

Photon tracking with GPUs in IceCube

Dmitry Chirkin^a, for the IceCube Collaboration¹

^aUniversity of Wisconsin, Madison, WI 53703, USA

Abstract

GPUs (graphics processing units) have become increasingly popular in the recent years for scientific calculations involving large numbers of similar steps. Photon propagation is a necessary part of simulating detector response to passing charged particles in IceCube that is an ideal application for use with GPUs. We discuss the principle ideas and practical issues of running such an application within the simulation chain used within our collaboration.

Keywords: photon propagation, IceCube

1. Photon tracking: introduction

Traditionally, tracking of the photons in IceCube is performed with photonics [1]. This involves tracking a large number of photons within the detector so that probability distribution functions can be built for all interesting initial and final photon coordinates and directions, and for all possible arrival times. These are tabulated or optionally parametrized into interpolation tables and saved on disk. The tables are then used during the simulation or reconstruction to estimate the mean numbers of photons arriving at given times, and the actual numbers of photons are sampled from Poisson distributions with those means.

This process is complicated by the fact that the tables typically need to be created in 7 or more dimensions and each dimension needs to be tabulated with sufficiently small granularity. This normally leads to tables that are so large that they cannot be fit into the memory of typical computing nodes on which the simulation is run. The table generation is slow, the resulting simulation suffers from a wide range of binning artifacts, and the simulation is slow with much time spent loading the tables into memory. In the recently developed approach some of these problems are reduced by parametrization and interpolation techniques, which not only reduce the binning artifacts and improve precision of the tables, but also speed up the simulation.

Despite the recent improvements the process remains a 2-step process, in which the photon tables are first created and then used during the simulation. The reason for this has traditionally been the fact that running the simulation this way was still significantly faster than a direct approach, in which photons are created and propagated as needed, during the course of the simulation itself.

Email address: dima@icecube.wisc.edu (Dmitry Chirkin)

¹<http://icecube.wisc.edu>

2. Photon Propagation Code

The photon propagation code (PPC) [2] was initially written to study the feasibility of direct photon propagation for simulation of events in IceCube. The simple nature of photon propagation physics allowed us to focus on the code optimization, to make sure the simulation ran as fast as possible. The simulation was written in C++, then re-written entirely in Assembly for the 32-bit i686 architecture with SSE vector optimizations. The Assembly version of the program used the SSE instructions for photon rotation and locating the optical sensor closest to the photon segment, while the calculation of the scattering angle was performed in one go using only the registers of the FPU stack.

A project called *i3mcm1* [3] demonstrated that significant acceleration of the photon propagation is possible by using the graphics processing units (GPUs). We confirmed this with a version of PPC that employs the NVIDIA GPUs (graphics processing units) via the CUDA programming interface [4]. Recently a new version was written that additionally uses OpenCL [4], supporting both NVIDIA and AMD GPUs, and also multi-CPU environments. The relative performance of these different implementations (for simulating both in-situ light sources, or flashers, and Cerenkov light from muons) is compared in Table 1. We have studied the simulation with these versions of PPC and *i3mcm1* and were able to demonstrate excellent agreement between them.

	C++	Assembly	GTX 295 GPU
flasher	1.00	1.25	147
muon	1.00	1.37	157

Table 1: Speedup factor of different implementations of PPC compared to the C++ version. The GPU used in this comparison was either of the two in the NVIDIA GTX 295 video card.

The reason for the substantial acceleration of PPC on GPUs is the highly parallel nature of the simulation of the photon propagation. All of the simulated photons go through the same

62 simulation steps (see Figure 1): photon propagation between
 63 the scattering points, calculation of the scattering angle and new
 64 direction, and evaluation of whether the current photon segment
 65 intersects any of the optical sensors of the detector array. The
 66 GPUs are designed to perform the same computational operation
 67 in parallel across multiple threads. Each thread works
 68 on its own photon for as long as the photon exists. When the
 69 photon is absorbed or hits the detector the thread receives the
 70 new photon from a pool of photons for as long as that pool is
 71 not empty. Although a single thread runs slower than a typical
 72 modern computer CPU core, running thousands of them in parallel
 73 results in the much faster processing of photons from the
 74 same pool on the GPU.




