João Madeiras Pereira
jap@inesc-id.pt

Importance Sampling of Many dynamic lights for Real Time Ray Tracing

João Madeiras Pereira
jap@inesc-id.pt

Motivation:

Monte Carlo ray tracing offers the capability of rendering scenes with large numbers of area light sources—lights can be sampled stochastically and shadowing can be accounted for by tracing rays, rather than using shadow maps or other rasterization based techniques that do not scale to many lights or work well with area lights. Current GPUs only afford the capability of tracing a few rays per pixel at real-time frame rates, making it necessary to focus sampling on important light sources. While state-of-the-art algorithms for offline rendering build hierarchical data structures over the light sources that enable sampling them according to their importance, they lack efficient support for dynamic scenes

Objective:

This thesis aims to implement the algorithm described in [1] which is capable to maintain the hierarchical light sampling data structures targeting real-time ray tracing. The results should be based on a real-time path tracer already implemented in Vulkan.

[1] https://research.nvidia.com/sites/default/files/pubs/2019-07_Dynamic-Many-Light Sampling//MPC19.pdf

Description:

The algorithm is based on a two-level BVH hierarchy that reduces the cost of partial hierarchy updates. Performance is further improved by updating lower-level BVHs via refitting, maintaining their original topology.

Vulkan-based Interactive Global Illumination development in Android

João Madeiras Pereira
jap@inesc-id.pt

Motivation:

The rapid increase of performance on low-power graphics processors has made a big impact. More advanced computer graphics algorithms are now possible on a new range of mobile devices. Indirect illumination is an important element for realistic image synthesis, but its computation is expensive and highly dependent on the complexity of the scene and of the BRDF of the involved surfaces. While off-line computation and pre-baking can be acceptable for some cases, many applications (games, simulators, etc.) require real-time or interactive approaches to evaluate indirect illumination.

Objective:

To compute, in an Android mobile device, real-time indirect lighting in a dynamic complex scene by developing a mobile version of the Image Space Photon Mapping (ISPM) algorithm: https://casual-effects.com/research/McGuire2009Photon/index.html

Description:

To implement, by using the low level Vulkan API, the ISPM algorithm which preserves the accurate and physically-based nature of photon mapping, supports arbitrary BSDFs, and captures both high- and low-frequency illumination effects such as caustics and diffuse color interreflection.

Interactive GPU-based Cloth Simulation

João Madeiras Pereira
jap@inesc-id.pt

Joaquim Jorge
jorgej@acm.org

Objectives:
This work aims to develop a parallel algorithm for cloth simulation that exploits GPUs for fast computation and the handling of very high resolution meshes. The research can be based on the work from https://sa2020.siggraph.org/en/attend/technical-papers/session_slot/86/13 which describes parallel algorithms for sparse matrix-vector multiplication (SpMV) and for dynamic matrix assembly on a multi-GPU workstation.
As outcome, it should reliably handle the collisions and generate vivid wrinkles and folds at 2-5 fps.

Requisites: Computer Graphics courses

External Entity: ALICE

Kinect Noise Removal with Deep Learning

João Madeiras Pereira jap@inesc-id.pt
Francisco Fernandes francisco.fernandes@ist.utl.pt

Motivation:

The introduction of Microsoft's Kinect system [1], currently on its 3rd generation, enabled 3D depth-sensing cameras to be available to the general public. These affordable RGBD sensors have a wide range of applications that go beyond their evident use in games, expanding to many other research domains. But one of its major problems is the "noise" in the depth sensor image, which only gets worse as the person or object's distance from the camera increases.This prevents the correct detection of objects' boundaries, as foreground and background are often mixed up together, due to incorrect depth values. Recently, NVIDIA researchers introduced a Deep Learning technique named Noise2Noise [2], which can very accurately remove noise from 2D images.Its most groundbreaking feature is that it accomplishes this by being trained entirely with noisy images only, i.e. without having previously seen any clean versions whatsoever. Many exciting possibilities emerge, by combining this technique with Kinect data.

