecture, its programming environment and management methodologies for diagnosis, testability and reliability have to be defined and implemented. The considered techniques are exempt from the use of hardened components for terrestrial and aeronautics applications for the design of low-cost solutions. For more details see <http://ardyt.irisa.fr>.

8.2.3. Labex CominLabs - BoWI (2012-2016)

Participants: Olivier Sentieys, Arnaud Carer.

The BoWi project (Body World Interactions) project aims at designing an accurate gesture and body movement estimation using very-small and low-power wearable sensor nodes, to propose pioneer interfaces for an emerging interacting world based on smart environments (house, media, information and entertainment systems...). Relying on Wireless Body Areas Sensor Networks, we propose an accurate Gesture and Body Movement estimation with extremely severe constraints in terms of footprint and energy consumption. The BoWi geolocation approach will combine radio communication distance measurement and inertial sensors and will also strongly benefit from cooperative techniques based on multiple observations and distributed computation. Different types of applications, such as health care, activity monitoring and environment control, are considered and prototyped. BoWI involves CAIRN, IRISA Granit (Lannion), IETR (Rennes), and Lab-STICC (Brest, Lorient, Vannes). For more details see <http://www.bowi.cominlabs.ueb.eu>.

8.2.4. Labex CominLabs - 3DCORE (2014-2018)

Participants: Olivier Sentieys, Daniel Chillet, Cédric Killian, Jiating Luo, Van Dung Pham, Ashraf El-Antably.

3DCORE (3D Many-Core Architectures based on Optical Network on Chip) is a project investigating new solutions based on silicon photonics to enhance by 2 to 3 magnitude orders energy efficiency and data rate of on-chip interconnect in the context of a many-core architecture. Moreover, 3DCore will take advantage of 3D technologies to design a specific optical layer suitable for a flexible and energy efficient high-speed optical network on chip (ONoC). 3DCORE involves CAIRN, FOTON (Rennes, Lannion) and Institut des Nanotechnologies de Lyon. For more details see <http://www.3d-opt-many-cores.cominlabs.ueb.eu>.

8.2.5. Labex CominLabs - RELIASIC (2014-2018)

Participants: Emmanuel Casseau, Arnaud Tisserand.

RELIASIC (Reliable Asic) will address the issue of fault-tolerant computation with a bottom-up approach, starting from an existing application as a use case (a GPS receiver) and adding some redundant mechanisms to allow the GPS receiver to be tolerant to transient errors due to low voltage supply. RELIASIC involves CAIRN, Lab-STICC (Lorient) and IETR (Rennes). For more details see <http://www.reliasic.cominlabs.ueb.eu> In this project, CAIRN is in charge of the analysis and design of arithmetic operators for fault tolerance. We focus on the hardware implementations of conventional arithmetic operators such as adders, multipliers and MACs but also higher level operators like butterfly computation operator for FFT algorithm.

8.2.6. *Labex CominLabs & Lebesgue - H-A-H (2014-2017)*

Participants: Arnaud Tisserand, Karim Bigou, Gabriel Gallin, Audrey Lucas.

H-A-H for *Hardware and Arithmetic for Hyperelliptic Curves Cryptography* is a project on advanced arithmetic representation and algorithms for hyper-elliptic curve cryptography. It will provide novel implementations of HECC based cryptographic algorithms on custom hardware platforms. H-A-H involves CAIRN (Lannion) and IRMAR (Rennes). For more details see <http://h-a-h.inria.fr/>.

8.3. European Initiatives

8.3.1. *H2020 ARGO*

Participants: Steven Derrien, Olivier Sentieys, Imen Fassi, Ali Hassan El-Moussawi.

Program: H2020-ICT-04-2015

Project acronym: ARGO

Project title: WCET-Aware Parallelization of Model-Based Applications for Heterogeneous Parallel Systems

Duration: Feb. 2016 - Feb. 2019

Coordinator: KIT

Other partners: KIT (DE), UR1/Inria/CAIRN (FR), Recore Systems (NL), TEI-WG (GR), Scilab Ent. (FR), Absint (DE), DLR (DE), Fraunhofer (DE)

Increasing performance and reducing cost, while maintaining safety levels and programmability are the key demands for embedded and cyber-physical systems, e.g. aerospace, automation, and automotive. For many applications, the necessary performance with low energy consumption can only be provided by customized computing platforms based on heterogeneous many-core architectures. However, their parallel programming with time-critical embedded applications suffers from a complex toolchain and programming process. ARGO will address this challenge with a holistic approach for programming heterogeneous multi- and many-core architectures using automatic parallelization of model-based real-time applications. ARGO will enhance WCET-aware automatic parallelization by a cross-layer programming approach combining automatic tool-based and user-guided parallelization to reduce the need for expertise in programming parallel heterogeneous architectures. The ARGO approach will be assessed and demonstrated by prototyping comprehensive time-critical applications from both aerospace and industrial automation domains on customized heterogeneous many-core platforms.

8.3.2. *ANR International ARTEFaCT*

Participants: Olivier Sentieys, Benjamin Barrois, Tara Petric, Tomofumi Yuki.

Program: ANR International France-Switzerland

Project acronym: ARTEFaCT

Project title: AppRoximaTivE Flexible Circuits and Computing for IoT

Duration: Feb. 2016 - Dec. 2019

Coordinator: CEA

Other partners: CEA-LETI (FR), CAIRN (FR), EPFL (SW)

The ARTEFaCT project aims to build on the preliminary results on inexact and exact near-threshold and sub-threshold circuit design to achieve major energy consumption reductions by enabling adaptive accuracy control of applications. ARTEFaCT proposes to address, in a consistent fashion, the entire design stack, from physical hardware design, up to software application analysis, compiler optimizations, and dynamic energy management. We do believe that combining sub-near-threshold with inexact circuits on the hardware side and, in addition, extending this with intelligent and adaptive power management on the software side will produce outstanding results in terms of energy reduction, i.e., at least one order of magnitude, in IoT applications. The project will contribute along three research directions: (1) approximate, ultra low-power circuit design, (2) modeling and analysis of variable levels of computation precision in applications, and (3) accuracy-energy trade-offs in software.

