

Opticks: GPU photon simulation via NVIDIA® OptiX™

Simon C. Blyth^{1,*}

¹Institute of High Energy Physics, CAS, Beijing, China.

Abstract. Opticks is an open source project that accelerates optical photon simulation by integrating NVIDIA GPU ray tracing, accessed via NVIDIA OptiX, with Geant4 toolkit based simulations. A single NVIDIA Turing architecture GPU has been measured to provide optical photon simulation speedup factors exceeding 1500 times single threaded Geant4 with a full JUNO analytic geometry automatically translated from the Geant4 geometry. Optical physics processes of scattering, absorption, scintillator reemission and boundary processes are implemented within CUDA OptiX programs based on the Geant4 implementations. Wavelength-dependent material and surface properties as well as inverse cumulative distribution functions for reemission are interleaved into GPU textures giving fast interpolated lookup or wavelength generation.

Opticks[1-5] enables Geant4[6-8]-based optical photon simulations to benefit from high performance GPU ray tracing made accessible by NVIDIA® OptiX™[9-11]. The Jiangmen Underground Neutrino Observatory (JUNO)[12] under construction in southeast China features the world's largest liquid scintillator detector, with a 20 kton spherical volume of 35 m diameter. The large size and high photon yield of the scintillator, makes the JUNO optical photon simulation extremely computationally challenging with regard to both processing time and memory resources. Opticks eliminates both these bottlenecks by offloading the optical photon simulation to the GPU.

Opticks auto-translates Geant4 detector geometries to GPU optimized forms without approximation. This translation was developed in the context of the JUNO detector. Any detector simulation limited by optical photons can benefit from Opticks. Drastically improved optical photon simulation performance can be transformative to the design, operation and understanding of diverse optical systems. Several groups from various neutrino experiments and dark matter search experiments are evaluating Opticks. Recent Opticks developments allow the optical photon simulation performance to benefit from ray trace dedicated processors, called RT cores[13], available in NVIDIA Turing architecture GPUs.

The most computationally demanding aspect of optical photon simulation is the calculation, at each step of the propagation, of intersection positions of rays representing photons with the geometry of the system. This ray tracing limitation of optical photon simulation is shared with the synthesis of realistic images in computer graphics. The computer graphics community has continuously improved ray tracing techniques. The Turing GPU architecture introduced by NVIDIA in 2018 is marketed as the world's first Ray-tracing GPU, with hardware "RT Cores" dedicated to the acceleration of ray geometry intersection. NVIDIA claims performance of more than 10 billion ray intersections per second, which is a factor 10 more than possible with earlier GPUs which perform the intersection acceleration in software.

*e-mail: simon.c.blyth@gmail.com.

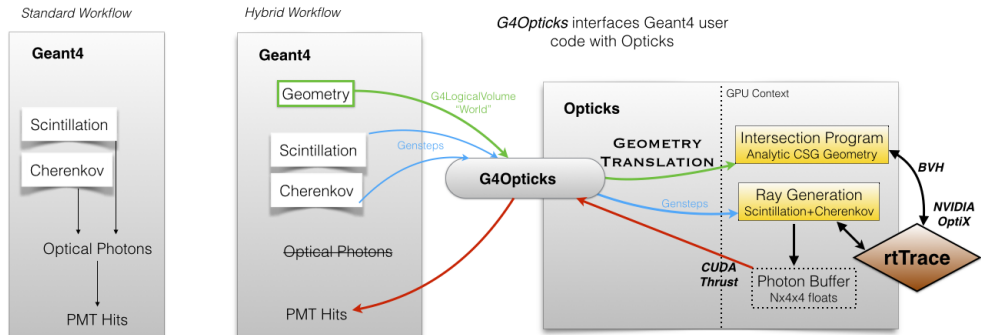


Figure 1. Comparison of the standard workflow of Geant4 optical photon simulation (left) with the hybrid Geant4 + Opticks workflow (right) using a single G4Opticks interface class.