Figure 1: Parallel nature of the photon propagation simulation: tracking of photons entails the computationally identical steps: propagation to the next scatter, calculation of the new direction after scatter, and evaluation of intersection points of the photon track segment with the detector array. These same steps are computed simultaneously for thousands of photons.

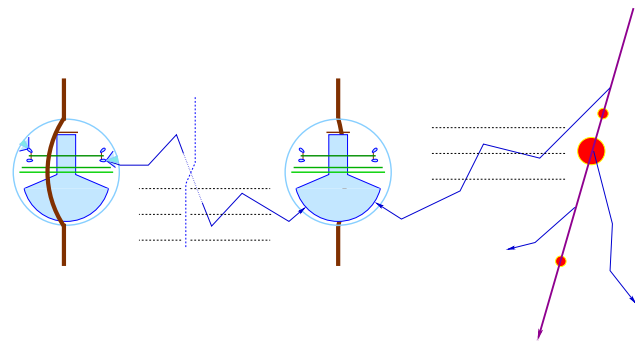


Figure 2: Typical simulation scenarios: photons emitted by the detector are tracked as part of the calibration procedure (left). Cerenkov photons emitted by a passing muon and cascades along its track are tracked to simulate the typical IceCube events (right).

101 is the tilt of the ice layers, i.e., dependence of the ice parameters
 102 not only on the depth, but also on the xy surface coordinates
 103 (shown in Figure 3).

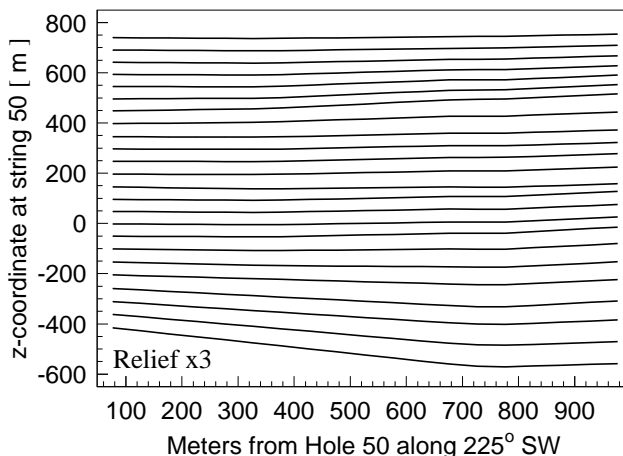


Figure 3: Extension of ice layers along the average gradient direction (taken from [6]). The y-axis shows the layer shift (relief) from its position at the location of a reference string at distance shown on the x-axis from this string along the average gradient direction (225 degrees SW). The relief shown is amplified by a factor of 3 for visual clarity of the ice layer tilt.

75 3. Simulation with photon propagation code

76 The direct photon simulation with PPC is typically used in
 77 the two scenarios shown in Figure 2. In the first the in-situ
 78 light sources of the detector are simulated for calibrating the
 79 detector and the properties of the surrounding ice. It is possible
 80 to very quickly re-simulate the detector response to a variety
 81 of ice scattering and absorption coefficients finely tabulated in
 82 depth bins. This allows for these coefficients to be fit directly,
 83 by finding the combination that is a best simultaneous fit to all
 84 of the in-situ light source calibration data [5]. For the 10 meter
 85 depth bins, 200 coefficients are fitted (with scattering and absorption
 86 defined in 100 layers spanning 1 km of depth of the
 87 detector), with nearly a million possible ice parameter configurations
 88 tested in less than a week on a single GPU-enabled
 89 computer. This method is intractable with the photonics-based
 90 simulation, as each new parameter set would require generation
 91 of the new set of photonics tables, each generation taking on
 92 the order of a week of computing time of a ~ 100 -CPU cluster.