8.4. International Initiatives

8.4.1. Inria Associate Teams

8.4.1.1. HARDIESSE

Title: Heterogeneous Accelerators for Reconfigurable Dynamic, Energy efficient, Secure Systems
International Partner (Institution - Laboratory - Researcher):

University of Massachusetts at Amherst (United States) - Reconfigurable Computing Group - Russel Tessier

Start year: 2014

See also: <https://team.inria.fr/cairn/hardiesse/>

Rapid evolutions of applications and standards require frequent in-the-field system modifications and thus strengthens the need for adaptive devices. This need for a strong flexibility, combined with technology evolution (and the so-called power wall) has motivated the surge towards the use of multiple processor cores on a single chip (MPSoC). While it is now clear that we have entered the multi-core era, it is however indisputable that, especially for energy-efficient embedded systems, these architectures will have to be heterogeneous, by combining processor cores and specialized accelerators. We foresee a need for systems able to continuously adapt themselves to changing environments where software updates alone will not be enough for tackling energy management and error tolerance challenges. We believe that a dynamic and transparent adaptation of the hardware structure is the key to success. Security will also be an important challenge for embedded devices. Protections against physical attacks will have to be integrated in all secured components. In this Associated Team, we study new reconfigurable structures for such hardware accelerators with specific focus on: energy efficiency, runtime dynamic reconfiguration, security, and verification.

8.4.2. Inria International Partners

8.4.2.1. Declared Inria International Partners

8.4.2.1.1. LRS

Title: Loop unRolling Stones: compiling in the polyhedral model

International Partner (Institution - Laboratory - Researcher):

Colorado State University (United States) - Department of Computer Science - Prof. Sanjay Rajopadhye

8.4.2.1.2. HARAMCOP

Title: Hardware accelerators modeling using constraint-based programming

International Partner (Institution - Laboratory - Researcher):

Lund University (Sweden) - Department of Computer Science - Prof. Krzysztof Kuchcinski

8.4.2.1.3. SPINACH

Title: Secure and low-Power sensor Networks Circuits for Healthcare embedded applications

International Partner (Institution - Laboratory - Researcher):

University College Cork (Ireland) - Department of Electrical and Electronic Engineering - Prof. Liam Marnane and Prof. Emanuel Popovici

Arithmetic operators for cryptography, side channel attacks for security evaluation, energy-harvesting sensor networks, and sensor networks for health monitoring.

8.4.2.2. Informal International Partners

Imec (Belgium), Fault-tolerant computing architectures.

Ecole Polytechnique Fédérale de Lausanne - EPFL (Switzerland), Optimization of embedded systems using fixed-point arithmetic, approximate computing.

Technical University of Madrid - UPM (Spain), Optimization of embedded systems using fixed-point arithmetic.

LSSI laboratory, Québec University in Trois-Rivières (Canada), Design of architectures for digital filters and mobile communications.

Department of Electrical and Computer Engineering, University of Patras (Greece), Wireless Sensor Networks, Worst-Case Execution Time, priority scheduling, loop transformations for memory optimizations.

Karlsruhe Institute of Technology - KIT (Germany), Loop parallelization and compilation techniques for embedded multicores.

Ruhr - University of Bochum - RUB (Germany), Reconfigurable architectures.

University of Science and Technology of Hanoi (Vietnam), Participation of several CAIRN's members in the Master ICT / Embedded Systems.

8.5. International Research Visitors

8.5.1. Visits of International Scientists

Prof. Maciej Cieselski, University of Massachusetts, Amherst, US, for three weeks in July. This visit was partly funded by HARDIESSE Inria Associate Team.

Prof. Daniel Massicotte, Université du Québec à Trois-Rivières, CA, for three weeks in December. This visit was funded by ISTIC.

Maroua Gam, LabTim (Technologie Imagerie Médicale), Monastir, Tunisia, for one month in March.

8.5.2. Visits to International Teams

Angeliki Kritikakou visited University of Patras, Greece, for 1 week in November. This visit was funded by U. Rennes 1.

Patrice Quinton visited University of Massachusetts, Amherst, US, for 1 week in December. This visit was funded by HARDIESSE Inria Associate Team.

Tomofumi Yuki visited University of Arizona, US, in June.

8.5.2.1. Sabbatical programme

Casseau Emmanuel

Date: Aug 2016 - Jul 2017

Institution: **University of Auckland** (New Zealand), Parallel and Reconfigurable Research Lab. of the Electrical and Computer Engineering department.

The goal of the project is to propose dynamic mapping and scheduling algorithms dedicated to unreliable heterogeneous platforms, enabling self-adaptive and resource-aware computing.

9. Dissemination

9.1. Promoting Scientific Activities

9.1.1. Scientific Events Selection

9.1.1.1. General Chair, Scientific Chair

E. Casseau was General Co-Chair of DASIP, Conference on Design and Architectures for Signal and Image Processing, October 12-14, 2016.

S. Derrien was Co-Chair of WRC, 10th HiPEAC Workshop on Reconfigurable Computing, January 18-20, 2016 (co-located with HiPEAC 2016).

T. Yuki was Co-Chair of IMPACT, 6th International Workshop on Polyhedral Compilation Techniques, January 18-20, 2016 (co-located with HiPEAC 2016).

9.1.1.2. Chair of Conference Program Committees

O. Sentieys was Track Chair at IEEE NEWCAS.

9.1.1.3. Member of the Conference Program Committees

D. Chillet was member of the technical program committee of HiPEAC RAPIDO, HiPEAC WRC, MCSoc, DCIS, ComPAS, DASIP, LP-EMS, ARC.

S. Derrien was a member of technical program committee of IEEE FPL and ARC conferences and of WRC and Impact workshops.

O. Sentieys was a member of technical program committee of IEEE/ACM DATE, IEEE FPL, ACM ENSSys, ACM SBCCI, IEEE ReConFig, CROWNCOM, FSP, FPGA4GPC.

T. Yuki was a member of technical program committee of SC'16, The International Conference for High Performance Computing, Networking, Storage and Analysis.

9.1.2. Journal

9.1.2.1. Member of the Editorial Boards

D. Chillet is member of the Editor Board of Journal of Real-Time Image Processing (JRTIP).

O. Sentieys is member of the editorial board of Journal of Low Power Electronics and International Journal of Distributed Sensor Networks.

A. Tisserand is Associate Editor of IEEE Transactions on Computers. He is a member of the editorial board of the International Journal of High Performance Systems Architecture, Inderscience.

9.1.3. Invited Talks

O. Sentieys gave an invited talk at FETCH (École d'hiver Francophone sur les Technologies de Conception des Systèmes embarqués Hétérogènes), Villard-de-Lans, France, in January 2016 on "Approximate Computing and Flexible Circuits for the IoT".

