

A vertical structure of the IceCube detector, showing a series of spherical modules connected by cables, with a blue and green glow.

IceCube's GPGPU's Cluster For Extensive MC Production

Heath Skarlupka

University of Wisconsin Madison





Outline

- What is IceCube?
- The IceCube Simulation Chain
- Using GPUs In Simulation
- GPU resources in Condor



IceCube



Installation



Life At The South Pole




Life At The South Pole



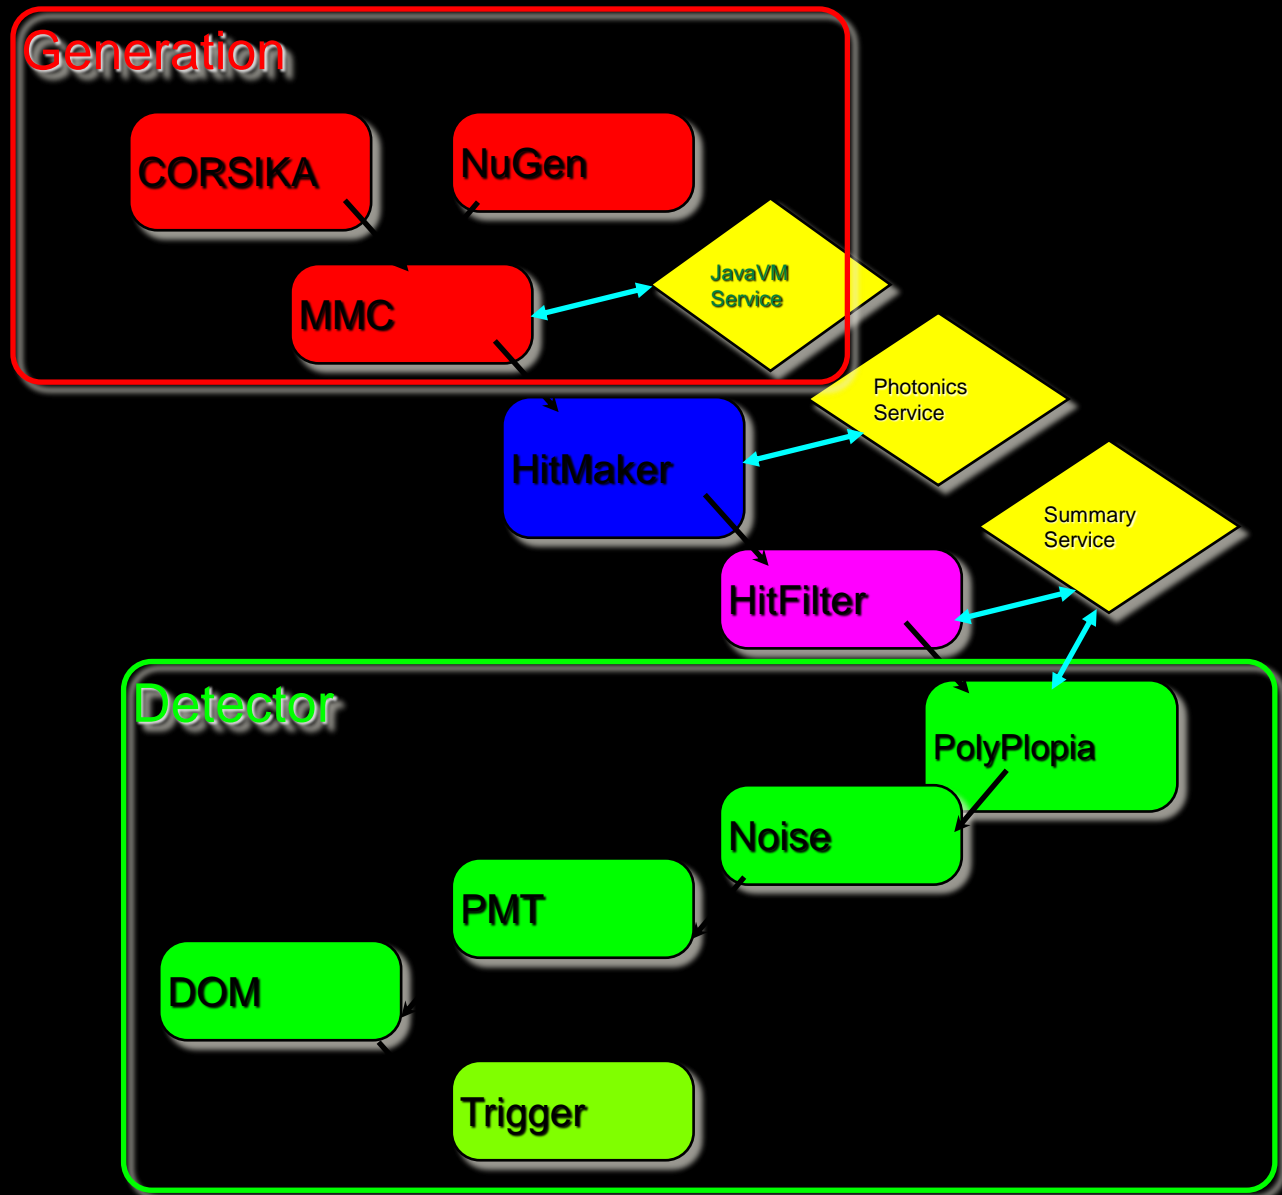
Life At The South Pole





Simulation

Simulation Chain




DAG (Directed Acyclical Graph) Based Simulation

Tasks

Name	Tray	Iter	Host	Status	Start	Finish
generator	0	0	cabinet-7-7-31.t2.ucsd.edu	OK	2010-01-08 11:39:03	2010-01-08 14:23:33
photronics	1	0	g10n05.hep.wisc.edu	OK	2010-01-08 17:51:53	2010-01-08 18:14:25
photronics	1	1	g16n33.hep.wisc.edu	OK	2010-01-08 17:51:54	2010-01-08 18:13:44
