

IceCube* Simulation

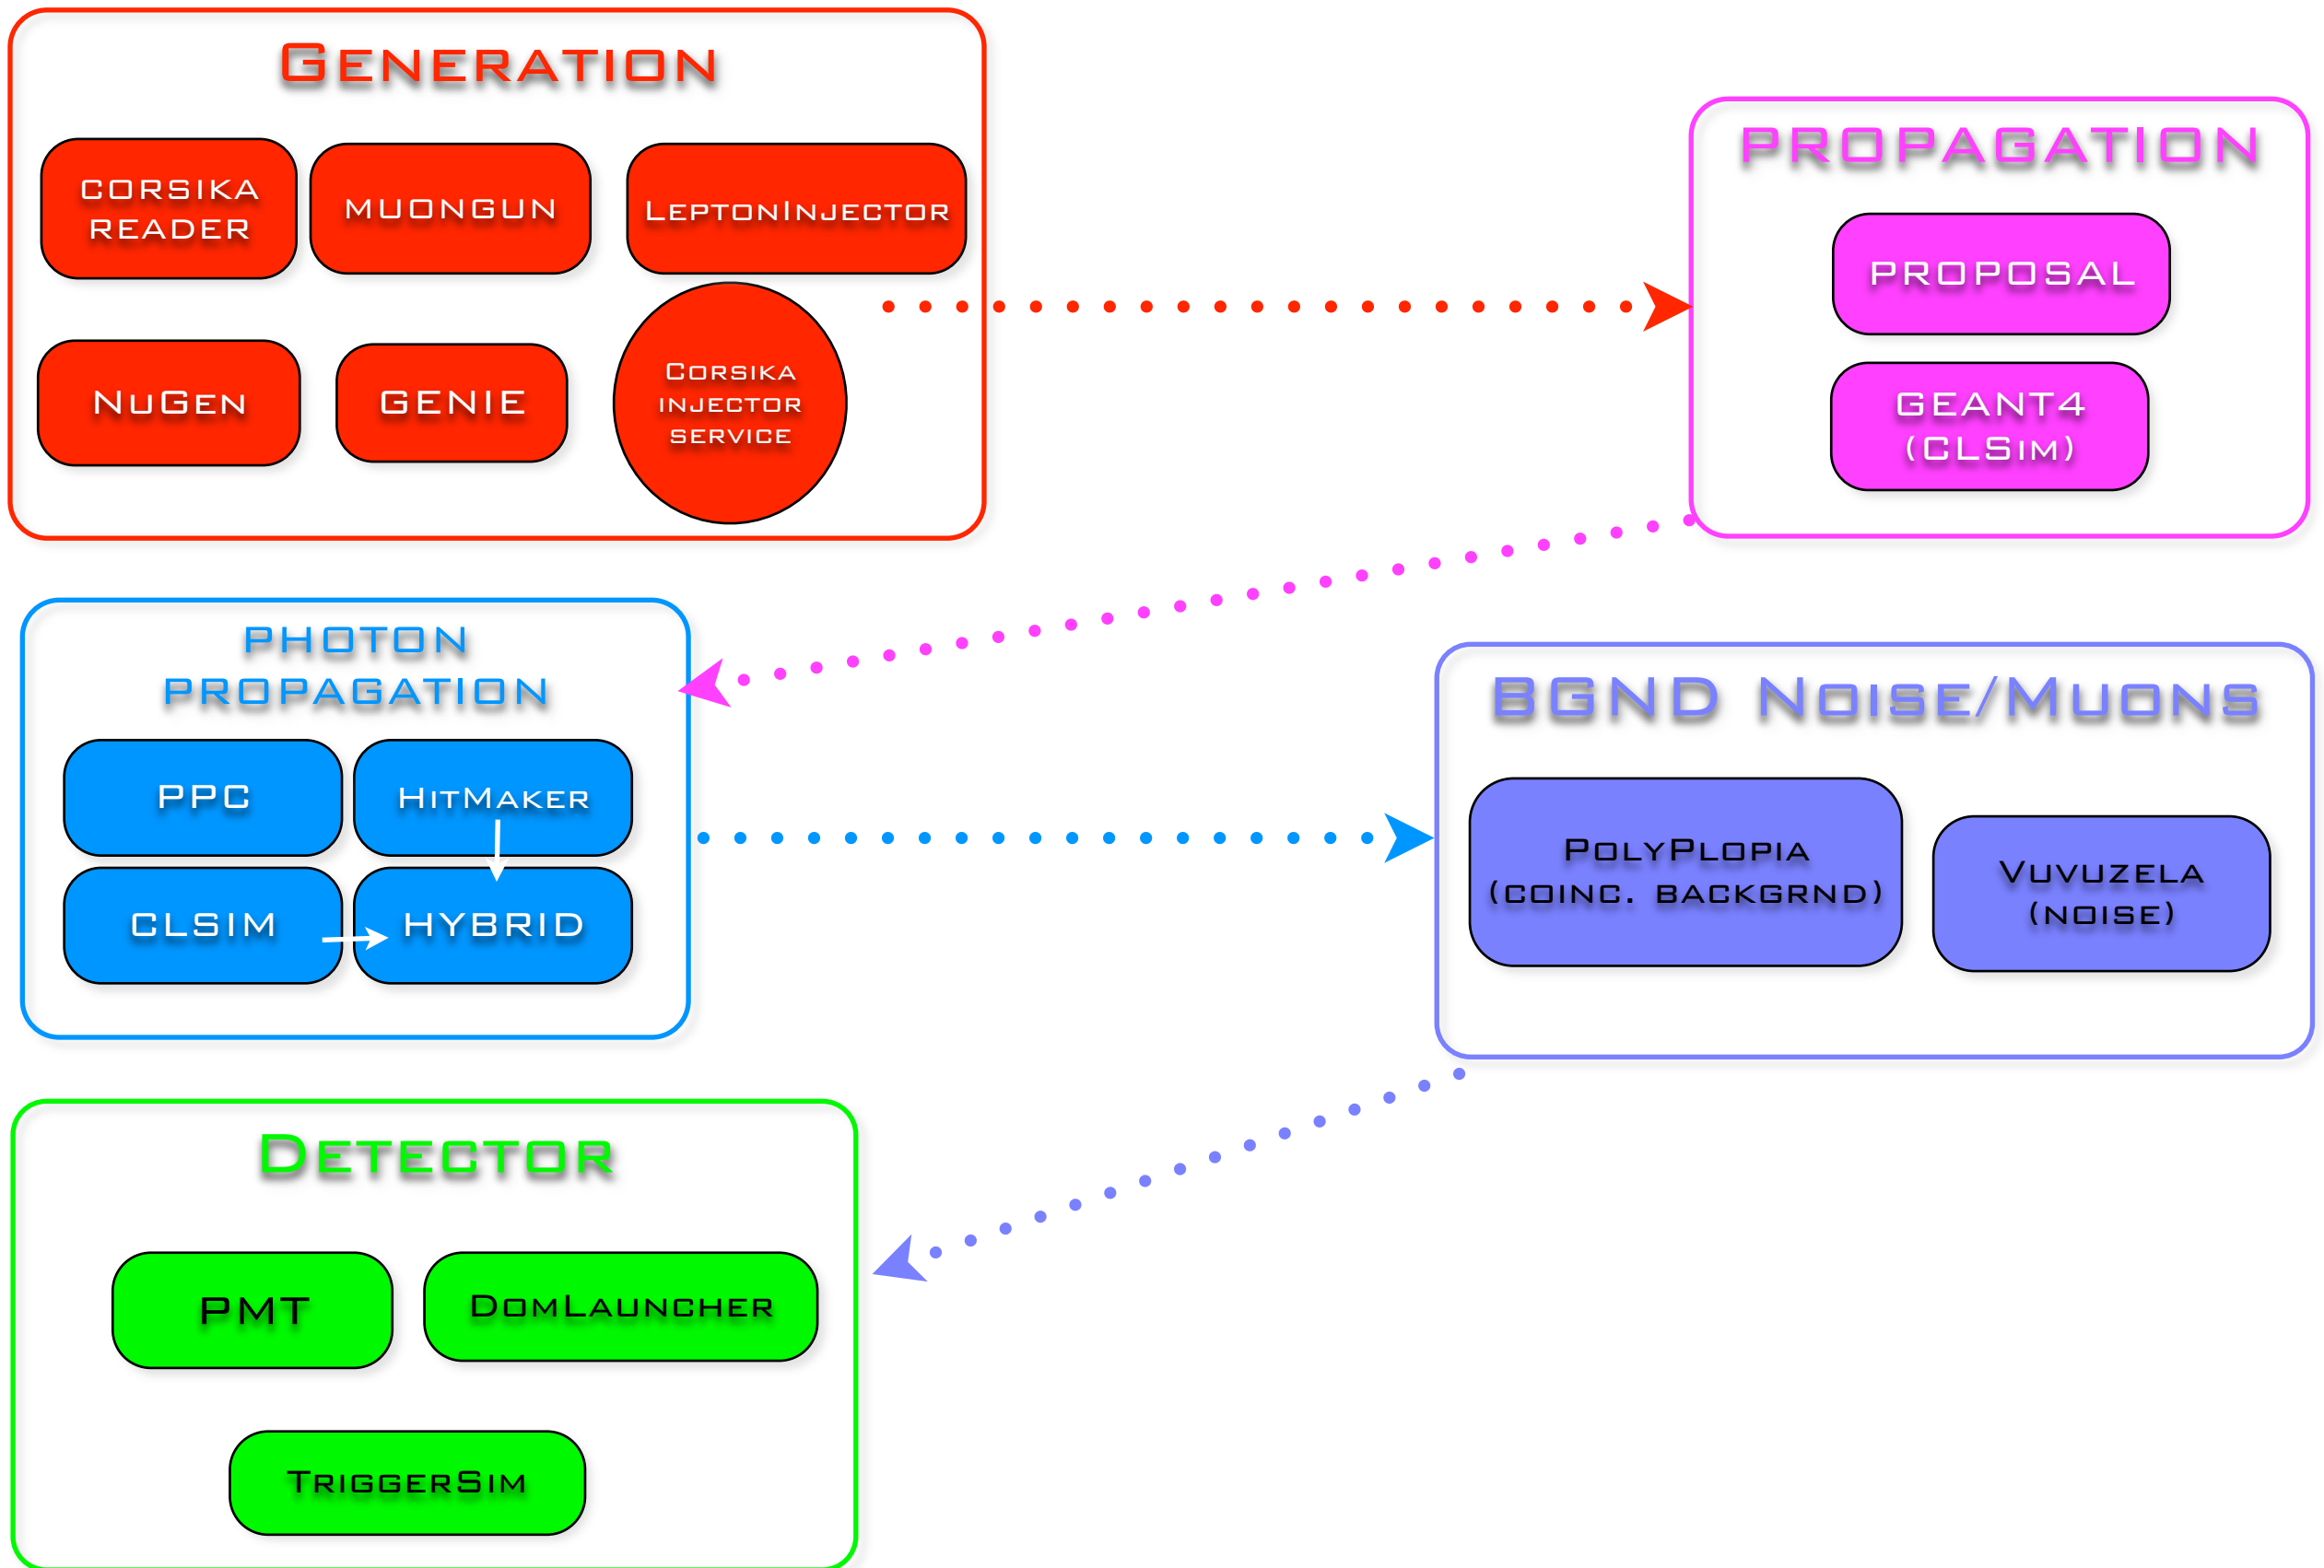
(*and beyond)