T. Yuki gave a half-day lecture at EJCP 2016, École Jeunes Chercheurs en Programmation, Lille.

T. Yuki gave an invited talk at University of Arizona in June 2016 on "Optimizing Compilers in High-Level Synthesis".

9.1.4. Leadership within the Scientific Community

D. Chillet is member of the Board of Directors of GretsI Association.

F. Charot, O. Sentieys and A. Tisserand are members of the steering committee of a CNRS spring school for graduate students on embedded systems architectures and associated design tools (ARCHI).

O. Sentieys and A. Tisserand are members of the steering committee of a CNRS spring school for graduate students on low-power design (ECOFAC).

A. Tisserand is co-organizer and president of scientific council of Seminar on Security of Embedded Electronic Systems (IRISA-DGA).

O. Sentieys is a member of the steering committee of the GDR SOC-SIP.

9.1.5. Scientific Expertise

O. Sentieys served as a jury member in the EDAA Outstanding Dissertations Award (ODA).

9.2. Teaching - Supervision - Juries

9.2.1. Teaching

- E. Casseau: signal processing, 16h, ENSSAT (L3)
- E. Casseau: low power design, 6h, ENSSAT (M1)
- E. Casseau: real time design methodology, 24h, ENSSAT (M1)
- E. Casseau: computer architecture, 36h, ENSSAT (M1)
- E. Casseau: system on chip and verification, 10h, Master by Research (SISEA) and ENSSAT (M2)
- E. Casseau: high level synthesis, 12h, Master by Research (SISEA) and ENSSAT (M2)
- E. Casseau: advanced processor architectures, 25h, Univ. of Science and Tech. of Hanoi (M2)
- S. Derrien: component and system synthesis, 20h, Master by Research (MRI ISTIC) (M2)
- S. Derrien: computer architecture, 12h, ENS Rennes (L3)
- S. Derrien: computer architecture, 24h, ISTIC(L3)
- S. Derrien: introduction to operating systems, 8h, ISTIC(M1)
- S. Derrien: embedded architectures, 48h, ISTIC(M1)
- S. Derrien: high-level synthesis, 6h, ISTIC(M1)
- S. Derrien: software engineering project, 40h, ISTIC(M1)
- F. Charot: processor architecture, 25h Univ. of Science and Tech. of Hanoi (M1)
- D. Chillet: embedded processor architecture, 20h, ENSSAT (M1)
- D. Chillet: multimedia processor architectures, 24h, ENSSAT (M2)
- D. Chillet: low-power digital CMOS circuits, 6h, Telecom Bretagne (M2)
- C. Killian: digital electronics, 62h, IUT Lannion (L1)
- C. Killian: signal processing, 36h, IUT Lannion (L2)
- C. Killian: automated measurements, 56h, IUT Lannion (L2)
- C. Killian: measurement chain, 35h, IUT Lannion (L2)
- C. Killian: embedded systems programming, 12h, IUT Lannion (L2)
- C. Killian: automatic control, 9h, IUT Lannion (L2)
- A. Kritikakou: computer architecture 1, 50h, ISTIC, Univ. Rennes 1 (L3)
- A. Kritikakou: computer architecture 2, 50h, ISTIC, Univ. Rennes 1 (L3)
- A. Kritikakou: operating systems 1, 24h, ISTIC, Univ. Rennes 1 (L3)
- A. Kritikakou: operating systems 2, 64h, ISTIC, Univ. Rennes 1 (L3)
- A. Kritikakou: multitasking operating systems, 45h, ISTIC, Univ. Rennes 1 (M1)
- O. Sentieys: digital signal processing, 40h, ENSSAT (M1)
- O. Sentieys: VLSI integrated circuit design, 40h, ENSSAT(M1)
- O. Sentieys: high level synthesis, 16h, Master by Research (SISEA) and ENSSAT (M2)
- A. Tisserand: multiprocessor architectures, 20h, ENSSAT and Master by Research (SISEA) (M2)
- C. Wolinski: computer architectures, 92h, ESIR (L3)
- C. Wolinski: design of embedded systems, 48h, ESIR (M1)
- C. Wolinski: signal, image, architecture, 26h, ESIR (M1)
- C. Wolinski: programmable architectures, 10h, ESIR (M1)
- C. Wolinski: component and system synthesis, 10h, Master by Research (MRI ISTIC) (M2)

9.2.2. Teaching Responsibilities

- C. Wolinski is the Director of ESIR.

S. Derrien is the responsible of the first year of the Master of Computer Science at ISTIC since Sep. 2012.
O. Sentieys is responsible of the "Embedded Systems" major of the SISEA Master by Research.
D. Chillet is the responsible of the ICT Master of University of Science and Technology of Hanoi.
C. Killian is the responsible of the second year of the Physical Measurement DUT at IUT of Lannion.

ENSSAT stands for "*École Nationale Supérieure des Sciences Appliquées et de Technologie*" and is an "*École d'Ingénieurs*" of the University of Rennes 1, located in Lannion.

ISTIC is the Electrical Engineering and Computer Science Department of the University of Rennes 1.

ESIR stands for "*École supérieure d'ingénieur de Rennes*" and is an "*École d'Ingénieurs*" of the University of Rennes 1, located in Rennes.

9.2.3. Supervision

PhD: Florent Berthier, Study and Design of an Ultra Low Power Asynchronous Core for Sensor Networks, Dec. 2016, O. Sentieys, E. Beigne.

PhD: Ali Hassan El-Moussawi, Performance/Accuracy Trade-Off in Automatic Parallelization for Embedded Many-Core Platforms, Dec. 2016, S. Derrien.

PhD: Jérémie Métairie, Reconfigurable Arithmetic Units for Secure Cryptoprocessors, May 2016, A. Tisserand, E. Casseau.

PhD in progress: Benjamin Barrois, Approximate Computing: a New Paradigm for Energy-Efficient Computing Architectures, Oct. 2014, O. Sentieys.

PhD in progress: Franck Bucheron, Secure Virtualization for Embedded Systems, Oct. 2011, A. Tisserand.

PhD in progress: Gaël Deest, Computing with Errors: Error-Tolerant Machine Code Generation for Unreliable Embedded Hardware, Oct. 2013, S. Derrien, O. Sentieys.

PhD in progress: Gabriel Gallin, Hardware Arithmetic Units and Crypto-Processor for Hyperelliptic Curves Cryptography, Oct. 2014, A. Tisserand.

