

Posudek diplomové práce

Matematicko-fyzikální fakulta Univerzity Karlovy

Thesis author Alisher Kenzhebayev

Thesis title Vectorized traversal of sparse volumes for GPU path tracing

Submitted 2024

Program Computer Science **Specialization** Computer Graphics and Games Development

Review author Tomáš Iser, Ph.D. **Role** supervisor

Position KSVI

Review text:

This review was written by Tomáš Iser, Ph.D., the thesis supervisor, in collaboration and with the agreement of Tobias Rittig, Ph.D., the thesis advisor.

Writing style The thesis is written in good English and consists of all the expected and important parts, with a good split into the individual chapters and sections.

However, the text uses sophisticated wording and complicated sentence structures that are rather colloquial and feel more like storytelling, unusual for an academic publication. Consider this example on p. 36: *“It wasn’t until then that mainly, the unit tests fired off that there was a mistake in the read access for an invalid memory offset. Why this occurred was then evident.”* The argumentation logic, that is how arguments are presented as building on each other, seems very prolonged and hard to follow.

The problem of sentences being unnecessarily long and complicated, hence difficult to read, is common throughout the text. Often, whole phrases could be removed without changing the original meaning. Consider this example on p. 11: *“However, single file cases aside, even for samples of small to moderate size that have temporal variance data, this memory issue again becomes important since they have to record the file in the animation frames, which is usually the case for animated or simulated volumes.”* Without losing its meaning, it could be shortened by half to: *“Even small volumes can have a large memory footprint in case of animations because individual keyframes have to be stored.”*

The complicated thinking is also emphasized in how things are named or referred to. A simpler nomenclature would have helped the understanding for first-time readers greatly.

Content It seems that the Introduction and Background chapters are somewhat repetitive with respect to each other. However, the background covers a wide array of topics, summarizes them in

an adequate length, and thus shows that the student has understood the materials he was working with. The section provides some references to previous or related work, but there could have been more research done in this context. In Figure 1.1, the caption refers to “two internal nodes” and “both internal nodes”, but it seems that 9 internal nodes are visible in the figure, not two. Could the author clarify that?

The Methodology chapter starts with a great overview Figure 2.1, which is very helpful in understanding the methodology, and it is well referred to throughout the text. Unfortunately, the chapter seems to be filled with text describing what follows in the upcoming sections or what was in the previous sections, but these “meta descriptions” are very long compared to actual new content. While referring to other parts of the thesis is very important, in some sections and subsections of the chapter, most of the text is just that and does not bring anything new.

The Results are insightful and there are many useful test scenarios. The presentation itself, e.g., consistent naming of cases, layout, and ordering, could be improved. Figure 4.4 could have been included more in the discussions. The image comparison is unfortunately not very meaningful, an MSE that is larger than 0.001 is significant. There also seems to be an alignment issue, and it is not clear whether that is a problem of the image, or of the implementation.

Implementation The implementation itself should be praised mainly because it worked with a very advanced C++ code of the OpenVDB and Mitsuba 3 projects. Especially Mitsuba 3 is based on Just-In-Time (JIT) compilation using LLVM/CUDA/Optix, and doing any work with such code requires a very careful approach. The thesis consists of many low-level unit tests, qualitative tests, and performance tests to ensure the implementation is correct at all levels. My main criticism is that the implementation only works for one specific data type (monochromatic float grids) and the thesis was not extended to other, more useful types.

Work with the student Overall the work with the student was focused and consisted of regular weekly meetings on which the student presented new progress, and technical and debugging questions and issues that were encountered. The student was cooperative and motivated to solve the very technical problem of this thesis, and proved that they are capable of working with advanced C++ projects.

I recommend the thesis for defense.

I suggest to not consider the thesis for the annual award.

29 August 2024, Tomáš Iser, Ph.D.