photronics	1	2	g16n25.hep.wisc.edu	OK	2010-01-08 17:51:27	2010-01-08 18:18:56
photronics	1	3	g16n05.hep.wisc.edu	OK	2010-01-08 17:51:34	2010-01-08 18:18:46
photronics	1	4	g16n36.hep.wisc.edu	OK	2010-01-08 17:52:00	2010-01-08 18:21:07
photronics	1	5	cabinet-7-7-20.t2.ucsd.edu	OK	2010-01-08 17:51:59	2010-01-08 18:34:37
photronics	1	6	g12n22.hep.wisc.edu	OK	2010-01-08 17:52:06	2010-01-08 18:08:11
photronics	1	7	g12n31.hep.wisc.edu	OK	2010-01-08 17:52:06	2010-01-08 18:03:37
photronics	1	8	g12n08.hep.wisc.edu	OK	2010-01-08 17:56:22	2010-01-08 18:09:10
ic59	2	0	g14n23.hep.wisc.edu	OK	2010-01-08 19:01:43	2010-01-08 19:17:36
ic59	3	0	g14n23.hep.wisc.edu	OK	2010-01-08 19:17:36	2010-01-08 19:36:15
verification	4	0	cabinet-4-4-25.t2.ucsd.edu	OK	2010-01-08 21:22:12	2010-01-08 21:34:24
trashcan	0	0	cabinet-6-6-28.t2.ucsd.edu	OK	2010-01-08 21:40:42	2010-01-08 21:41:35

