McGill | Graduate and Postdoctoral Studies

Internal Examiner Report – Doctoral Thesis - Revised

Report due date: Mar 10, 2020

Name of Student: Jing CHEN

Degree/Unit: Doctor of Philosophy

Thesis title: Hardware acceleration for elementary functions and RISC-V processor

Thank you for your valuable contribution to this student's examination. As a thesis examiner, you will complete this form and attach a written report providing a detailed justification of your evaluation. The deadline to send **this form and your written report** to thesiscoordinator.gps@mcgill.ca is Mar 10, 2020 .

Please note that a late report has serious academic and financial consequences for the student.

Evaluation of the Thesis: Complete the evaluation grid below and comment on the criteria in your written report.

Criteria for Evaluation of Thesis	Excellent Top 10%	Very Good	Good	Satisfactory	Unsatisfactory
1. Makes an original contribution to knowledge	○	X	○	○	○
2. Advances knowledge in the field	○	X	○	○	○
3. Is situated in a broader context and appropriately acknowledges the larger field of research (e.g., citations/references)	○	○	○	X	○
4. Details methodology and methods	○	X	○	○	○
5. Reports results clearly	○	X	○	○	○
6. Justifies analyses and conclusions	○	○	X	○	○
7. Discusses implications	○	○	○	X	○
8. Is presented appropriately for disciplinary norms (grammar, style, coherence, cohesion)	○	○	X	○	○
9. Complies with McGill's guidelines for thesis preparation	○	○	X	○	○

Overall Recommendation: Select one.

Recommendation	Select ONE
PASSED – The thesis is **ready to proceed** to the Oral Defence. • Your written report must include any recommendations for minor revisions to the thesis (i.e. stylistic or editorial changes). • Your written report **must include questions** to be asked of the student at the Oral Defence.	X
NOT PASSED – Major revisions to the thesis are required before proceeding to the Oral Defence. • At least one of the Criteria for the Evaluation of the Thesis must be judged as unsatisfactory if the thesis is NOT PASSED. • Your written report must include a detailed description of all the shortcomings that have informed your decision, including an itemized list of substantive issues to be addressed before the thesis can be given a PASS and proceed to the Oral Defence. Note: If this is the first "NOT PASSED" assessment, the student will be given one opportunity to revise and resubmit the thesis.	○

DATE : March 10, 2020 SIGNATURE: [redacted]

General Comments

I remain well pleased with the technical work done in this thesis research. There is excellent attention to detail in the design of the optimized elementary functions, and the work is quite relevant to continued efforts at using external accelerators for specialized purposes; recent interest in optimizing machine learning alone makes this a research area with significant practical contribution potential. There are other RISC-V implementations, with varying characteristics. An HLS-based approach, however, offers notable flexibility, and even if the performance and characteristics of the design given here (IPC, area, etc.) are not dramatically better than the gamut of other designs, showing an HLS approach is competitive has significant value. Additional effort at demonstrating the value of this flexibility would also be welcome of course.

My comments on the potential for extending this work are also still relevant. Actual application of the techniques to real workloads that could benefit from an external, pipelined accelerator, for example, would showcase the practical utility of the work, and also validate the assumptions driving the optimized designs. Integration of floating point and the elementary function designs into the soft processor seems a rather obvious progression for the work, and would have enabled much more extensive experimentation and validation of the relative contribution of the specific elementary functions previously optimized.

My prior concerns were focused on presentation issues, the extent of literature review, and clarification of the thesis contributions. In this areas Chen has greatly improved the thesis. More specifically,

- Presentation is dramatically better in this version. I no longer consider language issues a concern at all, and major restructuring has reduced repetition and greatly improved readability. Some typos remain of course, but the work is now at a reasonably professional level, and I do not have any major concerns with the thesis writing or structure.
- The extent of literature review is also significantly improved. Each chapter now includes a non-trivial discussion of related work. I appreciate the effort at concisely summarizing the different, comparable systems, and this is accompanied by detailed explanations of major differences in important characteristics. In this Chen demonstrates a good awareness of other work in the field. These sections, however, still tend to be much more descriptive more than insightful, and more in-depth analysis of the differences, making specific and clear arguments or comparisons with the thesis work at a high level (rather than just pointing out differences in raw metrics) would have made more interesting reading, and better demonstrated the broad and comprehensive understanding of how the thesis work is situated that is expected of a PhD thesis. Overall I find the improvements to related work are now sufficient, even if further improvements are possible.
- My final concern, explaining the actual thesis contributions in relation to prior publications, including MSc work, is also addressed. In particular, the extent to which Chen's MSc thesis work is used (extended) is much easier to see. There is some overlap here, as the PhD work builds on it, but there is also significant additional and novel contribution being made in this work.

I thus find the thesis is now acceptable to continue to the oral defence.

I do not have further substantial changes. Typos and minor grammatical issues are indicated on the returned thesis copy. Below is list a number of other minor issues that should also be addressed.

Minor Corrections

- p5: The arrows in Figure 1-1 imply the elementary functions are integrated into the RISC-V soft processor, but that is not actually done.
- p17: Figure 2-2 seems not the correct reference here—it is both quite far away (on p26), and does not seem appropriate for this reference to it.
- p37/table 3.1 (and some other tables too): Can you make this fit without needing to go sideways?

- p82: Table 4.6 is identical to table 4.1, other than including your own results. I realize you are trying to present a narrative, now showing your design in comparison, but less repetition would be welcome.
- p90/Section5.4: You include "Both accelerators do not support IEEE-754 subnormal inputs; instead, those inputs are flushed to zeros. We treat results from the GNU math.h software library as a baseline." as a contribution. How exactly is that a contribution?
- p95: what is the "non-restoring method" used for in sqrt?
- p95: what is "Quakes's algorithm"?
- p95: what is the "Vedic algorithm"?
- p105/S5.10.2: Why is the lower bound of FMax reduced (from 106 to 104MHz) by removing exception handling? The "midway" version has an even larger reduction in FMax (table 5-5, p106). That seems counter-intuitive.
- Chapter 6 describes your RISC-V implementation. However, it is not always clear what aspects are mandated by the RISC-V specs and where you have made significant design decisions. You mention low-level register uses and opcode bit allocations/encodings, but I assume those are required by the spec rather than a feature of your design. Some summary-clarification of what design choices you made beyond adhering to specifications would be useful.
- p144: I am unsure what the "recursive vector addition" program is for. You list your benchmarks, and then state "When measuring the performance of our HLS-generated processors, they are all loaded with a recursive vector addition program shown below." How exactly is that code being used in your other tests?
- p146: Your pipelined processor does not handle data hazards. Does this affect correctness of the output then?
- p151: Using line-counts as a measure of complexity has many caveats, magnified even further if you want to compare different languagesm such as C with Verilog. Making strong statements based on that comparison is thus not possible, and you should minimally discuss some of these issues in relation to how meaningful a code-length comparison actually is.
- p153: If I recall correctly, most financial apps avoid floating point and use integer/fixed-point due to imprecision concerns.