PhD in progress: Aymen Gammoudi, New Visual Adaptive Real-Time OS for Embedded Multi-Core Architecture, Oct. 2015, D. Chillet, M.Khalgui.

PhD in progress: Mael Gueguen, Improving the performance and energy efficiency of complex heterogeneous manycore architectures with on-chip data mining, Nov. 2016, O. Sentieys, A. Termier.

PhD in progress: Xuan Chien Le, Indirect Monitoring in Self-Powered Wireless Sensor Networks for Smart Grid and Building Automation, Oct. 2013, O. Sentieys, B. Vrigneau.

PhD in progress: Audrey Lucas, Software support resistant to passive and active attacks for asymmetric cryptography on (very) small computation cores, Jan. 2016, A. Tisserand.

PhD in progress: Jiating, Luo, Communication protocol exploration in the context of 3D integration of multiprocessors interconnected by Optical Network-on-Chip with energy constraints, Nov. 2014, D. Chillet, C. Killian, S. Le-Beux.

PhD in progress: Genevieve Ndour, Approximate Computing with High Energy Efficiency for Internet of Things Applications, Apr. 2016, A. Tisserand, A. Molnos (CEA LETI).

PhD in progress: Joel Ortiz Sosa, Study and design of a digital baseband transceiver for wireless network-on-chip architectures, Nov. 2016, O. Sentieys, C. Roland (Lab-STICC).

PhD in progress: Kleanthis Papachatzopoulos, Predictable and fault-tolerant multicore architecture, Oct. 2016, A. Kritikakou, O. Sentieys.

PhD in progress: Tara Petric, Approximate@runtime: Playing with accuracy at run-time for low-power flexible circuits in IoT nodes, Nov. 2016, T. Yuki, O. Sentieys.

PhD in progress: Van Dung Pham, Design space exploration in the context of 3D integration of multiprocessors interconnected by Optical Network-on-Chip, Dec 2014, O. Sentieys, D. Chillet, C. Killian, S. Le-Beux.

PhD in progress: Rafail Psiakis, A Self-Healing Reconfigurable Accelerator Structure for Fault-Tolerant Multi-Cores, Oct. 2015, A. Kritikakou, O. Sentieys.

PhD in progress: Rengarajan Ragavan, Ultra-Low Power Reconfigurable Architectures for Computing and Control in Wireless Sensor Networks, Oct. 2013, O. Sentieys, C. Killian.

PhD in progress: Simon Rokicki, Hybrid Hardware/Software Dynamic Compilation for Adaptive Embedded Systems, Oct. 2015, S. Derrien.

PhD in progress: Baptiste Roux, Architectural Exploration of a Low-Power Flexible Radio Embedded on Drones, Oct. 2014, O. Sentieys, M. Gautier.

PhD in progress: Nicolas Roux, Sensor-aided Non-Intrusive Appliance Load Monitoring: Detecting Activity of Devices through Low-Cost Wireless Sensors, Oct. 2016, O. Sentieys, B. Vrigneau.

PhD in progress: Mai-Thanh Tran, Hardware Synthesis of Flexible and Reconfigurable Radio from High-Level Language Dedicated to Physical Layer of Wireless Systems, Oct. 2013, E. Casseau, M. Gautier.

10. Bibliography

Major publications by the team in recent years

- [1] R. DAVID, S. PILLEMENT, O. SENTIEYS. *Energy-Efficient Reconfigurable Processors*, in "Low Power Electronics Design", C. PIGUET (editor), Computer Engineering, Vol 1, CRC Press, August 2004, chap. 20
- [2] S. DERRIEN, S. RAJOPADHYE, P. QUINTON, T. RISSET. *High-Level Synthesis of Loops Using the Polyhedral Model: The MMA α Software*, in "High-Level Synthesis From Algorithm to Digital Circuit", P. COUSSY, A. MORAWIEC (editors), Springer Netherlands, 2008, pp. 215-230, <http://dx.doi.org/10.1007/978-1-4020-8588-8>
- [3] C. GUY, B. COMBEMALE, S. DERRIEN, J. STEEL, J.-M. JÉZÉQUEL. *On Model Subtyping*, in "8th European Conference on Modelling Foundations and Applications (ECMFA)", Kgs. Lyngby, Denmark, July 2012, <http://hal.inria.fr/hal-00695034>
- [4] C. HURIAUX, A. COURTAY, O. SENTIEYS. *Design Flow and Run-Time Management for Compressed FPGA Configurations*, in "IEEE/ACM Design, Automation and Test in Europe (DATE)", March 2015, <https://hal.inria.fr/hal-01089319>
- [5] J.-M. JÉZÉQUEL, B. COMBEMALE, S. DERRIEN, C. GUY, S. RAJOPADHYE. *Bridging the Chasm Between MDE and the World of Compilation*, in "Journal of Software and Systems Modeling (SoSyM)", October 2012, vol. 11, n^o 4, pp. 581-597 [DOI : 10.1007/s10270-012-0266-8], <https://hal.inria.fr/hal-00717219>
- [6] B. LE GAL, E. CASSEAU, S. HUET. *Dynamic Memory Access Management for High-Performance DSP Applications Using High-Level Synthesis*, in "IEEE Transactions on VLSI Systems", 2008, vol. 16, n^o 11, pp. 1454-1464
- [7] K. MARTIN, C. WOLINSKI, K. KUCHCINSKI, A. FLOCH, F. CHAROT. *Constraint Programming Approach to Reconfigurable Processor Extension Generation and Application Compilation*, in "ACM transactions on Reconfigurable Technology and Systems (TRETTS)", June 2012, vol. 5, n^o 2, pp. 1-38, <http://doi.acm.org/10.1145/2209285.2209289>