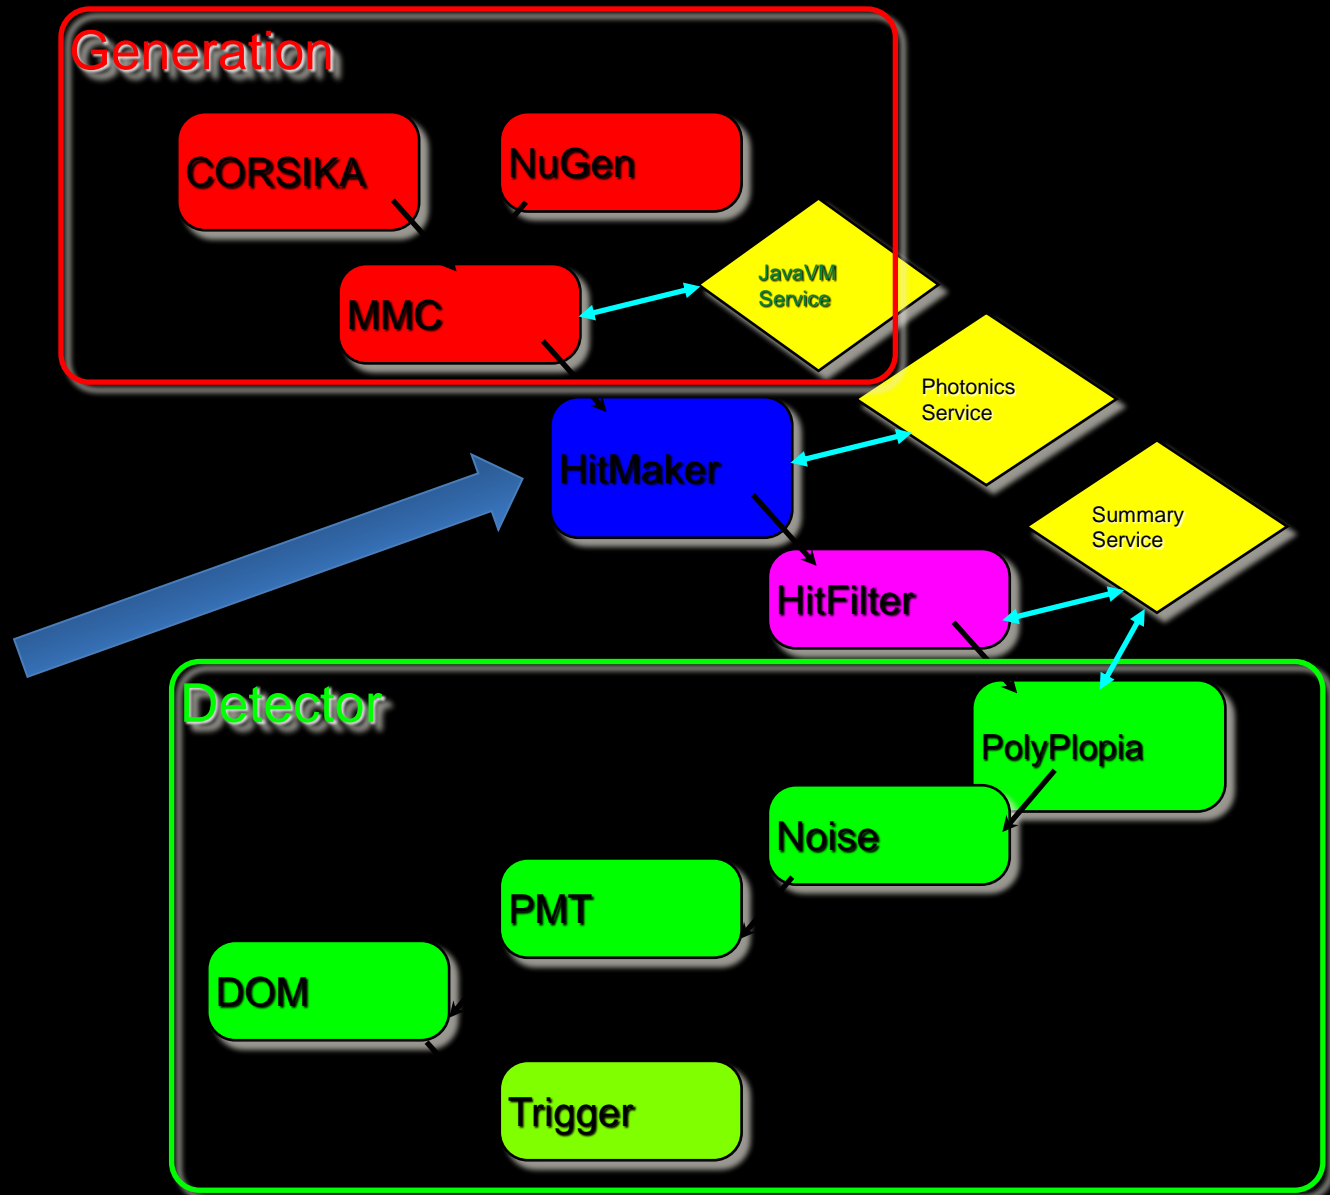
Task Graph



Separate simulation segments into tasks


Assign task to a node in DAG

Problems



Photon Tracking With Tables

- Generating tables is slow
 - Slows changes to detector simulation
- 25 GB table size
 - Can't run on the GRID
 - Copying into memory is slow
- Moving new tables into distributed resources is time consuming