Juan Carlos Díaz Vélez
Science and Computing Advisory Panel

Software Goals

- Quality:
 - improve physics in simulation
 - improve data/Monte Carlo agreement
- Quantity:
 - improve throughput and increase statistics
 - reduce resource requirements
 - improve utilization of existing resource
- Two objectives are sometimes at odds
 - Better simulations often require more CPU, GPU, Memory
 - Balance the needs for more statistics vs. better simulation
- Yearly simulation workshops aimed at strategies for achieving these goals

The IceCube Simulation Chain



Generators

- ▶ **Cosmic-ray Air Showers:**

- ▶ ~~CORSIKA (FORTRAN stand-alone)~~
- ▶ ~~corsika-reader: IceTray reader for standard format~~
- ▶ ~~CorsikaInjectorService (IceTop)~~
- ▶ DYNSTACK CORSIKA - Optimized Module with C++11 plugin

- ▶ **Muons:**

- ▶ **MuonGun:** parametrization of flux of atm. muons under the ice.
 - ▶ Requires updated parametrization

- ▶ **Neutrinos:**

- ▶ ~~neutrino-generator: injects neutrinos, propagates them through Earth, forces interaction in detector volume.~~
- ▶ ~~genie-icetray: detailed simulation of neutrino interactions with GENIE. (Used for low-energy simulations)~~
- ▶ **LeptonInjector / NuFSGen:** weighted leptons+weights to account for flux models, interaction models, in-earth propagation, etc.

CORSIKA

DYNSTACK (Kevin Meagher UW-Madison) -

Current generation module and weighting checked against standard (legacy) simulation. Not yet in production. Testing in grid. Included in next software release.

Server socket configuration fixed; simulation now running efficiently with 4:1 CPU:GPU cores. S-frame contains all needed information for weighting (no need to query simprod database). Open call to get involved in developing biased sampling.

IceTop -

To do: implementation of DYNSTACK module to do G4 tank simulation directly before discarding EM shower component.


CORSIKA8 (Antonio Augusto Alves Junior - KIT) -

Development guidelines outlined in *R. Engel et al., Comp.Soft. Big Sci. 3 (2019) 2*. Once released, will replace current DYNSTACK library.

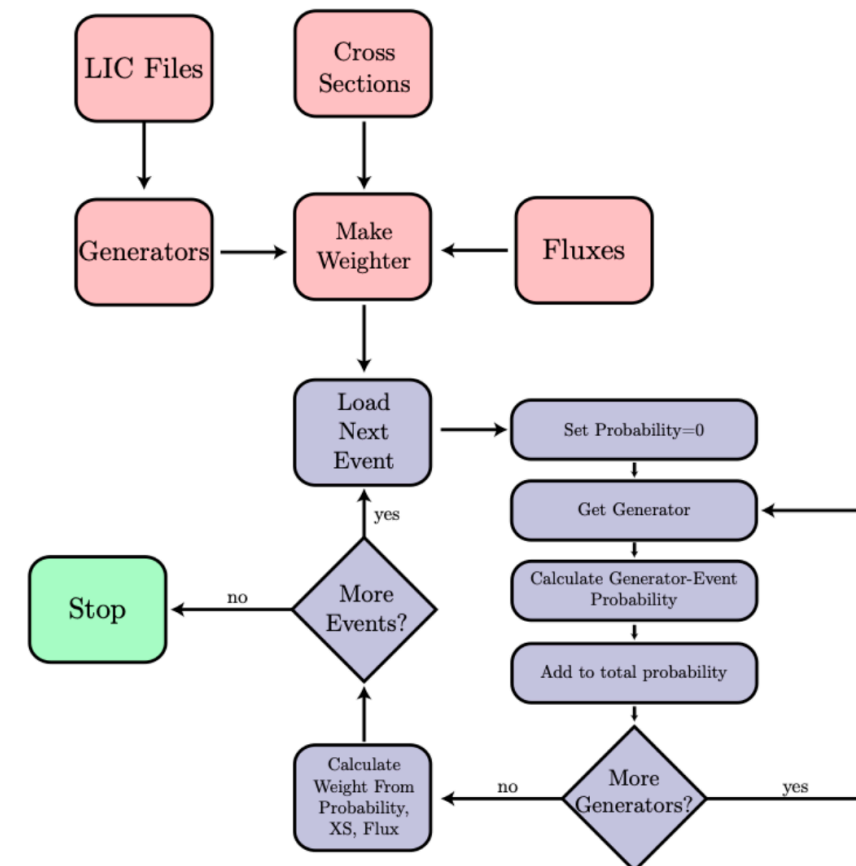
Pending Tasks:

- Implement IceTop module. We discussed a dedicated workshop with people involved.
- Implement DYNSTACK biasing schemes.
- Update DYNSTACK module to support CORSIKA8 (when released)

Lepton Injector / Lepton Weighter (Carlos A. Argüelles , C.Weaver, B. Smithers, and A. Schneider)



- Current inconsistencies between LW and NuGen samples (possibly due to a bug in calibration file for LW).
- Currently weighting Range+Volume mode simulations of the same flavor is not trivial.
- LW needs some work in terms of documentation and user friendliness.