- [8] D. MENARD, D. CHILLET, F. CHAROT, O. SENTIEYS. *Automatic Floating-point to Fixed-point Conversion for DSP Code Generation*, in "Proc. ACM/IEEE CASES", October 2002
- [9] D. MENARD, O. SENTIEYS. *Automatic Evaluation of the Accuracy of Fixed-point Algorithms*, in "IEEE/ACM Design, Automation and Test in Europe (DATE-02)", Paris, March 2002
- [10] S. PILLEMENT, O. SENTIEYS, R. DAVID. *DART: A Functional-Level Reconfigurable Architecture for High Energy Efficiency*, in "EURASIP Journal on Embedded Systems (JES)", 2008, pp. 1-13
- [11] R. ROCHER, D. MÉNARD, O. SENTIEYS, P. SCALART. *Analytical Approach for Numerical Accuracy Estimation of Fixed-Point Systems Based on Smooth Operations*, in "IEEE Transactions on Circuits and Systems. Part I, Regular Papers", October 2012, vol. 59, n^o 10, pp. 2326 - 2339 [DOI : 10.1109/TCSI.2012.2188938], <http://hal.inria.fr/hal-00741741>
- [12] C. WOLINSKI, M. GOKHALE, K. MCCABE. *A polymorphous computing fabric*, in "IEEE Micro", 2002, vol. 22, n^o 5, pp. 56–68
- [13] C. WOLINSKI, K. KUCHCINSKI, E. RAFFIN. *Automatic Design of Application-Specific Reconfigurable Processor Extensions with UPaK Synthesis Kernel*, in "ACM Trans. on Design Automation of Elect. Syst.", 2009, vol. 15, n^o 1, pp. 1–36, <http://doi.acm.org/10.1145/1640457.1640458>
- [14] S. WULIANG, B. COMBEMALE, S. DERRIEN, R. FRANCE. *Using Model Types to Support Contract-Aware Model Substitutability*, in "9th European Conference on Modelling Foundations and Applications (ECMFA)", Montpellier, France, P. VAN GORP, T. RITTER, L. ROSE (editors), LNCS, Springer-Verlag Berlin Heidelberg, 2013, vol. 7949, pp. 118-133 [DOI : 10.1007/978-3-642-39013-5_9], <http://hal.inria.fr/hal-00808770>

Publications of the year

Doctoral Dissertations and Habilitation Theses

- [15] F. BERTHIER. *Design of an ultra low power processor for wireless sensor nodes*, Université de Rennes 1, France, December 2016, <https://hal.inria.fr/tel-01423146>
- [16] A. H. EL MOUSSAWI. *SIMD-aware Word Length Optimization for Floating-point to Fixed-point Conversion targeting Embedded Processors*, Université de Rennes 1, December 2016, <https://hal.inria.fr/tel-01425642>
- [17] J. MÉTAIRIE. *Contributions to $GF(2^m)$ arithmetic operators for elliptic curve cryptography*, Université Rennes 1, May 2016, <https://hal.archives-ouvertes.fr/tel-01324924>
- [18] J. MÉTAIRIE. *Contributions to $GF(2^m)$ Operators for Cryptographic Purposes*, Université Rennes 1, May 2016, <https://tel.archives-ouvertes.fr/tel-01387919>

Articles in International Peer-Reviewed Journals

- [19] M. M. ALAM, E. BEN HAMIDA, O. BERDER, O. SENTIEYS, D. MENARD. *A Heuristic Self-Adaptive Medium Access Control for Resource-Constrained WBAN Systems*, in "IEEE Access", April 2016, vol. 4, pp. 1287-1300, <https://hal.archives-ouvertes.fr/hal-01396104>

- [20] F. BERTHIER, E. BEIGNE, F. HEITZMANN, O. DEBICKI, J.-F. CHRISTMANN, A. VALENTIAN, O. BILLOINT, E. AMAT, D. MORCHE, S. CHAIRAT, O. SENTIEYS. *UTBB FDSOI suitability for IoT applications: Investigations at device, design and architectural levels*, in "Solid-State Electronics", 2016, vol. 125, pp. 14 - 24 [DOI : 10.1016/J.SSE.2016.09.003], <https://hal.inria.fr/hal-01423144>
- [21] K. BIGOU, A. TISSERAND. *Binary-Ternary Plus-Minus Modular Inversion in RNS*, in "IEEE Transactions on Computers", November 2016, vol. 65, n^o 11, pp. 3495-3501 [DOI : 10.1109/TC.2016.2529625], <https://hal.inria.fr/hal-01314268>
- [22] R. BONAMY, S. BILAVARN, D. CHILLET, O. SENTIEYS. *Power Modeling and Exploration of Dynamic and Partially Reconfigurable Systems*, in "Journal of Low Power Electronics", September 2016, n^o September, <https://hal.archives-ouvertes.fr/hal-01345664>
- [23] M. FYRBIAK, S. ROKICKI, N. BISSANTZ, R. TESSIER, C. PAAR. *Hybrid Obfuscation to Protect against Disclosure Attacks on Embedded Microprocessors*, in "IEEE Transactions on Computers", 2017, <https://hal.inria.fr/hal-01426565>
- [24] V. KELEFOURAS, A. KRITIKAKOU, I. MPORAS, V. KOLONIAS. *A high performance Matrix-Matrix Multiplication Methodology for CPU and GPU architectures*, in "Journal of Supercomputing", 2016, pp. 1-41 [DOI : 10.1007/s11227-015-1613-7], <https://hal.archives-ouvertes.fr/hal-01255183>
- [25] A. KRITIKAKOU, F. CATTHOOR, V. KELEFOURAS, C. GOUTIS. *Array Size Computation under Uniform Overlapping and Irregular Accesses*, in "ACM Transactions on Design Automation of Electronic Systems (TODAES)", 2016, <https://hal.archives-ouvertes.fr/hal-01239705>
- [26] P. A. M. OLIVEIRA, R. J. CINTRA, F. M. BAYER, S. KULASEKERA, A. MADANAYAKE. *Low-complexity Image and Video Coding Based on an Approximate Discrete Tchebichef Transform*, in "IEEE Transactions on Circuits and Systems for Video Technology", January 2016 [DOI : 10.1109/TCSVT.2016.2515378], <https://hal.inria.fr/hal-01319513>
- [27] L.-Q.-V. TRAN, A. DIDIOUI, C. BERNIER, G. VAUMOURIN, F. BROEKAERT, A. FRITCH. *Co-Simulating Complex Energy Harvesting WSN Applications: An In-Tunnel Wind Powered Monitoring Example*, in "International Journal of Sensor Networks", 2016, <https://hal.inria.fr/hal-01264265>
- [28] S. WANG, C. XIAO, W. LIU, E. CASSEAU. *A comparison of heuristic algorithms for custom instruction selection*, in "Microprocessors and Microsystems: Embedded Hardware Design (MICPRO)", August 2016, vol. 45, n^o A, 11 p. , <https://hal.inria.fr/hal-01354991>