The GPUs



Dell Cloud Products

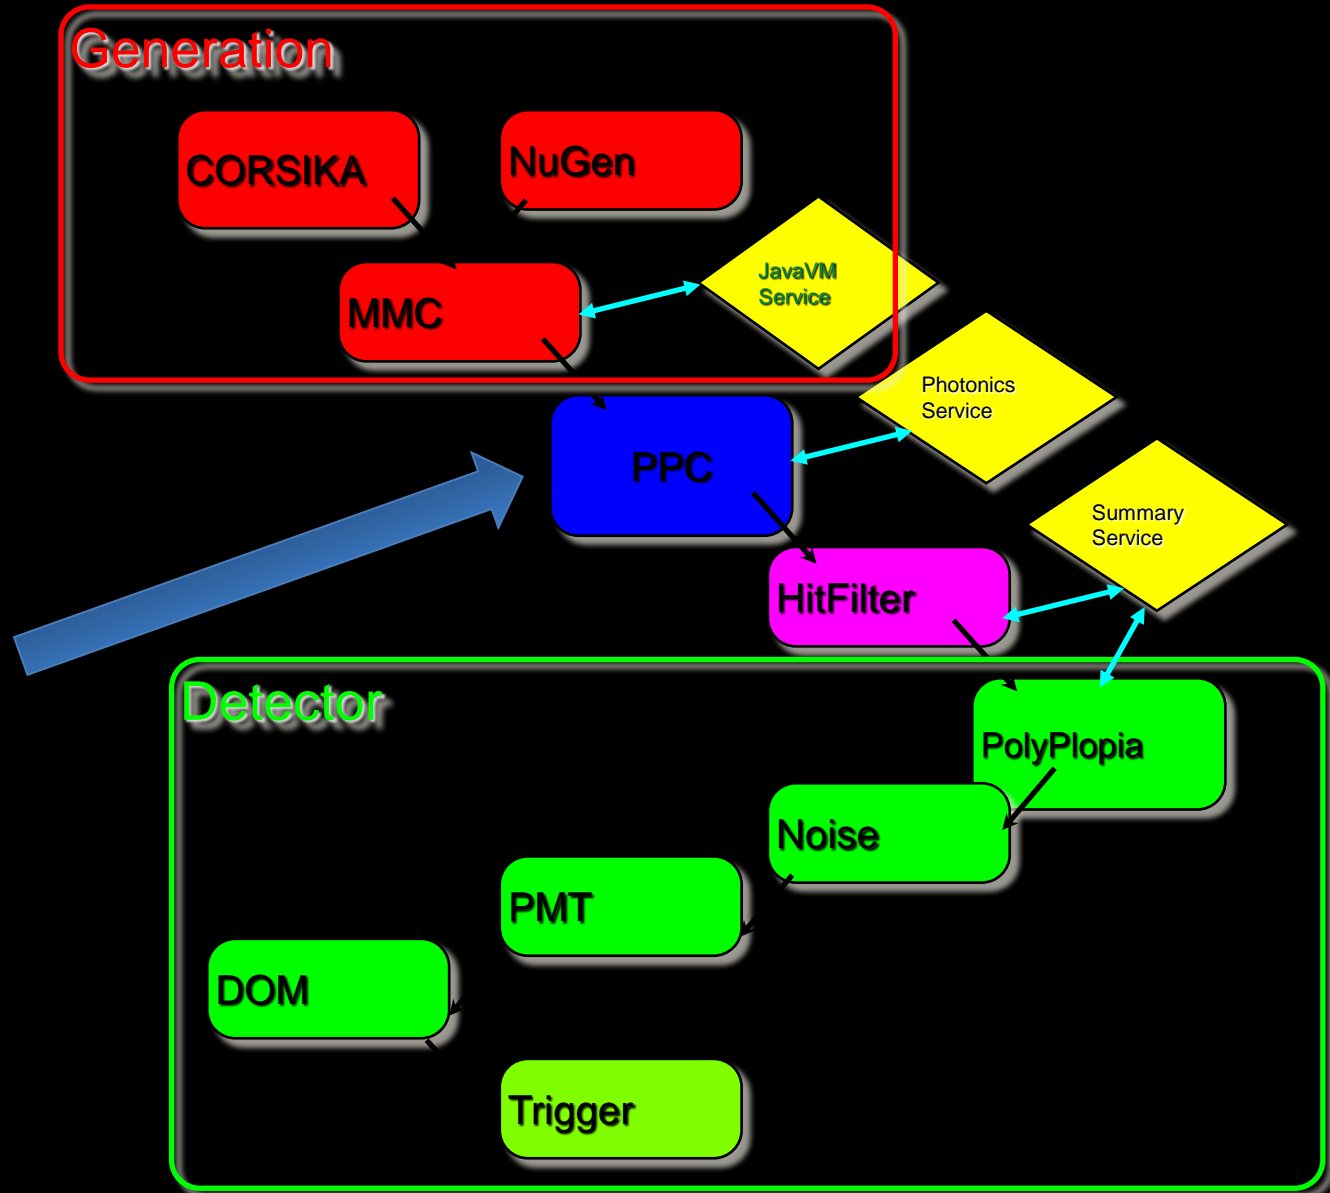
3 C410x




6 C5145s




Photon Propagation Code



Photon Propagation




Each thread runs on
a different GPU core





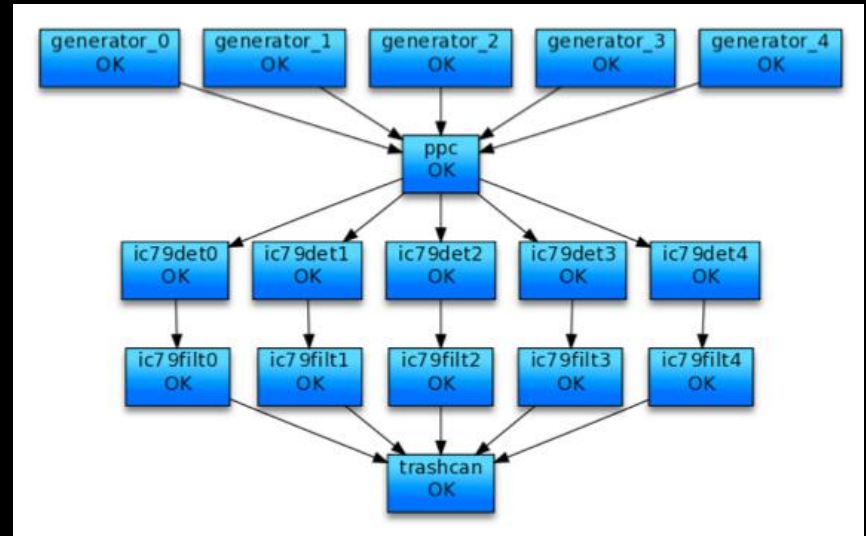
GPU-based simulation

- Execution of photon propagation simulation on dedicated GPU nodes.
- For many simulations GPU segment of chain is much faster than the rest of the simulation.
- Small number GPU-enabled machines can consume the data from large pool of CPU cores.

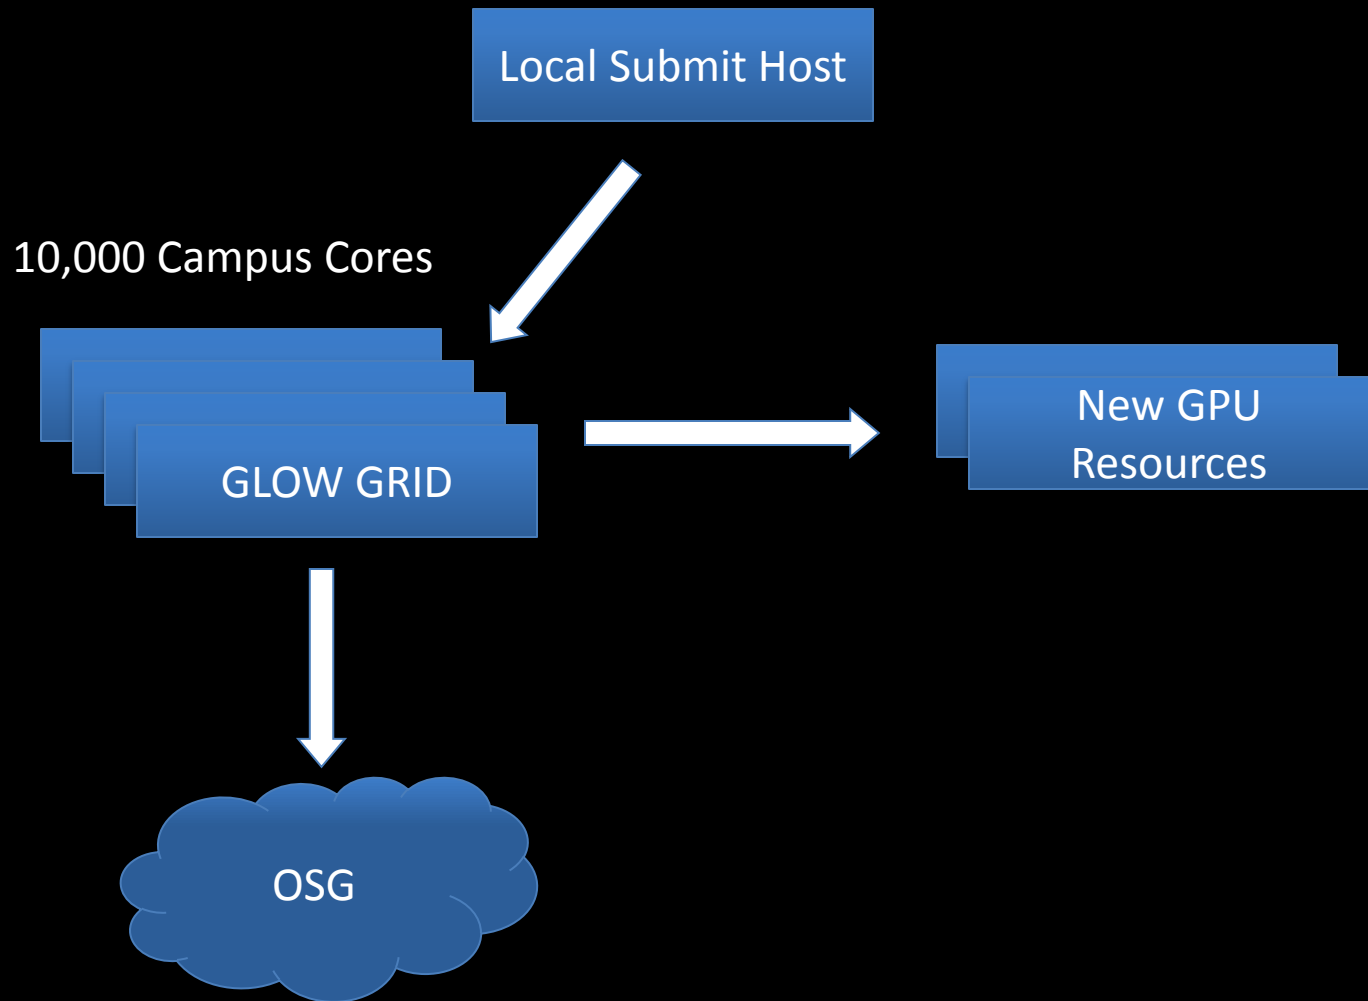


GPU-based simulation

- Optimal DAG differ depending on the specific simulation



Moving To The GRID



Condor Node Configuration

```
START = (My.GZK_GPU_SLOT == TRUE  
&& TARGET.GZK_GPU_JOB == TRUE)
```

```
#GPU Info
```

```
SLOT1_GPU_DEV=0
```

```
...
```

```
SLOT4_GPU_DEV=3
```



Condor Submit File

Executable = nbody.sh

Universe = vanilla

...

environment =

"CUDA_VISIBLE_DEVICES=\$\$(GPU_DEV)"

+GZK_GPU_JOB = TRUE

Working With Condor Team

- Developing ways to abstract GPU resources
- Automation of GPU slot creation

Conclusions

- 150x increase in performance over photon tables method
- Better utilization of GRID resources
- Changes to simulation are incorporated faster