Objective:

We intend to apply the Noise2Noise technique to Kinect's depth data, in order to obtain noise-free depth images.

Description:

Modify or re-implement the Noise2Noise technique to work with Kinect depth images;
Create a proof-of-concept app in Unity; Possibly integrate the approach in other frameworks;
Test and evaluate the approach in different scenarios: objects at different distances, different background conditions, etc.;Compare to other depth noise-removal approaches;
Explore other ways to improve the algorithm, e.g. using temporal coherency.

Requisites: "Machine Learning" course, or similar knowledge; Interest for working with Kinect.

Force Field calculation for Biomolecules

João Madeiras Pereira
jap@inesc-id.pt

Francisco Fernandes
francisco.fernandes@ist.utl.pt

Motivation:

At the molecular level, almost every biological function, process and interaction is guided by attraction and repulsion forces between the atoms. Think of it as a giant pool of spherical magnets with different sizes, forces and charges. A Force Field comprises not only these atoms' energies, but also their type, radius, bonds, distances, angles, etc. But to determine the direction, speed and overall movement of each atom, we need to calculate the contribution of the forces of every other atom over this one. If we consider a system composed of thousands of atoms, this problem, as well as the time necessary to computationally calculate it, grows immensely. Therefore, more efficient ways to solve it are required.

Objective:

This work aims to use space-partitioning techniques and data structures, such as Octrees or Bounding Volume Hierarchies, to tackle this problem in acceptable running times. We can divide the 3D space into volumes of different sizes and store in each one the average of the forces of all the atoms inside each one. Then, to calculate the influence of the Force Field over each atom, we can simply query the volumes where he is included into, instead of querying all the other atoms.

Description:

You will have to study and test different Computer Graphics 3D Acceleration Structures and their performance; implement a Particle-System-like concept tool to calculate the forces over each atom and to simulate their movements across time; come up with improvements and optimizations to the Force Field calculation method; acceleration via GPUs is not required, but welcome.

Fast Rendering of Protein Representations

João Madeiras Pereira
jap@inesc-id.pt

Francisco Fernandes
francisco.fernandes@ist.utl.pt

Motivation:

Many molecular biology software tools require the visualization of protein structures using different representations. When displaying proteins with hundreds of thousands of atoms, their visualization becomes incredibly slow due to the complexity of the mesh and the high number of triangles involved.

Objective:

The proposed work will focus on developing fast and clever ways to display the proteins by relying on non-mesh-based approaches similar to splat rendering or the ZSpheres technique used in ZBrush.

Description:

The visualizations should incorporate photorealistic effects. By using the latest RTX platform provided by Nvidia, it should be possible to navigate through these visualizations at interactive framerates. Technologies like OpTix and Falcor from Nvidia, as well as DXR from Microsoft or Vulkan, will be potential candidates for this thesis' work development.

A Virtual Reality platform for Scientific Visualization of Plasma Physics Simulation

João Madeiras Pereira jap@inesc-id.pt
Alfredo Ferreira alfredo.ferreira@tecnico.ulisboa.pt

Motivation:

At the Group of Lasers and Plasmas (GoLP) of Instituto de Plasmas e Fusão Nuclear at Instituto Superior Técnico, researchers use Particle-in-Cell (PIC) method-based simulations to model plasma with little or no approximations. PIC simulations solve the smallest scales of the plasma, so simulating large systems (where dynamical scales are much larger than the smallest scales) is computationally demanding. For this reason, these simulations are massively parallel, and run on some of the largest supercomputers available worldwide. These simulations produce multidimensional grid and particle data sets that generally exceed several gigabytes per file, making their analysis and visualization non-trivial.

Objective:

This thesis aims to produce a platform for scientific visualization of results of the mentioned numerical simulations resorting to virtual reality (VR) experiences. VR glasses and haptic devices will be available for the development.

