Blender based Rendering-as-a-Service Platform for High Performance Computing Clusters

IT4Innovations national supercomputing center

Milan Jaros, Petr Strakos, Lubomir Riha

milan.jaros@vsb.cz, petr.strakos@vsb.cz, lubomir.riha@vsb.cz

Introduction

- Open-source rendering solutions that are capable of utilizing a large number of computational resources are rare.
- Usually made for render farms but not directly for High Performance Computing (HPC) clusters.
- One of such solutions is Flamenco [1] that is a render manager based on Blender 3D creation suite [2]. It supports offline rendering only.
- We are currently developing a rendering-as-a-service platform that efficiently utilizes HPC resources in the supercomputing centres.
- HPC cluster can be equipped with classical CPU nodes or accelerated nodes or their combination.
- We decided to build our solution on Blender since it gathers on popularity and offers realistic renderers and variety of extension possibilities in terms of plug-ins.

Our approach

- Our platform is based on Blender renderers and upgrades them with HPC technologies such as:
  - MPI for distributed rendering.
  - OpenMP for parallelization for multi-core CPUs.
  - Support for modern architectures such as Intel Xeon Phi in form of accelerator or stand-alone processor which include extended support for wide SIMD units (up to 512 bits).
  - Integration of EMBREE library [3] into Blender Cycles, the physically based production renderer.
  - Integration of OpenSWR library into Blender Eevee, the physically based rendering engine for real-time visualization.
- In this way we can offer not only standard offline but also interactive rendering mode which relies on fast HPC interconnecting networks.

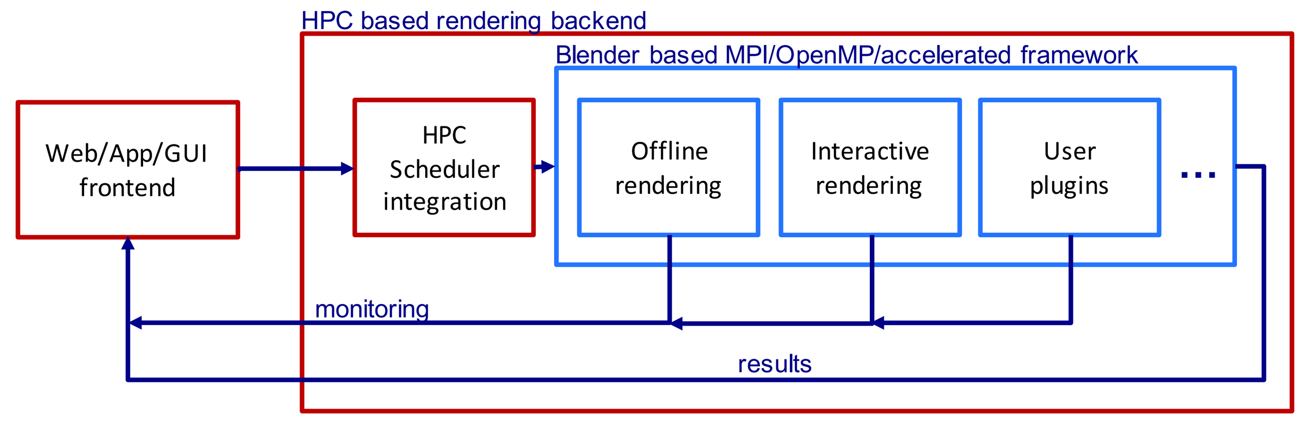

Figure 1: Rendering-as-a-service platform using HPC and Blender (already developed parts are highlighted by blue, the red parts are under development).

Offline & Interactive rendering scalability

- Scalability performance in offline rendering mode using Cycles renderer [4]:
  - Scene duplicated on all nodes.
  - Support for load balancing.
  - OMP24 stands for 2×Haswell and Symmetric for 2×Haswell + 2×KNC.

