THE UNIVERSITY of York

Dr Martin A. Trefzer
Dipl.-Phys., PhD, Senior Member IEEE
Intelligent Systems and Nanoscience Research Group

Department of Electronic Engineering
HESLINGTON, YORK, YO10 5DD
Direct line +44 (0)1904 32-2346
Email : martin.trefzer@york.ac.uk

York, September 19, 2017

Review of Doctoral Thesis submitted to Brno University of Technology, Faculty of Information Technology, Department of Computer Systems

Title: Automated Multi-objective Parallel Evolutionary Circuit Design and Approximation
Author: Radek Hrbáček

Subject, contributions and scope

One of the main drivers of modern technology is power efficiency, closely followed by reliability. At the same time, both features are notoriously hard to achieve in electronic systems. All levels, including design, implementation, test and fabrication, are facing major challenges when creating reliable ultra-low power systems enabling today's connected world. This doctoral thesis addresses some of these challenges by combining evolutionary circuit design techniques with multi-objective optimisation (MOO) to create a novel kind of circuits, approximate circuits, where power saving and potentially increased robustness is achieved by trading off exact computation. An automated design methodology that is capable to balance accuracy and power with circuit robustness and complexity would be of immense value to expand research into low-power, distributed mobile systems, as well as to enable real-world industrial applications involving smart sensor systems, mobile technologies and internet-of-things (IoT) applications. I therefore consider the topic of the thesis as timely and relevant to the author's field of study.

Goals and objectives

The overall hypothesis and research objectives of the thesis are formulated in Chapter 1.3. The main hypothesis encompasses all three aspects of evolutionary design automation, multi-objective optimisation and approximation of digital circuits. This is then broken down into five smaller, more specific research objectives. The first three aim to develop the three aforementioned ingredients separately, the fourth objective then aims to bring the outcomes together into a design framework that can be applied to a range of real-world benchmark problems, and the fifth objective aims to validate the approach by interfacing with an industry-standard design flow.

The hypothesis and research objectives are clearly stated and formulated. Metrics of what would be considered a successful outcome of the different objectives are discussed in the

relevant sections and papers that form the body of the thesis. The five objectives nicely lead up to the main hypothesis.

Methodology

The research methodology adopted in this thesis is sound. The author reviews and critically discusses essential components of his envisaged system before selecting and further developing appropriate methodologies and algorithms.

The thesis begins with a discussion aimed at understanding and defining the notion and terminology of “approximate circuits”. It is nice to see a discussion involving both aspects accuracy and error. One thing that could have been made even clearer are the differences between precision and accuracy in computing. While a system can easily operate with a reduced precision, the user of a system usually expects a guaranteed level of overall accuracy. In my opinion, it is an interesting discussion to be had of what the implications of reducing either one in a system might be. Metrics that later form the basis of the quality metric (fitness function) of the evolved circuits are also discussed here.

With regards to selecting an appropriate representation for circuit evolution, the methodology followed is to select Cartesian genetic programming (CGP). This is a reasonable and safe choice as the candidate’s group has significant experience in this field, allowing the candidate to capitalize on proven methods and focussing on more novel aspects such as the design of useful multi-objective fitness functions and developing extensions to parallelise CGP to run on Xeon Phi. The choice of using SAT solvers to speed-up fitness evaluation also makes a lot of sense, particularly since one of the stated goals of the thesis is to parallelise and speed-up evolutionary search.

The author discusses a number of methods, including hardware accelerators on ZYNQ, to speed-up evolutionary search. However, it is not made entirely clear in the summary which combination of parallelisation/hardware accelerators have been used to conduct the experiments. My understanding from the papers is that ultimately multiple processors have been used to solve the large MO populations in parallel?

The author then investigates a number of multi-objective optimisation algorithms and selects the one he deems most appropriate. This is also a sensible approach. One thing that would have been useful to know is whether (and if yes, which) modifications to standard NSGA-II were made. If no changes were made, then it would have been good to see at least a short discussion why not.

Despite the thesis being a collection of papers framed with background, summary and conclusion chapters. It would have been nice to see a bit more discussion in the body of text tying the papers together. In Chapter 1.2, the summary and discussion of open problems could have been expanded a little bit and possibly relate open research questions to real-world applications.

In Chapter 4 it would have been useful to provide more detailed discussion/critical reflection and put the contributions into context, refer back to the different papers that contributed to achieving specific contributions and critically reflect on the approach taken and the results obtained.

There are only minor typos etc. in the thesis, for which I will provide a marked-up PDF.

Organisation, style and language

The thesis is cumulative, i.e. a collection of papers with an accompanying body of text introducing and motivating the field of study and the background in Chapters 1 and 2. This is followed by a summary of all papers included in the thesis in Chapter 3 and a critical discussion of the findings and contributions of the thesis in Chapter 4. The structure and presentation of the thesis is very clear. Generally, the use of language and grammar is excellent throughout the thesis and the contributed papers. The writing style is clear and concise. There are a good overall number of citations including the ones from Chapters 1-4 and the seven contributed papers. The amount of introduction and background is fine, but the level of detail and amount of discussion, mainly in Chapters 1.2 and 4 as mentioned before, could be increased.

Author's publications

The author's list of publications is impressive, and it is a great achievement to obtain recognition from a wider audience in the form of awards, e.g. DATE'17 and "Humies" at GECCO'14. There are three journal publications and four conference publications, all published in high-ranking journals or international conferences. In addition to the seven papers forming this thesis, the author has authored or was involved in ten further publications in internationally acclaimed conferences.

All core ideas and contributions of the thesis have been published and I consider the number and quality standard of the publications as more than sufficient to defend this doctoral dissertation.

Questions

(1) The author discusses a number of methods, including hardware accelerators on ZYNQ, to speed-up evolutionary search. Could you explain why you have used the multi-core approach, rather than hardware acceleration or a combination of both? Would there be any benefit in looking at parallel evolutionary algorithms, e.g. island models?

(2) NSGA-II was adopted as the most suitable for multi-objective circuit optimisation. Could you explain which modifications to standard NSGA-II were made. If no changes were made, then could you explain why the standard implementation was sufficient?

(3) Could you expand on current open research questions and relate them to real-world applications. For instance, what are the challenges in practice when designing with limited precision or accuracy? What guarantees can be given for approximate circuits?

(4) Can you give any examples of how your approach or your resulting circuits compete with other approaches from the field of approximate circuits. For instance, how does the power consumption of your circuits compare with ANNs approximating functions, etc.?

Summary

This is a very well-written doctoral thesis presenting novel contributions and scientific results in the area of multi-objective evolutionary design of approximate circuits. In particular, it is very impressive to see an automated approach developed and integrated into industry-standard design flow. It is great to see intellectual property generated from research, and good practice

to disseminate these results in the form of a library of approximate circuits that can be used in synthesis.

In my professional opinion, the author of the thesis has proven his ability to conduct research and achieve original scientific results. The thesis contains both new and original results that have already been published in high-quality international journals or conferences. Therefore, **I recommend the thesis for presentation and defence**.

Dr. Martin A. Trefzer
Reviewer of the thesis