SnowStorm

Efficient propagation of systematic uncertainties from calibration to analysis with the SnowStorm method in IceCube

M.G. Aartsen *et al* JCAP10(2019)048

DOI: [10.1088/1475-7516/2019/10/048](https://doi.org/10.1088/1475-7516/2019/10/048)


Erik Ganster

Why new MC and why SnowStorm?

- Lots of changes/updates since last large-scale MC production (at least DiffuseNuMu is based on 2012 MC):
 - New ice models: Spice3.2.1
 - New SPE templates
 - New software tools
 - ...
- Large collaborative effort to produce new MC sets

| Systematic | Baseline Value | Sampling Distribution | Range |
|------------------------|------------------|-----------------------|------------------------------------|
| IceWavePlusModes | --- | 2x 12 Gaussians | --- |
| Scattering | 1.0 | uniform | [0.9, 1.1] |
| Absorption | 1.0 | uniform | [0.9, 1.1] |
| AnisotropyScale | 1.0 | uniform | [0.0, 2.0] (= 0-15%) |
| DOMEfficiency | 1.0 | uniform | [0.9, 1.1] |
| HoleIceForward_Unified | p0 0.0 p1 0.0 | uniform | p0 [-1.0, +1.0] p1 [-0.2, +0.2] |

Table: Overview of all available systematic perturbations for SnowStorm and their default sampling distributions



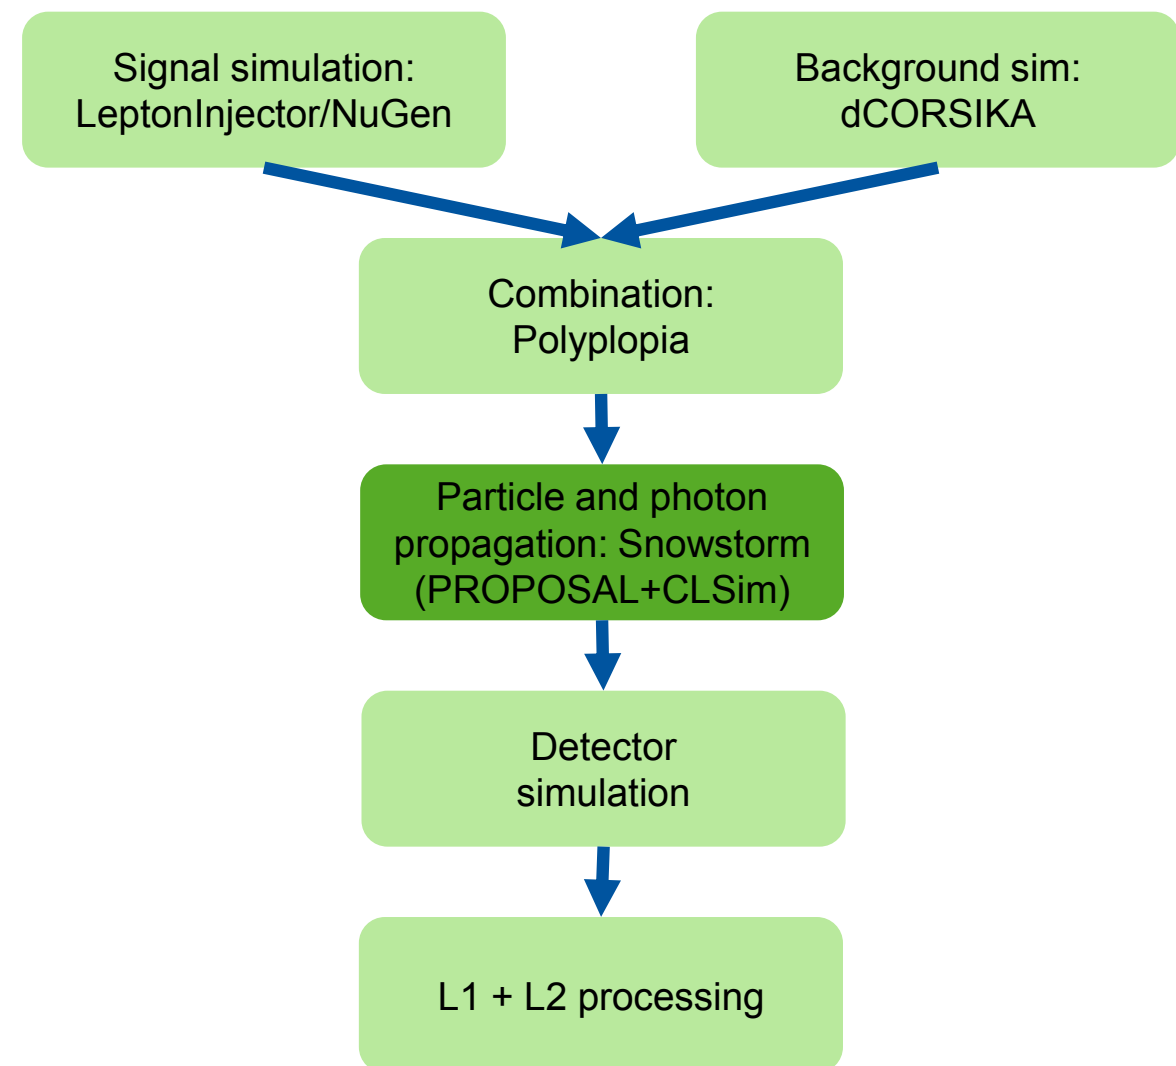
- SnowStorm short: Continuous variation of nuisance parameters (detector systematics) (blue) instead of discrete sets for specific values (red)