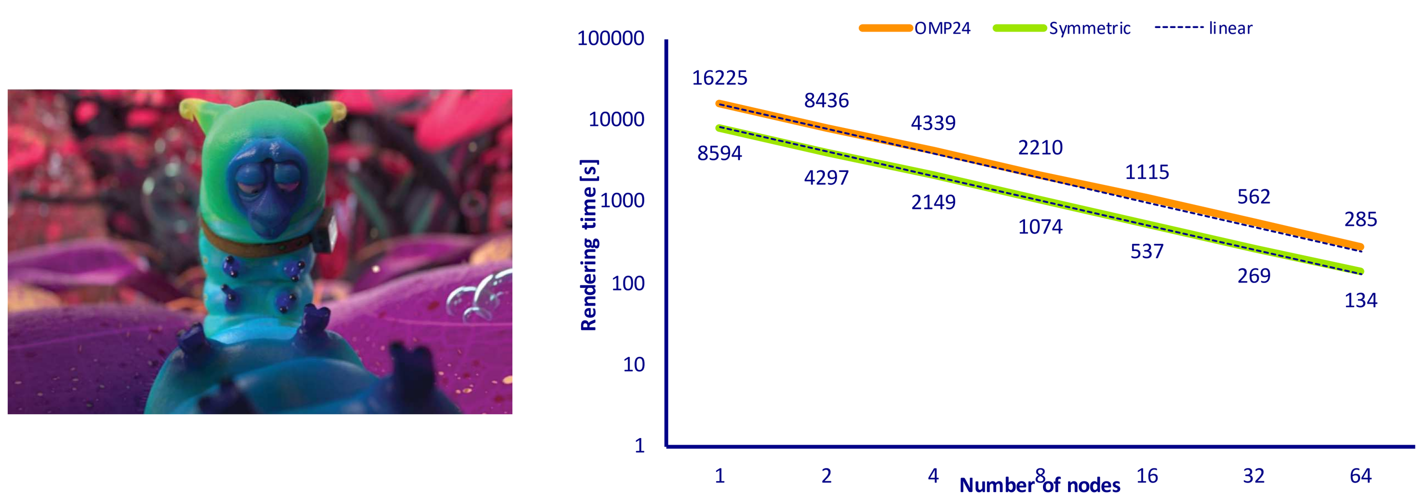

Figure 2: Rendering of Cosmos Laundromat scene (by Blender Institute), Scene example (left), Scaling performance (right).

- Scalability performance in interactive rendering mode using Cycles renderer [4]:
  - OMP24 stands for 2×Haswell and Symmetric for 2×Haswell + 2×KNC.