Invited Conferences

- [29] O. SENTIEYS, J. SEPÚLVEDA, S. LE BEUX, J. LUO, C. KILLIAN, D. CHILLET, I. O'CONNOR, H. LI. *Design Space Exploration of Optical Interfaces for Silicon Photonic Interconnects*, in "2th International Workshop on Optical/Photonic Interconnects for Computing Systems (OPTICS Workshop), co-located with IEEE/ACM Design Automation and Test in Europe (DATE'16)", Dresden, Germany, March 2016, <https://hal.inria.fr/hal-01293506>

International Conferences with Proceedings

- [30] B. BARROIS, K. PARASHAR, O. SENTIEYS. *Leveraging Power Spectral Density for Scalable System-Level Accuracy Evaluation*, in "IEEE/ACM Conference on Design Automation and Test in Europe (DATE)", Dresden, Germany, March 2016, 6 p. , <https://hal.inria.fr/hal-01253494>
- [31] B. BARROIS, O. SENTIEYS, D. MENARD. *The Hidden Cost of Functional Approximation Against Careful Data Sizing – A Case Study*, in "IEEE/ACM Design Automation and Test in Europe (DATE)", Lausanne, France, 2017, <https://hal.inria.fr/hal-01423147>
- [32] K. BIGOU, A. TISSERAND. *Hybrid Position-Residues Number System*, in "ARITH: 23rd Symposium on Computer Arithmetic", Santa Clara, CA, United States, J. HORMIGO, S. OBERMAN, N. REVOL (editors), IEEE, July 2016, <https://hal.inria.fr/hal-01314232>
- [33] G. DEEST, N. ESTIBALS, T. YUKI, S. DERRIEN, S. RAJOPADHYE. *Towards Scalable and Efficient FPGA Stencil Accelerators*, in "6th International Workshop on Polyhedral Compilation Techniques (IMPACT'16), held with HIPEAC'16", Prague, Czech Republic, Proceedings of the IMPACT series, <http://impact.gforge.inria.fr/>, January 2016, <https://hal.inria.fr/hal-01254778>
- [34] A. H. EL MOUSSAWI, S. DERRIEN. *Superword Level Parallelism aware Word Length Optimization*, in "DATE - Design, Automation & Test in Europe Conference & Exhibition", Lausanne, Switzerland, D. ATIENZA, G. D. NATALE (editors), IEEE, March 2017, <https://hal.inria.fr/hal-01425550>
- [35] N. ESTIBALS, G. DEEST, A. EL-MOUSSAWI, S. DERRIEN. *System level synthesis for virtual memory enabled hardware threads*, in "Design, Automation & Test in Europe Conference & Exhibition", Dresden, France, March 2016, <https://hal.inria.fr/hal-01424772>
- [36] A. GAMMOUDI, A. BENZINA, M. KHALGUI, D. CHILLET. *New Reconfigurable Middleware for Adaptive RTOS in Ubiquitous Devices*, in "10th International Conference on Mobile Ubiquitous Computing, Systems, Services and Technologies", Venise, Italy, October 2016, <https://hal.inria.fr/hal-01401716>
- [37] A. GAMMOUDI, A. BENZINA, M. KHALGUI, D. CHILLET. *Real-Time Scheduling of Reconfigurable Battery-Powered Multi-Core Platforms*, in "28th International Conference on Tools with Artificial Intelligence", San Jose, United States, November 2016, <https://hal.inria.fr/hal-01401712>
- [38] A. GAMMOUDI, A. BENZINA, M. KHALGUI, D. CHILLET, A. GOUBAA. *Reconf-Pack: A Simulator for Reconfigurable Battery-Powered Real-Time Systems*, in "30th European Simulation and Modelling Conference", Las Palmas, Spain, October 2016, <https://hal.inria.fr/hal-01401706>
- [39] C. HURIAUX, O. SENTIEYS, R. TESSIER. *Effects of I/O Routing through Column Interfaces in Embedded FPGA Fabrics*, in "FPL - 26th International Conference on Field Programmable Logic and Applications", Lausanne, Switzerland, IEEE, August 2016, <https://hal.inria.fr/hal-01341156>
- [40] J. LUO, A. ELANTABLY, D. D. PHAM, C. KILLIAN, D. CHILLET, S. LE BEUX, O. SENTIEYS, I. O'CONNOR. *Performance and Energy Aware Wavelength Allocation on Ring-Based WDM 3D Optical NoC*, in "Design, Automation & Test in Europe Conference & Exhibition (DATE) 2017", Lausanne, Switzerland, March 2017, <https://hal.inria.fr/hal-01416958>
- [41] T.-H. NGUYEN, P. SCALART, M. GAY, L. BRAMERIE, C. PEUCHERET, T. NGUYEN-TI, M. GAUTIER, O. SENTIEYS, J.-C. SIMON, M. JOINDOT. *Blind Adaptive Transmitter IQ Imbalance Compensation in M-QAM Optical Coherent Systems*, in "2016 IEEE International Conference on Communication (ICC)

2016)", Kuala Lumpur, Malaysia, Communications (ICC), 2016 IEEE International Conference on, May 2016 [DOI : 10.1109/ICC.2016.7510925], <https://hal.archives-ouvertes.fr/hal-01337225>

- [42] T. H. NGUYEN, P. SCALART, M. GAY, L. BRAMERIE, C. PEUCHERET, O. SENTIEYS, J.-C. SIMON, M. JOINDOT. *Bi-harmonic decomposition-based maximum loglikelihood estimator for carrier phase estimation of coherent optical M-QAM*, in "Optical Fiber Communication Conference (OFC 2016)", Anaheim, CA, United States, Optical Fiber Communication Conference 2016, OSA (ISBN: 978-1-943580-07-1), March 2016, vol. DSP for Coherent Systems (Tu3K), Tu3K.3 [DOI : 10.1364/OFC.2016.Tu3K.3], <https://hal.archives-ouvertes.fr/hal-01309175>
- [43] R. RAGAVAN, B. BARROIS, C. KILLIAN, O. SENTIEYS. *Pushing the Limits of Voltage Over-Scaling for Error-Resilient Applications*, in "Design, Automation & Test in Europe Conference & Exhibition (DATE 2017)", Lausanne, Switzerland, March 2017, <https://hal.archives-ouvertes.fr/hal-01417665>
- [44] R. RAGAVAN, C. KILLIAN, O. SENTIEYS. *Adaptive Overclocking and Error Correction Based on Dynamic Speculation Window*, in "ISVLSI", Pittsburgh, United States, 2016 IEEE Computer Society Annual Symposium on VLSI (ISVLSI), July 2016, pp. 325 - 330 [DOI : 10.1109/ISVLSI.2016.13], <https://hal.inria.fr/hal-01416945>
- [45] S. ROKICKI, E. ROHOU, S. DERRIEN. *Hardware-Accelerated Dynamic Binary Translation*, in "IEEE/ACM Design, Automation & Test in Europe Conference & Exhibition (DATE)", Lausanne, Switzerland, March 2017, <https://hal.inria.fr/hal-01423639>
- [46] B. ROUX, M. GAUTIER, O. SENTIEYS, S. DERRIEN. *Communication-Based Power Modelling for Heterogeneous Multiprocessor Architecture*, in "IEEE 10th International Symposium on Embedded Multicore /Many-core Systems-on-Chip (MCSoc 2016)", Lyon, France, September 2016, <https://hal.inria.fr/hal-01412835>
- [47] M.-T. TRAN, M. GAUTIER, E. CASSEAU. *On the FPGA-based implementation of a flexible waveform from a high-level description: Application to LTE FFT case study*, in "EAI International Conference on Cognitive Radio Oriented Wireless Networks (Crowncom16)", Grenoble, France, May 2016, <https://hal.inria.fr/hal-01302652>