SnowStorm

Erik Ganster

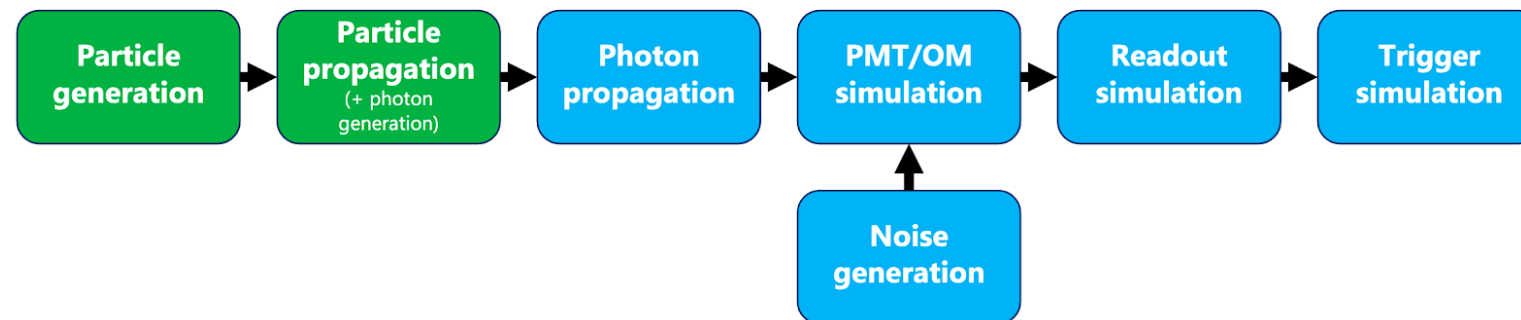
SnowStorm Simulation Chain – SnowStorm

- Based on “standard” simulation chain
- Merge of signal+background I3MCTrees before any particle or photon propagation
→ Ensures that all particles get treated/propagated with the exact same parameters/settings further on
- Main SnowStorm simulation step:
 - Particle (muon) propagation with PROPOSAL
 - Photon propagation using CLSim
- Perturbing the ice model properties for chunks of frames using the *SnowStorm perturber*



Upgrade/Gen2 Simulation chain (Tom Stuttard)

There is currently a working end-to-end simulation for upgrade but there are plenty of placeholders and hacks.



CISim/PPC (DEgg, mDOM):

Currently limited to spherical sensors of comparable size to IceCube DOMs.

- mDOM simulation modelled as flat disks on the surface of a sphere with wavelength and angular acceptance per PMT glass and gel taken into account.
- DEgg: currently modeled as 2 spherical PTMs.