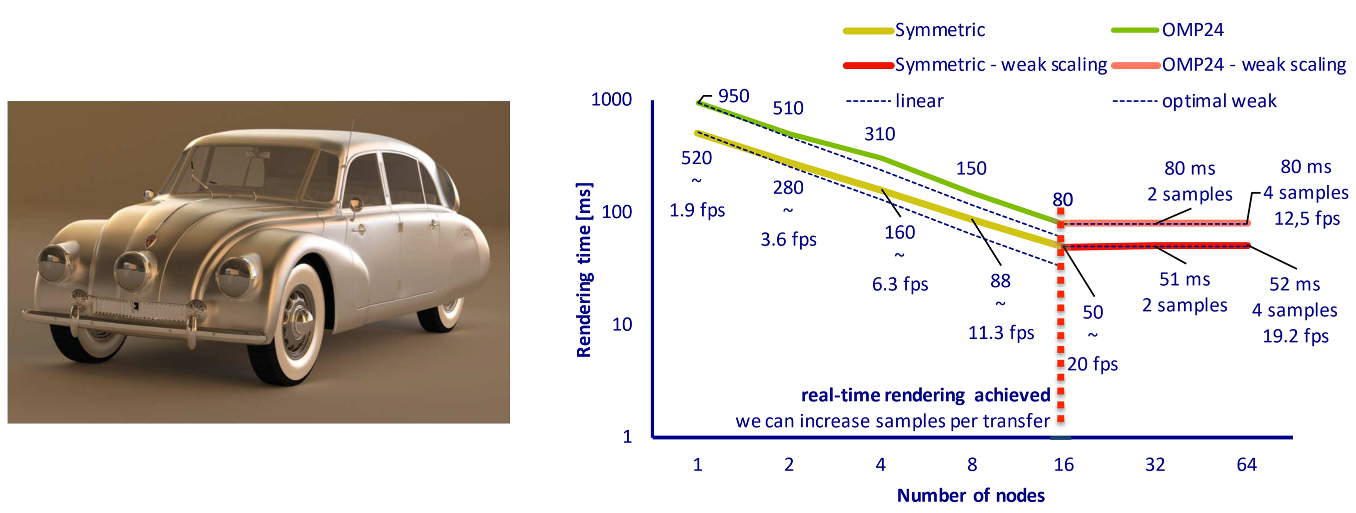

Figure 3: Rendering of Tatra T87 scene, Scene example (left), Scaling performance (right).

Architecture performance

- Performance evaluation in terms of rendering time in Cycles renderer offline mode for different processor architectures.
- General Blender Cycles benchmark scenes have been used.
- Evaluated architectures:
  - NVIDIA GeForce GTX TITAN X.
  - Intel Xeon Phi 7120P (KNC).
  - Intel Xeon E5-2680v3 (Haswell).
  - Intel Xeon Phi 7250 (KNL).
  - Intel Xeon 8160 (SKL).
- The results show promising performance on:
  - new Intel Xeon Phi 7250 (68 CPU cores) with 512 bit wide SIMD units.
  - Intel Xeon 8160 (24 CPU cores) architectures with 512 bit wide SIMD units.
- Parallel efficiency in our implementation has been reached.