Description:

Navigating through these interactive visualizations in VR should be an immersive experience, and provide a completely new way of visualizing and interpreting scientific results. Different VR prototypes will be developed and will show numerical simulations of plasmas in the vicinity of compact astrophysical objects such as neutron stars and black holes, as well as of plasmas produced with intense lasers in the laboratory and used to create compact particle accelerators.

Análise de estruturas em ambientes imersivos

João Madeiras Pereira
jap@inesc-id.pt

Alfredo Ferreira
alfredo.ferreira@tecnico.ulisboa.pt

Enquadramento:

A análise de estruturas é uma disciplina fundamental no processo de conceção e dimensionamento dos mais variados tipos de construções, desde pontes a barragens ou edifícios de grande altura. A avaliação da segurança de estruturas é tradicionalmente feita recorrendo a aplicações do tipo desktop especializadas. No entanto, as tecnologias de realidade virtual podem trazer novas possibilidades. As suas capacidades imersivas aliadas ao uso de elementos gráficos ricos e relevantes podem representar vantagens tangíveis.

Objectivos:

Neste trabalho de mestrado propõe-se o desenvolvimento da componente gráfica de uma ferramenta de análise de estruturas que tire partido de ambientes imersivos. Pretende-se analisar as vantagens e limitações e contribuir para uma melhor compreensão das particularidades da aplicação de óculos de realidade virtual neste domínio.

Ray Tracing Explorer in VR

João Madeiras Pereira
jap@inesc-id.pt

Alfredo Ferreira
alfredo.ferreira@tecnico.ulisboa.pt

Enquadramento:

O algoritmo de ray-tracing é um pilar fundamental da Computação Gráfica. No entanto, a sua explicação é quase exclusivamente realizada recorrendo a diagramas em 2D, apesar de este algoritmo ser aplicado em 3D.

Objectivos:

Neste trabalho de mestrado propomos o desenvolvimento de uma ferramenta de visualização do funcionamento do algoritmo de ray-tracing interativa. Recorrendo à realidade virtual esta ferramenta permitirá estudar e explorar o funcionamento deste algoritmo de uma forma mais simples e natural.

Descrição:

Através do uso de óculos de realidade virtual (HTC Vive) e respectivos controladores, irá ser criado um ambiente imersivo que permite ao utilizador uma forma natural de visualizer o algoritmo em funcionamento como se estivesse no interior do espaço a ser desenhado. Este poderá ver os raios a serem disparados, as intersecções a acontecerem, os novos raios a serem criados e os pixeís a serem coloridos. Tudo isto acontecerá a uma velocidade que o utilizador controla. Este pode ainda navegar livremente na cena e consultar painéis com informação relevante.

Observation: Tese a ser atribuída ao aluno nº 89492, Lídia Custódio

Rigid Body Simulation with Extended Position Based Dynamics

João Madeiras Pereira

jap@inesc-id.pt

Motivation:

Traditional rigid body simulators linearize constraints because they operate on the velocity level or solve the equations of motion implicitly thereby freezing the constraint directions for multiple iterations. A state-of-the-art method [1] always works with the most recent constraint directions. This allows to trace high speed motion of objects colliding against curved geometry, to reduce the number of constraints, to increase the robustness of the simulation, and to simplify the formulation of the solver.

Objective:

This thesis aims to implement a fully fledged rigid body solver that handles contacts, a variety of joint types and the interaction with soft objects. In particular, the Extended Position Based Dynamics (XPBD) algorithm, described in [1], will be implemented in order to resolve accurately small temporal and spatial detail .

[1] https://matthias-research.github.io/pages/publications/PBDBodies.pdf

[2] https://www.youtube.com/watch?v=zzy6u1z_l9A&t=0s

Description:

The student will develop a XPBD-based rigid body simulation algorithm to handle infinitely stiff joints. This substepping should increase accuracy and energy conservation. It should also allow the handling of large mass ratios and fast directional changes within a single time step.

Observation: Tese a ser atribuída ao aluno nº 98659, Miguel Seabra