Assessment of the diploma thesis entitled

„Acceleration of Ray-Casting for CSG scenes“

submitted by: Petr zajíček

This thesis deals with the problem of efficient ray-scene intersection for the specific case of scenes modeled using the Constructive Solid Geometry (CSG) technique. A novel data structure, the Operation KD-tree, is proposed for this purpose, along with a corresponding traversal algorithm. The advantage of the data structure over other possible approached (based on CSG and triangle meshes) in terms of rendering speed is evaluated on a few simple example scenes.

In justifying the usefulness of the CSG representation, the author limits himself to arguments about the scenes size (CSG leads to more concise representations than triangle meshes) and the corresponding benefits in memory and runtime efficiency. However, I find that the real importance of the ability to directly render CSG models lies especially in the increased accuracy and robustness in engineering applications, where CSG models are commonplace. With a direct CSG rendering approach, the need to transform a CSG tree into a boundary representation - an inherently unstable operation - is altogether avoided. For that reason, and in spite of the common trends in "real-time" ray tracing research, I find the problem addressed in the thesis of great importance.

Regarding the technical solution itself I find the idea of the Operation KD-tree quite interesting. The thesis does provide convincing arguments for the design of this data structure, nonetheless, the OKD-tree is clearly unable to efficiently cope with what I judge the main problem in efficient ray casting of CSG models, that is the high degree of overlap among the individual primitives. Previous research in ray casting has addressed this problem either by breaking primitives into smaller pieces or by storing references to "large" primitives in tree nodes closer to the root. I find it disappointing that the author did not explore this line of research. That said, I do consider the OKD-tree a valid and sound approach.

I found the evaluation quite limited in its scope. Indeed, the number and type of models tested is not a representative sample of real-world (engineering) applications of the proposed method. For that reason, the measured results cannot be considered a convincing evidence for the benefit offered by the proposed OKD-tree.

The thesis text itself is very well structured, the solved problem with all its nuances is well explained and the solution chosen is well motivated. The text does contain quite a number of spelling and grammar errors but I do not consider them a significant flaw of the work.

In summary, the submitted thesis reports on a solid piece of work and proves the candidate's ability carry out research and implementation work. I do recommend the thesis for acceptance.

Praha, 20.4. 2012

Jaroslav Křivánek