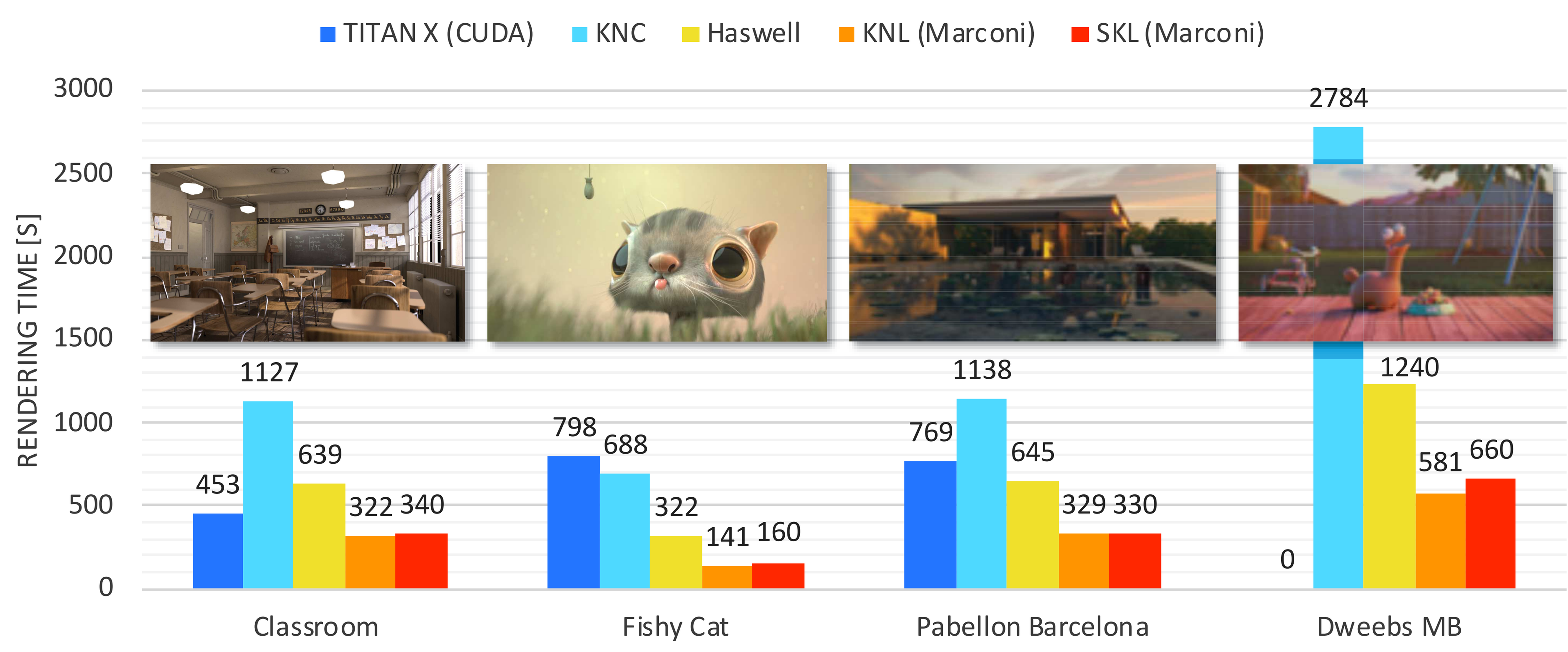

Figure 4: Performance evaluation - Salomon IT4Innovations: Haswell, KNC, TITAN X. Marconi CINECA: KNL, SKL

Rendering examples

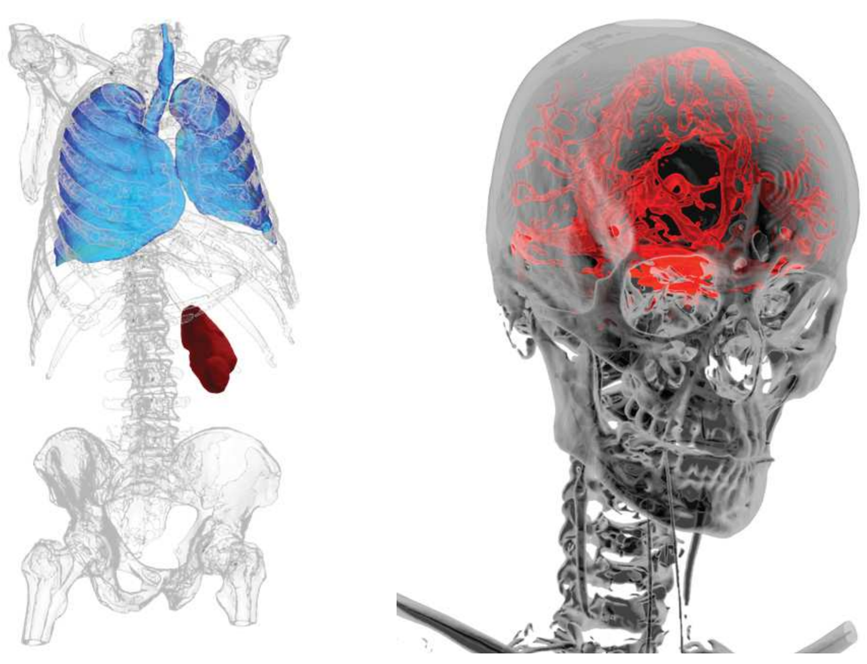

(a) Medical imaging

(b) Engineering

(c) Animation (Agent327 by Blender Institute)

Figure 5: Various use cases rendered in Blender by Rendering-as-a-Service

Acknowledgement

This work was supported by The Ministry of Education, Youth and Sports from the Large Infrastructures for Research, Experimental Development and Innovations project "IT4Innovations National Supercomputing Center - LM2015070".

References

[1] *Flamenco.* (5 2018). https://www.flamenco.io/
[2] *blender.org - Home of the Blender project - Free and Open 3D Creation Software.* (5 2018). https://www.blender.org/
[3] *Embree.* (4 2018). https://embree.github.io/
[4] Jaros, M., Riha, L., Karasek, T., Strakos, P. and Krpelik, D. *Rendering in Blender Cycles using MPI and Intel Xeon Phi.* (2017). ACM International Conference Proceeding Series.