

# Posudek diplomové práce

Matematicko-fyzikální fakulta Univerzity Karlovy

**Autor práce** Federico Forti  
**Název práce** Efficient GPU Path Tracing in Solid Volumetric Media  
**Rok odevzdání** 2017  
**Studijní program** Informatika      **Studijní obor** Počítačová grafika a vývoj počítačových her

**Autor posudku** Oskar Elek      **Role** Vedoucí  
**Pracoviště** KSVI

## Text posudku:

\*\*\* Summary \*\*\*

This thesis investigates an efficient implementation of the path tracing (PT) rendering algorithm on graphics processing units (GPU) - specifically, the volumetric variant of the algorithm.

PT is a de-facto standard for conducting physically-based light transport simulations, both in research and in production, so its optimal implementation can be widely beneficial. Given the relative simplicity of the algorithm, it seems suitable for being implemented on massively parallel computing platforms like GPUs. The difficulty of that however lies in the specific execution flow and random memory access patterns that PT tends to have. While notable effort has been invested into optimizing the surface-oriented PT, its volumetric counterpart has received little attention so far.

\*\*\* Exposition and clarity \*\*\*

The text doesn't provide a particularly enjoyable read - the English is basic and contains frequent mistakes. It is apparent (and i know for a fact) that it has been written in haste. However, it still reads quite fluently, and it's usually not difficult to follow the flow of arguments. The high-level organization also makes sense. In that regard i see it as a one-off work report rather than a thesis, which is fine and as such it does function.

The first part of the text presents a lot of theory. It feels unfinished in many places, leaving the impression that the reader already needs to know most, otherwise they would get lost. This is not very useful - the text would work better if the theory was skeletal, with frequent references to standard literature, and more focus was invested into the latter chapters, in particular the evaluation part. That all said, i do appreciate the effort the student put into mapping the area - it certainly does document the knowledge he gained during the work on the thesis.

\*\*\* Technical soundness \*\*\*

On the conceptual level, the investigation is quite sensible. A reasonable number of optimization techniques is evaluated, and the results (although sometimes tough to understand due to confusing presentation) seem to generally match the conclusions made by the student, and importantly, the expectations from the configurations.

The validation of the algorithm against a reference is completely absent. However, i know that the student has successfully validated his implementation, so it's not quite clear to me

(unless it's simply due to time shortage) why it hasn't been included in the text.

The implementation is very solid. It's usually quite difficult to design a modular and extensible GPU code and managing that (mainly thanks to the well designed encapsulation, i believe) is a success. Its progressive nature is also a small plus.

\*\*\* Fulfillment of goals \*\*\*

This is a mixed bag. On one hand, the evaluation is in places confusing and/or inconsistent, so it seems difficult to derive conclusions from the presented data. In general, the main conclusion at the end is very basic, and merely confirms the expectations. The resolution of the datasets used in the tests is also very limited (probably due to the lack of better available hardware), so it's not clear how the algorithm scales to datasets that would be about an order of magnitude bigger, yet still fit into the device memory on more advanced/newer architectures.

That said, in the basic sense, i think the thesis reaches the originally posed goals. It indeed proves that the GPU implementation can be significantly faster than an established CPU one, and evaluates a multitude of optimization techniques to achieve that.

\*\*\* Conclusion \*\*\*

To conclude, i believe the thesis should be defended. The presentation and used methodology is wonky in places, but overall i believe the student has proven he's able to conduct independent work, do his own literature search and implement it. The questionable evaluation can be ascribed to the student's lack of exposure to academic work, but since as far as i know he doesn't plan to continue with academic research career, this is acceptable, from my perspective.

Questions i would like to discuss on the defense (in that order of priority):

- Better explain the performance comparison to Mitsuba. Why are the presented numbers so strange and inconsistent? Or in other words, the presented numbers should be justified or re-measured.
- What are actually the properties that make volumetric PT difficult for a GPU implementation, in detail. Do author's own measurements support this?
- Have you tried any code-level optimization, mainly for reducing the number of registers?

**Práci doporučuji k obhajobě.**

**Práci nenavrhují na zvláštní ocenění.**

*Pokud práci navrhuje na zvláštní ocenění (cena děkana apod.), prosím uveďte zde stručné zdůvodnění (vzniklé publikace, významnost tématu, inovativnost práce apod.).*

-

**Datum** 22. January 2018

**Podpis**