Implementing an efficient GPU optical photon simulation equivalent to the Geant4 simulation requires that all aspects of the Geant4 context relevant to optical photon generation and propagation are translated into an appropriate form and uploaded to the GPU. Figure 1 illustrates the hybrid simulation workflow. At initialization the Geant4 top volume pointer is passed to Opticks which translates the geometry and constructs the OptiX GPU context including intersection, bounding box and closest hit CUDA programs and buffers that these programs access. GPUs contain hardware dedicated to fast texture lookup and interpolation, that is exploited via a single 2D float4 "boundary" texture containing interleaved material and surface properties as a function of wavelength for all unique boundaries. The boundary index returned from a ray traced primitive intersection enables four wavelength interpolated material or surface properties to be obtained from a single hardware optimized texture lookup.

The Opticks geometry model is based upon the observation that many elements of a detector geometry are repeated demanding the use of instancing for efficient representation, see Figure 2. Geometry instancing is a technique used in computer graphics libraries including OpenGL and NVIDIA OptiX that avoids duplication of information on the GPU by storing repeated elements only once together with 4x4 transform matrices that specify the locations and orientations of each instance. A digest string for every structure node is formed from the transforms and shape indices of its progeny nodes. Subsequently repeated sub-trees and their placement transforms are identified using the digests, after disqualifying repeated sub-trees that are contained within other repeats. All structure nodes passing instancing criteria are assigned an instance index with the remainder forming the global non-instanced group. These instanced sub-trees are used for the creation of the NVIDIA OptiX analytic geometry instances, and OpenGL mesh geometry instances.

Photons are generated on the GPU via NVIDIA OptiX ray generation programs, using CUDA ports of Geant4 photon generation loops and "genstep" buffers collected within modified scintillation and Cherenkov processes. Instead of generating photon secondary tracks, "genstep" parameters such as the process type code, the number of photons to generate and the line segment along which to generate them are collected. Collecting and copying gensteps to the GPU rather than photons avoids allocation of CPU memory for the photons, only collected photon hits require CPU memory allocation.

Opticks aims to provide GPU accelerated optical photon simulation for any detector. Achieving this requires physicists from many experiments to use and improve Opticks. Snowmass can assist by introducing Opticks to a wider community.




Figure 2. Cutaway OpenGL rendering of millions of simulated optical photons from a 200 GeV muon crossing the JUNO liquid scintillator. Each line corresponds to a single photon with line colors representing the polarization direction. Primary particles are simulated by Geant4, scintillation and Cerenkov "gensteps" are uploaded to the GPU and photons are generated, propagated and visualized all on the GPU. Representations of some of the many thousands of photomultiplier tubes that instrument the liquid scintillator are visible. The acrylic vessel that contains the liquid scintillator is not shown.

References

- [1] Opticks Repository, <https://bitbucket.org/simoncblyth/opticks/>
- [2] Opticks References, <https://simoncblyth.bitbucket.io>
- [3] Opticks Group, <https://groups.io/g/opticks>
- [4] S. Blyth, EPJ Web Conf. **214**, 02027 (2019)
<https://doi.org/10.1051/epjconf/201921402027>
- [5] Blyth Simon C 2017 J. Phys.: Conf. Ser. **898** 042001
<https://doi.org/10.1088/1742-6596/898/4/042001>
- [6] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce et al., Nucl. Instrum. Methods. Phys. Res. A **506**, 250 (2003)
- [7] J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Dubois, M. Asai et al., IEEE Trans Nucl Sci, **53**, 270 (2006)
- [8] J. Allison, K. Amako, J. Apostolakis, P. Arce, M. Asai, T. Aso et al., Nucl. Instrum. Methods. Phys. Res. A **835**, 186 (2016)
- [9] OptiX: a general purpose ray tracing engine
S. Parker, J. Bigler, A. Dietrich, H. Friedrich, J. Hoberock et al., ACM Trans. Graph.: Conf. Series **29**, 66 (2010)
- [10] OptiX introduction, <https://developer.nvidia.com/optix>
- [11] OptiX API, <http://raytracing-docs.nvidia.com/optix/index.html>
- [12] Neutrino physics with JUNO F. An et al., J. Phys. G. **43**, 030401 (2016)
- [13] NVIDIA RTX, <https://developer.nvidia.com/rtx>