


IceCube T1 Parallel

Max, Omar

2026.10.01



Benchmarking?

- $1e2-1e4$ GeV: 100 events, $\gamma=-1.5$
 - 127GPUs/100events $\rightarrow$ 70k GPUh for full dataset (RTX 2080 Ti)
- $1e4-1e6$ GeV: 100 events, $\gamma=-1.5$
 - 310GPUs/100events $\rightarrow$ 34k GPUh for full dataset (on RTX 2080 Ti)
 - 270GPUs/100events $\rightarrow$ 30k GPUh for full dataset (on A100)
- $1e6-1e8$ GeV: 10 events, $\gamma=-1$.
 - 1106GPUs/10 events $\rightarrow$ 120k GPUh for full dataset (on A100)


Target $2e8$ events

Target $4e7$ events

Target $4e6$ events

Wall time of Prometheus

Stacked histograms for ice-like detector
Photon propagation done by ppc



Execution time
for 1k events

Rough wall time estimates based on IceCube datasets

Energy range	Simulated power law index	DOM oversizing	CPU kh / flavor	GPU kh / flavor	N Events generated
10^2 - 10^4 GeV	-1.5	1	1.2	3.1	$2\text{E}+08$
10^4 - 10^6 GeV	-1.5	1	0.6	5.1	$4\text{E}+07$
10^6 - 10^8 GeV	-1.0	3	3.2	25.2	$4\text{E}+06$