93 In the second scenario the Cerenkov photons created by
 94 the passing muons and cascades are simulated as part of the
 95 larger simulation of the detector response to atmospheric and
 96 other fluxes of muons and neutrinos. The simulation is able
 97 to account for some effects that are difficult to implement with
 98 the photonics-based simulation, because their simulation would
 99 lead to additional degrees of freedom, thus increasing the size
 100 of the parametrization effort and tables many-fold. One of these

Effects that are treated precisely with PPC (and only approximately with photonics) include the simulation of the longitudinal profile of light generation by cascades and the angular distribution of the Cerenkov photons around the emitting muons or cascades.

Other effects implemented recently only with PPC include the direct simulation of the somewhat different ice properties in the column of ice refrozen around the detector strings after they have been deployed; and the slight azimuthal dependence of the scattering function.

4. Concurrent execution and runtime optimization

The PPC keeps track of several execution time counters that help judge the performance of the GPU code. One counter is es-

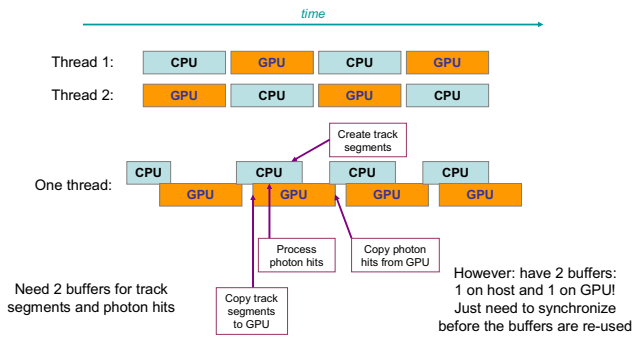


Figure 4: Concurrent execution on CPU and GPU sides. Two possible solutions are shown: the first (top) requires running at least two ppc threads was eventually replaced with the solution (bottom) that is enabled in a single thread.

117 timed on the CPU side, by calculating the time elapsed wait-172
 118 ing for the GPU code to return. While waiting for the GPU,173
 119 code the CPU can run other parts of the simulation (e.g., trigger174
 120 simulation), but to maximize the use of the GPUs these parts,175
 121 should finish quicker than the GPU part. The typical utilization,176
 122 of the GPUs achieved in our tests is $\gtrsim 90\%$. The implemen-177
 123 tation of the concurrent execution of the program on both CPU,178
 124 and GPU sides is facilitated by the fact that the GPU side works,179
 125 on its own memory buffer, which is exchanged with the main
 126 computer memory buffer only at the beginning or at the end of
 127 the execution on the GPU side. So, while the CPU processes
 128 photon detector hits created by the previous run on the GPU,
 129 the GPU is running through the photons previously prepared by
 130 the CPU (see Figure 4).

131 The other two execution time counters calculate the mini-183
 132 mum and maximum time spent by different threads running on
 133 the GPU. If these are close to each other, all execution units of
 134 the GPU have been equally loaded. The difference we observe
 135 is typically on the order of $\sim 0.5\%$.

5. Hardware considerations: our GPU cluster

137 The simulation of a single day of experimental background
 138 data of the full IceCube detector completes in about 10 days,
 139 of calculation on a small 3-computer GPU cluster equipped
 140 with 18 GPUs (3 NVIDIA GTX 295 cards per computer, each
 141 card contains 2 GPUs). This 3-computer cluster was built from
 142 consumer-available parts for approximately \$10k. This cluster
 143 is assisted by a larger CPU-only computing cluster that runs the
 144 simulation of cosmic rays and several other tasks.

