

## HABILITATION THESIS REVIEWER'S REPORT

### Masaryk University

<b>Candidate</b>	RNDr. Jiří Filipovič, Ph.D.
<b>Habilitation thesis</b>	Software Performance Optimization in Scientific Computing
<b>Reviewer</b>	Doc. RNDr, Milan Češka, Ph.D.
<b>Reviewer's home unit, institution</b>	Faculty of Information Technology, Brno University of Technology

The research presented in this habilitation thesis falls into the area of high-performance computing with the focus on performance optimization in scientific computing. The overall topic of the habilitation is important and is growing in relevance as computational demands of many important applications exceed the capabilities of available hardware architectures.

The contributions of the candidate in this area are split into three main threads that are presented in three sections: Approximate modelling (Section 2), Application of GPU acceleration (Section 3), and Automatic code optimisation (Sections 4).

**Section 2** briefly summarizes two key results: i) CaverDock, a novel tool-supported method for computational analysis of ligand transports and ii) a high-performance computing environment for haptic rendering.

CaverDock showcases an excellent interdisciplinary research combining computational biology and high-performance computing. The related achievements have been presented in two top-ranked journals (D1/Q1) in the area of computational biology. It is evident that the candidate played the key role on the side of high-performance computing and significantly contributed to this exceptional result.

The second result has been also published in a top-ranked journal. Although the candidate significantly contributed to this result, it is related to quite old work (2009-2011, before obtaining Ph.D.) and thus I found it less relevant with respect to the scope and timeline of the habilitation thesis.

The exposition of these results is unfortunately very minimalistic. I miss related work that would help me to understand the context of these results. I would also appreciate more details on the computational/algorithmic aspects of this work. Furthermore, I found the title "approximate modelling" very misleading – I understand the motivation to unify these two results, but it would require more effort including broader related work.

**Section 3** briefly summarizes an application of a GPU acceleration in computational chemistry and image processing in microscopy. This section is also very short and provides only a very high-level overview of the results. I miss a more detailed landscape of the related work. Although this section provides some technical details, I found difficult to understand the key ideas – maybe a schematic illustration of the proposed parallelisation strategies would help here.

This section builds on two decent publications (a CORE B conference paper and a Q2/Q3 journal paper) but also refers to two debatable publications (a 1<sup>st</sup> edition workshop paper and a paper published in a very general MDPI journal) – please see below my comments related to the paper selection.

**Section 4** summarizes the results achieved in the area of automated optimisation methods for GPU programs including kernel fusion and autotuning methods. This section, especially the second part, provides a very nice summary of the achieved results including a broader context of this work as well as an appropriate description of the technical details.

The section builds on strong publications including two Q1 journals in area of high-performance computing. In my opinion, these results as well as their exposition clearly demonstrate the ability of the candidate to conduct successful and independent research in the highly competitive field.

**Selection of the papers in the collection:** The selection includes very strong papers (see namely [32,82, 29, 110]), to which the candidate significantly contributed, and thus it demonstrates high-quality of the candidate's research work. However, I have the following reservations:

- i) The candidate should provide information about the quality of the conferences and journals where the results have been published – this would significantly simplify the evaluation process.
- ii) The selected papers should showcase the top results relevant to the habilitation and thus including some of the papers is debatable: some papers are very old (see [75, 79]) and some publication venues are of a low quality (see [37,96, 98]). In these cases, the candidate should explain the impact and importance of these publications.

**Questions for the habilitation thesis defence:**

Q1) Can you please comment the key computational/algorithmic aspects that made CaverDock a unique tool for the analysis of ligand transports?

Q2) Can you please comment the parallelisation strategies used in the results described in Section 3? Are these strategies somehow novel with respect to state of the art?

Q3) Can you please comment the selection of the papers in the collection (see above)?

**Conclusion**

Despite of the aforementioned drawbacks that lower the quality of the thesis itself, I believe that the body of work presented in the collection of the papers is very good, demonstrating a high level of scientific quality and an important contribution to the field. Therefore, in my opinion, the habilitation thesis **fulfils** the standard requirements on habilitation theses in the field of computer science.

Date: March 14<sup>th</sup>, 2024

Signature: