FACHGUTACHTEN
PROJEKTNUMMER: P29119-N35

Reviewer A

1.) Scientific/scholarly quality (including innovative aspects and originality) with special attention to strengths and weaknesses:

The architecture of semiconductor devices is increasingly based on three-dimensional component layouts approaching fully three-dimensional layouts. Examples are recent processors as Intel's Haswell or IBM's Power8. As it is essential to understand the physical characteristics of a device during the design process and long before it is actually built, device simulation has become of high importance for the semoconductor industries. Three-dimensional device simulation is in its infancy, however, and therefore, the research proposed is of high relevance.

The present project proposal adresses simulation technologies with the intention to meet the requirements of three-dimensional layouts of the devices. It's ambitious aim is to perform the "first large-scale, highly accurate 3D semocinductor device simulations based on *deterministic solutions* of the stationary Boltzmann Transport Equations". This idea is highly original.

In order to achieve the goal, the proposal is addressing quite a few innovative problems and methods in quite different domains of physics, numerical mathematics and computer science:

- The Boltzmann Transport Equation (BTE);
- Truncated expansions in spherical harmonics (SHE);
- Memory saving data structures;
- Fast linear Krylov-solvers;
- Parallel ILU preconditioners;
- Algebraic multi-grid methods;
- Tensor-product structural representation of the higher-order SHE;
- GPUs and XEON Phi implementations.

While this domain-transcending philosophy appears as *the* strength of the project it also carries the risk that all elements of the approach need to comply with each other in order to become an effective machinery. This requires the smooth interplay between the physical scheme, its discretized model, the choice of the numerical solvers, perconditioners and multi-grid schemes and their successful implementation on novel processing elements like GPUs or other type of many-core processors. Nevertheless, the intended efficient interplay of all these various elements is what exerts the true fascination and potential of the proposal.

☐ excellent	☒ very good	☐ good	☐ average	☐ poor

2.) Approaches/methods and feasibility of the proposal with special attention to strengths and weaknesses:

The proposal presents a very impressive account on methods and innovative approaches to tackle the problem of three-dimensional device simulations on modern parallel architectures. As mentioned above, the risky part of the proposal lies in the proper interplay of these methods and implementations with control of systematic errors. It is evident that these potential problems cannot be fully known in advance and have to be clarified or solved in the course of the project. Special attention should be given to the following issues:

Charge transport in semiconductors can be described in a semi-classical manner by the BTE. While this apparently allows to create a computationally feasible model, one has to remember that the solution of the BTE is a compromise between a fully quantum mechanical approach and macroscopic classical models; therefore, the systematic errors of the approximation need to studied very carefully.

The dimensionality of the problem can be reduced from 6 to 4 by means of the truncated SHE. The proposal presents as innovation and advantage the automatized adaption of the truncation order for a given accuracy. On the one hand, the volume and energy dependence of the truncation order at given accuracy might become increasingly demanding on memory and compute resources – this remains a risk –, on the other hand, the required accuracy itself has to be assessed by means of a proper validation scheme.

The memory benefits of the data structure introduced to avoid non-relevant kinetic energy entries of the system matrix vs. the computational expenses for extensive index computations need to be clarified in a quantitative manner, i.e. in terms of real full costs of using the parallel systems. The same question should be answered for the effectivity using the tensor-product structure of the higher-order SHE for memory savings vs. compute efforts.

It should be clarified in an early phase in how far algebraic multi-grid technologies might really become superior to fast linear Krylov-solvers with efficient parallel preconditioners. Algebraic multi-grid methods exploiting explicit knowledge of the physical properties of the system are very difficult research projects *sui generis*.

Furthermore, it might become difficult to achieve truly large-scale and fast implementations on GPUs and XEON Phi architectures. Again this remains a risk.

All in all, the approach presented in the proposal seems well structured and carefully reflected. While some risk remains as is always the case in interesting research, I consider the project to be feasible,.

☒ excellent	☐ very good	☐ good	☐ average	☐ poor

3.) Qualifications of the researchers involved (based on their academic age) with special attention to strengths and weaknesses:

Dr. Rupp can present a quite impressive vita, with educational exams passed with distinction in a row as well as awards ranging back to his high school graduation as winner of both the Styrian Mathematics and Physics Olympiad. In this respect, Dr. Rupp can present a reputable academic history.

Dr. Rupp counts 63 publications since 2010. 22 are cited in the Web of Science, 8 of these are journal papers. While the paper record shows that Dr. Rupp has worked steadily and has presented his results in international workshops and symposia, his citation count is – with 24 citations excluding self-citations – a bit low for the number of publications and his academic age.

In this light, one can assess Dr. Rupp good academic qualifications by international standards, still I consider his potential as higher than suggested by the citation analysis.

☐ excellent	☐ very good	☒ good	☐ average	☐ poor

4. Ethical issues:

Does the project give rise to any ethical issues?	☐ yes	☒ no	☐ do not know

FACHGUTACHTEN
PROJEKTNUMMER: P29119-N35

5. Overall evaluation with regard to key strengths and weaknesses and final funding recommendation:

Overall, the proposal deserves the rating Very Good, given its excellency in the methods – the key strength of the proposal –, its very good scientific quality and the good academic standing of the PI. I recommend full funding of the proposal with high priority.

☐	Excellent - funding with highest priority
☒	Very Good - funding with high priority
☐	Good - resubmission with some revisions
☐	Average - resubmission with major revisions
☐	Poor - rejection

Please note that the FWF places high demands on the quality of the projects it funds and thus predominantly supports projects rated as 'very good' or 'excellent'.

FACHGUTACHTEN
PROJEKTNUMMER: P29119-N35

Reviewer B

1.) Scientific/scholarly quality (including innovative aspects and originality) with special attention to strengths and weaknesses:

This research proposal aims at providing a significant contribution to the field of numerical methods for the simulation of the next generation of **fully three-dimensional** semiconductor devices. Instead of using standard Monte Carlo approaches, or over-simplified drift-diffusion approximations, this project will be based on the massively parallel solution of the more accurate and computationally more efficient Boltzmann Transport Equation (BTE) on shared and distributed memory supercomputers, for which the model accuracy is expected to improve linearly with computational effort, while the accuracy of Monte Carlo simulations improves only with the square root of the invested computational effort.

However, the numerical solution of the BTE is extremely challenging, due to the **nonlinearity** and due to the **high dimensionality** of the problem. Particular problems lie in the solution of the nonlinear problem via Newton-type techniques and the finding of appropriate initial guesses, as well as in the efficient solution of the resulting sequence of linear problems and the design of preconditioners that are suitable for a massively parallel implementation. The present project aims at adressing all these issues in a very original way.

First, the applicant proposes to solve for the logarithm of the probability distribution function *f*, instead of solving directly for *f*, in order to make better use of the resolution properties of the used computational grid. Also, the number of dimensions of the problem is going to be reduced at the aid of an expansion in terms of spherical harmonics. Second, the applicant proposes to further reduce the necessary memory requirements (due to the high dimensionality of the problem), by introducing a special datastructure that allows to store only blocks of grid cells that are located around kinetic energies that are relevant for the simulation and by exploiting the tensor-product structure of the system matrix generated by the use of spherical harmonics. Concerning the search for an improved initial guess of the nonlinear Newton solver, the present proposal contains a **very innovative** idea based on a new and **highly original multi-model approach**. There, the initial guess of the nonlinear Newton solver will be provided by the numerical solution of an appropriate - but significantly simplified - drift-diffusion model, which can be derived from the underlying BTE under several assumptions (which need not to be all valid, since the drift-diffusion model is only expected to provide an initial guess for the nonlinear solver of the governing BTE model). This idea can be clearly seen as a **very original contribution** of this project. Finally, the project aims at providing **new efficient and massively parallel preconditioners** that are needed for the iterative solution of the linear systems arising in each Newton iteration. Such efficient preconditioners are difficult to design in a massively parallel environment, but the present project contains several innovative ideas, also based on previous work of the applicant. It has to be stressed that the applicant has already consolidated experience with massively parallel linear algebra packages like PETSc, which makes also this part of the proposal very strong. It has to be underlined that this part of the project alone will be a major contribution to scientific computing and numerical linear algebra.

There are no significant weaknesses in the overall structure of the project.

☒ excellent	☐ very good	☐ good	☐ average	☐ poor

FACHGUTACHTEN
PROJEKTNUMMER: P29119-N35

2.) Approaches/methods and feasibility of the proposal with special attention to strengths and weaknesses:

The proposal is well written and well planned. It is very ambitious, but given the track record of the PI and of his senior collaborators (Prof. Ansgar Jüngel and Prof. Tibor Grasser), it is feasible with the requested resources. The work plan is sound and the proposed methodology is valid. Several of the proposed work packages are independent of each other, which makes a success more likely. For some critical tasks, alternative fallback strategies have been outlined. The host institution (TU Wien) has all the necessary infrastructure (including the supercomputer Vienna Scientific Cluster) that is needed to carry out this research project successfully. The project promotes **Open Science**, i.e. open access to publications and software. All the above-mentioned aspects are a clear **strength** of this proposal.

In my view, there is only one minor **weakness** concerning the planned adaptive mesh that has been sketched on page 13 of the proposal. The applicant proposes to store only grid cells for relevant kinetic energies, in order to save memory. However, even the adapted grid shown in the right panel of Fig. 4 of the proposal, is only a quasi-uniform equidistant (?) grid. Here, the use of **hierarchic** adaptive meshes with multiple levels (AMR) might be more appropriate and lead to even more memory savings, especially considering the high dimensionality of the problem. It could also have been further combined with local p-adaptation. In other words, according to the referee, a more sophisticated hp-adaptive scheme might be even more appropriate than the simple strategy proposed in Fig. 4.

In the proposal on page 10, the applicant mentions other approaches like the Discontinuous Galerkin method [62] or multigroup equations [63], but he does not have a track record on these methods, so it is not clear how the new approaches of this project that are based on spherical harmonics can be directly transferred by the PI or his group to DG schemes, apart from the rather general techniques concerning preconditioners for iterative linear solvers or the innovative initial guess for the Newton method.

☐ excellent	☒ very good	☐ good	☐ average	☐ poor

3.) Qualifications of the researchers involved (based on their academic age) with special attention to strengths and weaknesses:

All involved researchers, i.e. the PI and his senior collaborators, have an **excellent** academic track record based on their academic age and also compared to **the highest international standards**. All researchers are highly productive concerning international publications in leading journals and conferences in the field: **48** publications of the PI **Dr. Karl Rupp**, **153** publications of Prof. **Ansgar Jüngel** and **364** publications of **Prof. Tibor Grasser** are currently listed in Scopus.

There are definitely **no weaknesses** concerning the qualification and international reputation of the PI and of his senior collaborators.

☒ excellent	☐ very good	☐ good	☐ average	☐ poor

4. Ethical issues:

Does the project give rise to any ethical issues?	☐ yes	☒ no	☐ do not know

5. Overall evaluation with regard to key strengths and weaknesses and final funding recommendation:

Overall an excellent project that is highly competitive also according to international standards. If successful, it will have a very important international impact in the field of scientific computing, high performance computing and numerical methods for PDE in high dimensions. The possible practical applications concerning numerical 3D semi-conductor design are of great importance for big international companies like Intel, Nvidia or AMD. I therefore **strongly recommend** this project for funding.

☒	Excellent - funding with highest priority
☐	Very Good - funding with high priority
☐	Good - resubmission with some revisions
☐	Average - resubmission with major revisions
☐	Poor - rejection

Please note that the FWF places high demands on the quality of the projects it funds and thus predominantly supports projects rated as 'very good' or 'excellent'.