PostDoc Position at Inria
Exploration of Energy Efficient Designs with Approximation

Inria project Team:	CAIRN
Scientific advisors:	Olivier Sentieys (olivier.sentieys@inria.fr)
	Tomofumi Yuki (tomofumi.yuki@inria.fr)
Duration:	18 months
Starting Date:	1 October 2017

Context and Objective

The context of this post-doctoral position is energy efficient computing with the emphasis on techniques for approximation. In embedded system design, the output quality is often sacrificed for energy efficiency and area cost. Having accuracy of the computation as a tunable parameter opens the door for many possibilities to reduce energy consumption by performing less computations and/or transferring less data.

Approximate computing has gained a lot of attention in the recent years, and there are many techniques for approximation available in the literature. However, methods to select the appropriate technique for a given application context is still lacking. Furthermore, there is little study on how approximating a kernel influences the entire system behavior, especially when more than one kernel is approximated. This project is not necessarily about developing new techniques for approximation, but to better understand/model the impact of the approximation. The goal is having energy efficient design, and we seek to find systematic methods to answer the key questions: when to approximate, what to approximate, and how to approximate.

The scope of the project is relatively open and applicants are expected to identify the direction that suits them the most as a function of their background and interest. Our target architecture is System on Chip platforms that are equipped with FPGAs and embedded processors. The project is expected to have both software and hardware components, and the communication between the FPGA and the cores creates additional opportunities. Examples of specific topics include:
- Studying the interaction between approximating the input feature set to a classifier and its output quality.
- Can we dynamically adapt the method of approximation at run-time?
- Can we reduce the cost of data transfer between the FPGA and the cores by possibly adding extra (redundant) computations?

Expected Profile:
- PhD in Computer Science
- Background in FPGA design and/or compilers.
- Familiarity with the C language and High-Level Synthesis tools.

What is valued the most is autonomy. We expect the postdoc to be motivated and capable of composing short and mid-term objectives themselves.

Bibliography

[1] Marc Baboulin, Alfredo Buttari, Jack Dongarra, Jakub Kurzak, Julie Langou, Julien Langou, Piotr Luszczek, and Stanimire Tomov. Accelerating scientific computations with mixed precision algorithms. Computer Physics Commu- nications, 180(12):2526–2533, 2009.

[2] Benjamin Barrois, Olivier Sentieys, and Daniel Menard. The Hidden Cost of Functional Approximation Against Careful Data Sizing – A Case Study. In Design, Automation & Test in Europe Conference & Exhibition (DATE 2017), Lausanne, France, 2017.

[3] Rengarajan Ragavan, Benjamin Barrois, Cedric Killian, and Olivier 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.