National Conferences with Proceedings

- [48] J. LUO, D. CHILLET, C. KILLIAN, S. LE BEUX, I. O 'CONNOR, O. SENTIEYS. *Crosstalk noise aware wavelength allocation in WDM 3D ONoC*, in "Colloque National du GDR SoC-SiP", Nantes, France, June 2016, <https://hal.inria.fr/hal-01406355>
- [49] J. LUO, D. CHILLET, C. KILLIAN, S. LE BEUX, I. O 'CONNOR, O. SENTIEYS. *Wavelength spacing optimization to reduce crosstalk in WDM 3D ONoC*, in "Conférence d'informatique en Parallélisme, Architecture et Système", Lorient, France, July 2016, <https://hal.inria.fr/hal-01406341>
- [50] V. D. PHAM, D. CHILLET, C. KILLIAN, S. LE BEUX, I. O 'CONNOR, O. SENTIEYS. *Gestion de la consommation d'un ONoC intégré dans un MPSoC*, in "Colloque National du GDR SoC-SiP", Nantes, France, June 2016, <https://hal.inria.fr/hal-01414341>
- [51] V. D. PHAM, C. KILLIAN, D. CHILLET, S. LE BEUX, O. SENTIEYS, I. O 'CONNOR. *Gestion de la consommation d'un réseau optique intégré dans un MPSoC*, in "Conférence d'informatique en Parallélisme, Architecture et Système", Lorient, France, July 2016, <https://hal.inria.fr/hal-01406347>

- [52] S. ROKICKI, E. ROHOU, S. DERRIEN. *Hybrid-JIT : Compilateur JIT Matériel/Logiciel pour les Processeurs VLIW Embarqués*, in "Conférence d'informatique en Parallélisme, Architecture et Système (Compas)", Lorient, France, July 2016, <https://hal.archives-ouvertes.fr/hal-01345306>

Conferences without Proceedings

- [53] G. DEEST, N. ESTIBALS, T. YUKI, S. DERRIEN, S. RAJOPADHYE. *Towards Scalable and Efficient FPGA Stencil Accelerators*, in "IMPACT'16", Prague, Czech Republic, January 2016, <https://hal.inria.fr/hal-01425018>
- [54] A. KRITIKAKOU, T. MARTY, C. PAGETTI, C. ROCHANGE, M. LAUER, M. ROY. *Multiplexing Adaptive with Classic AUTOSAR? Adaptive Software Control to Increase Resource Utilization in Mixed-Critical Systems*, in "Workshop CARS 2016 - Critical Automotive applications : Robustness & Safety", Göteborg, Sweden, CARS 2016 - Critical Automotive applications : Robustness & Safety, September 2016, <https://hal.archives-ouvertes.fr/hal-01375576>

Patents and standards

- [55] F. BERTHIER, E. BEIGNE, F. HEITZMANN, O. DEBICKI, O. SENTIEYS. *Cœur de processeur asynchrone et microcontrôleur de nœud de capteur communicant comportant un tel cœur de processeur*, 2016, n° 2016, <https://hal.inria.fr/hal-01423133>

Other Publications

- [56] P. GUILLOUX, A. TISSERAND. *Accurate Modeling of Fault Impact in Arithmetic Circuits*, October 2016, DASIP: Conference on Design and Architectures for Signal and Image Processing (Demo Night), Poster, <https://hal.inria.fr/hal-01404772>
- [57] P. GUILLOUX, A. TISSERAND. *Plateforme matérielle–logicielle d’émulation de fautes pour des opérateurs arithmétiques*, July 2016, 8 p. , Compas 2016 : Conférence d'informatique en Parallélisme, Architecture et Système, <https://hal.inria.fr/hal-01313051>
- [58] P. GUILLOUX, A. TISSERAND. *Plateforme matérielle–logicielle à bas coût pour l’émulation de fautes*, June 2016, Colloque du GDR SoC-SiP, Poster, <https://hal.inria.fr/hal-01346576>
- [59] J. LUO, V. D. PHAM, C. KILLIAN, D. CHILLET, S. LE BEUX, I. O'CONNOR, O. SENTIEYS. *POSTER: Wavelength Allocation for Efficient Communications on Optical Network-on-Chip*, October 2016, pp. 1656 - 1658, Conference on Design and Architectures for Signal and Image Processing, Poster [DOI : 10.1145/2810103.2810122], <https://hal.inria.fr/hal-01406328>
- [60] O. SENTIEYS, D. MENARD, K. PARASHAR, D. NOVO. *Fixed-point refinement, a guaranteed approach towards energy efficient computing*, January 2016, Tutorial, <https://hal.inria.fr/hal-01423184>
- [61] A. TISSERAND, G. GALLIN. *Hardware and Arithmetic for Hyperelliptic Curves Cryptography*, November 2016, CominLabs Days 2016, Poster, <https://hal.inria.fr/hal-01404755>
- [62] M.-T. TRAN, E. CASSEAU, M. GAUTIER. *Demo abstract : FPGA-based implementation of a flexible FFT dedicated to LTE standard*, October 2016, 2 p. , Conference on Design and Architectures for Signal and Image Processing (DASIP), Demo Night, Poster, <https://hal.inria.fr/hal-01354992>

- [63] Y. UGUEN, F. DE DINECHIN, S. DERRIEN. *Arithmetic Optimizations for High-Level Synthesis*, September 2016, working paper or preprint, <https://hal.inria.fr/hal-01373954>