145 Unfortunately we experienced problems with 3 out of the
 146 24 GPUs that we have experimented with. The problems in-
 147 cluded jobs exiting with errors, GPUs locking up, or sometimes
 148 even the entire computer locking up and requiring rebooting.
 149 Further investigation revealed that the jobs that fail also occa-
 150 sionally produce NaNs or INFs in their GPU threads. With
 151 a small bit of in-line assembly we found that only threads run-
 152 ning on specific hardware multiprocessors (MPs) on each of the
 153 faulty GPUs (each of the GTX 295 GPUs has 30 MPs) are af-
 154 fected. By checking within threads during execution whether
 155 they are running on the faulty MP and immediately stopping

156 those threads it was possible to use the faulty GPUs for calcu-
 157 lation at 29/30 % capacity. The jobs running on these GPUs
 158 are 3% slower, but run without further problems. We do con-
 159 tinuously monitor and record the unexpected NAN or INF con-
 160 ditions in the GPU threads or problems reported by the CUDA
 161 interface; however these became extremely rare after disabling
 162 the 3 faulty MPs as described above.

163 To optimize the use of the GPU-enabled computers fur-
 164 ther we are experimenting the Directed Acyclical Graph (DAG)
 165 tools [7]. This involves separating simulation segments into
 166 separate tasks, and assigning these tasks to DAG nodes. DAG
 167 assigns separate tasks to different computer nodes; execution of
 168 photon propagation simulation is performed on dedicated GPU
 169 nodes.

170 For many simulations the GPU segment of the simulation
 171 chain is much faster than the rest of the simulation. For these,
 172 a small number of GPU-enabled machines can consume the data
 173 from a large pool of CPU cores. However, the optimal DAG
 174 configuration differs depending on the specific simulation.

175 We are currently running the GPU simulation routinely on
 176 our cluster at UW-Madison, which is being upgraded to include
 177 48 more Tesla M2070 GPUs (built around the 3 x DELL Pow-
 178 erEdge C410x and 6 x DELL PowerEdge C6145 for \sim \$200k).
 179 We are experimenting with running the PPC-based simulation
 180 on other IceCube sites in U.S., Canada, and Germany.

6. Concluding remarks

181 We have developed a photon propagation tool that can re-
 182 place the older two-step photon tracking paradigm in certain
 183 situations, while achieving more precision, better description
 184 of the physics of the process and shorter run time. The program
 185 is capable of running on both CPU cores and GPU hardware,
 186 achieving very significant speed up (factors in excess of ~ 100)
 187 on the latter.

188 Although we have encountered some hardware problems
 189 while running on 3 out of 24 of our consumer-grade GPU cards,
 190 it was possible to identify and disable the faulty parts of the
 191 GPUs and continue to use the problem cards at 97 % capacity.
 192 We are in a process of building a professional-grade cluster of
 193 high-end GPU nodes, which will contribute to further produc-
 194 tion of the simulated data.

References

195 [1] J. Lundberg, et al., Light tracking through ice and water - scattering and
 196 absorption in heterogeneous media with photonics, Nucl. Instrum. Methods
 197 A 581 (2007) 619. arXiv:astro-ph/0702108v2.
 198 [2] D. Chirkin, photon propagation code:
 199 <http://icecube.wisc.edu/~dima/work/WISC/ppc>.
 200 [3] T. Abuzayyad, i3mcm1: <http://wiki.icecube.wisc.edu/index.php/i3mcm1>.
 201 [4] NVIDIA CUDA: <http://www.nvidia.com/cuda>.
 202 [5] D. Chirkin, et al., Study of south pole ice transparency with icecube
 203 flashers, Contribution to the 32st International Cosmic Ray Conference,
 204 Beijing, China, 11-18 August 2011, session HI.2.3, contribution 0333.
 205 [6] R. C. Bay, et al., South pole paleowind from automated synthesis of ice
 206 code records, J. Geophys. Res. 115 (2010) D14126.
 207 [7] DAG: http://en.wikipedia.org/wiki/directed_acyclic_graph.