References in notes

- [64] S. HAUCK, A. DEHON (editors). *Reconfigurable Computing: The Theory and Practice of FPGA-Based Computation*, Morgan Kaufmann, 2008
- [65] V. BAUMGARTE, G. EHLERS, F. MAY, A. NÜCKEL, M. VORBACH, M. WEINHARDT. *PACT XPP — A Self-Reconfigurable Data Processing Architecture*, in "The Journal of Supercomputing", 2003, vol. 26, n^o 2, pp. 167–184
- [66] C. BECKHOFF, D. KOCH, J. TORRESEN. *Portable module relocation and bitstream compression for Xilinx FPGAs*, in "24th Int. Conf. on Field Programmable Logic and Applications (FPL)", 2014, pp. 1–8
- [67] C. BOBDA. *Introduction to Reconfigurable Comp.: Architectures Algorithms and Applications*, Springer, 2007
- [68] D. BOGGS, G. BROWN, N. TUCK, K. VENKATRAMAN. *Denver: NVIDIA's First 64-bit ARM Processor*, in "Micro", 2015
- [69] S. BORKAR, A. A. CHIEN. *The Future of Microprocessors*, in "Commun. ACM", May 2011, vol. 54, n^o 5, pp. 67–77, <http://doi.acm.org/10.1145/1941487.1941507>
- [70] J. M. P. CARDOSO, P. C. DINIZ, M. WEINHARDT. *Compiling for reconfigurable computing: A survey*, in "ACM Comput. Surv.", June 2010, vol. 42, 13:1 p. , <http://doi.acm.org/10.1145/1749603.1749604>
- [71] K. COMPTON, S. HAUCK. *Reconfigurable computing: a survey of systems and software*, in "ACM Comput. Surv.", 2002, vol. 34, n^o 2, pp. 171–210, <http://doi.acm.org/10.1145/508352.508353>
- [72] J. CONG, H. HUANG, C. MA, B. XIAO, P. ZHOU. *A Fully Pipelined and Dynamically Composable Architecture of CGRA*, in "IEEE Int. Symp. on Field-Program. Custom Comput. Machines (FCCM)", 2014, pp. 9–16, <http://dx.doi.org/10.1109/FCCM.2014.12>
- [73] G. CONSTANTINIDES, P. CHEUNG, W. LUK. *Wordlength optimization for linear digital signal processing*, in "IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems", October 2003, vol. 22, n^o 10, pp. 1432- 1442
- [74] M. COORS, H. KEDING, O. LUTHJE, H. MEYR. *Fast Bit-True Simulation*, in "Proc. ACM/IEEE Design Automation Conference (DAC)", Las Vegas, June 2001, pp. 708-713
- [75] J. C. DEHNERT, B. K. GRANT, J. P. BANNING, R. JOHNSON, T. KISTLER, A. KLAIBER, J. MATTSON. *The Transmeta Code Morphing Software: Using Speculation, Recovery, and Adaptive Retranslation to Address Real-Life Challenges*, in "International Symposium on Code Generation and Optimization: Feedback-Directed and Runtime Optimization", 2003

- [76] R. H. DENNARD, F. H. GAENSSLEN, V. L. RIDEOUT, E. BASSOUS, A. R. LEBLANC. *Design of ion-implanted MOSFET's with very small physical dimensions*, in "IEEE Journal of Solid-State Circuits", 1974, vol. 9, n^o 5, pp. 256–268
- [77] A. HORMATI, M. KUDLUR, S. MAHLKE, D. BACON, R. RABBAH. *Optimus: efficient realization of streaming applications on FPGAs*, in "Proc. ACM/IEEE CASES", 2008, pp. 41–50
- [78] H. KALTE, M. PORRMANN. *REPLICA2Pro: Task Relocation by Bitstream Manipulation in Virtex-II/Pro FPGAs*, in "3rd Conference on Computing Frontiers (CF)", 2006, pp. 403–412
- [79] J.-E. LEE, K. CHOI, N. D. DUTT. *Compilation Approach for Coarse-Grained Reconfigurable Architectures*, in "IEEE Design and Test of Computers", 2003, vol. 20, n^o 1, pp. 26–33, <http://doi.ieeecomputersociety.org/10.1109/MDT.2003.1173050>
- [80] H. LEE, D. NGUYEN, J.-E. LEE. *Optimizing Stream Program Performance on CGRA-based Systems*, in "52nd IEEE/ACM Design Automation Conference", 2015, pp. 110:1–110:6, <http://doi.acm.org/10.1145/2744769.2744884>
- [81] B. MEI, S. VERNALDE, D. VERKEST, H. DE MAN, R. LAUWEREINS. *ADRES: An architecture with tightly coupled VLIW processor and coarse-grained reconfigurable matrix*, in "Proc. FPL", Springer, 2003, pp. 61–70
- [82] N. R. MINISKAR, S. KOHLI, H. PARK, D. YOO. *Retargetable Automatic Generation of Compound Instructions for CGRA Based Reconfigurable Processor Applications*, in "Proc. ACM/IEEE CASES", 2014, pp. 4:1–4:9, <http://doi.acm.org/10.1145/2656106.2656125>
- [83] Y. PARK, H. PARK, S. MAHLKE. *CGRA express: accelerating execution using dynamic operation fusion*, in "Proc. Int. Conf. on Compilers, Architecture, and Synthesis for Embedded Systems", New York, NY, USA, CASES'09, ACM, 2009, pp. 271–280, <http://doi.acm.org/10.1145/1629395.1629433>
- [84] A. PUTNAM ET AL.. *A reconfigurable fabric for accelerating large-scale datacenter services*, in "ACM/IEEE 41st International Symposium on Computer Architecture (ISCA)", June 2014, pp. 13–24, <http://dx.doi.org/10.1109/ISCA.2014.6853195>
- [85] G. THEODORIDIS, D. SOUDRIS, S. VASSILIADIS. 2, in "A survey of coarse-grain reconfigurable architectures and CAD tools", Springer Verlag, 2007
- [86] G. VENKATARAMANI, W. NAJJAR, F. KURDAHI, N. BAGHERZADEH, W. BOHM, J. HAMMES. *Automatic compilation to a coarse-grained reconfigurable system-on-chip*, in "ACM Trans. on Emb. Comp. Syst.", 2003, vol. 2, n^o 4, pp. 560–589, <http://doi.acm.org/10.1145/950162.950167>