Noise:

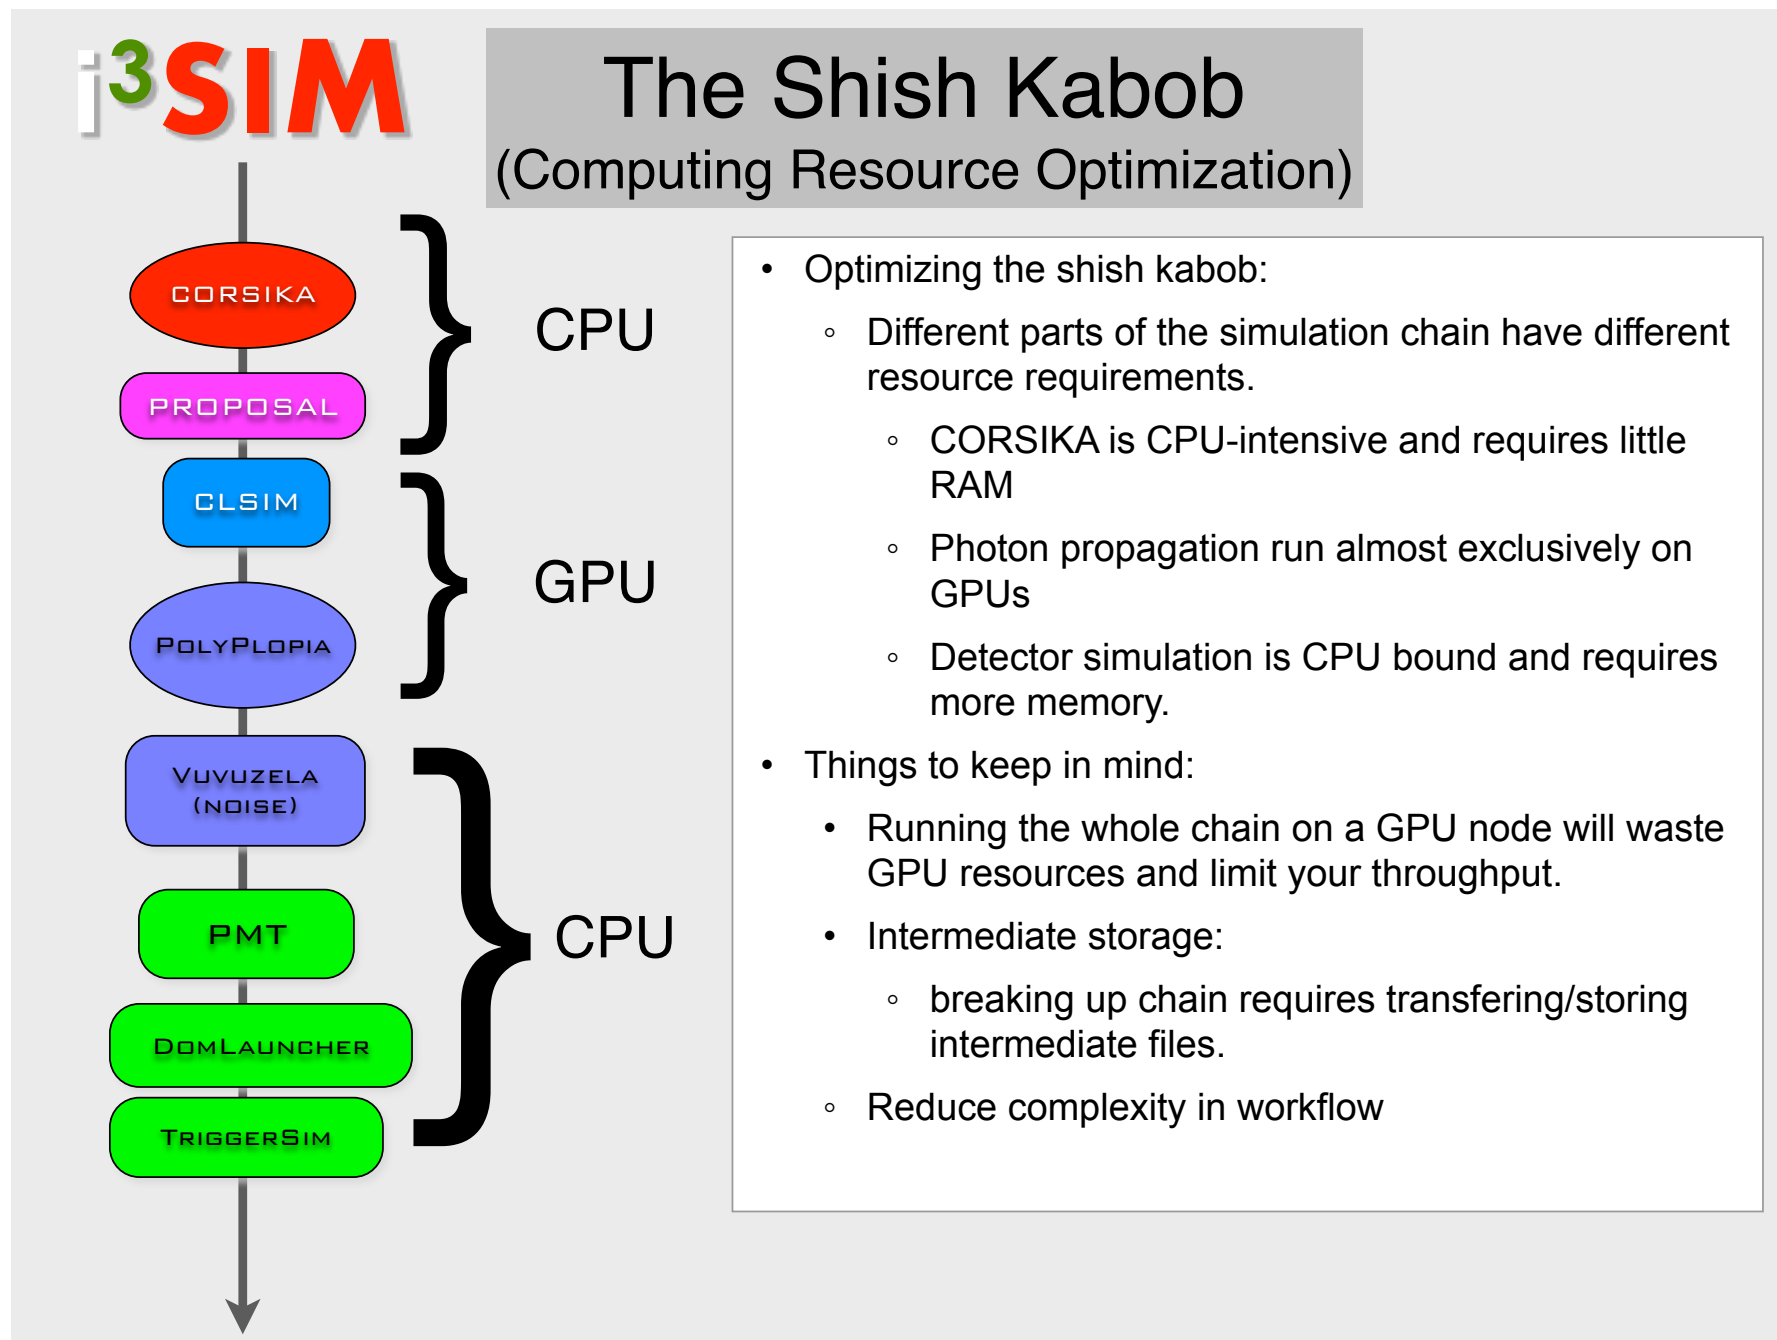
No current plans to develop noise parameterization (as we did for IceCube DOMs). Instead generate GEANT4 noise simulations of radioactive decay and glass scintillation on the PTM and pressure sphere glass. The noise module would then just sample hits from these simulations.

Tasks: Lots of work to do. See task list in backup slides

We are incorporating upgrade simulation/software into general software calls. Alternating on topics.



SCAP Recommendation 2018-9

Efforts should continue in a highly focused manner in order to maintain workflows which can run efficiently on systems where IceCube can request resources. This requires work on the workflows themselves, but also on monitoring and job scheduling and on the handling of intermediate results.




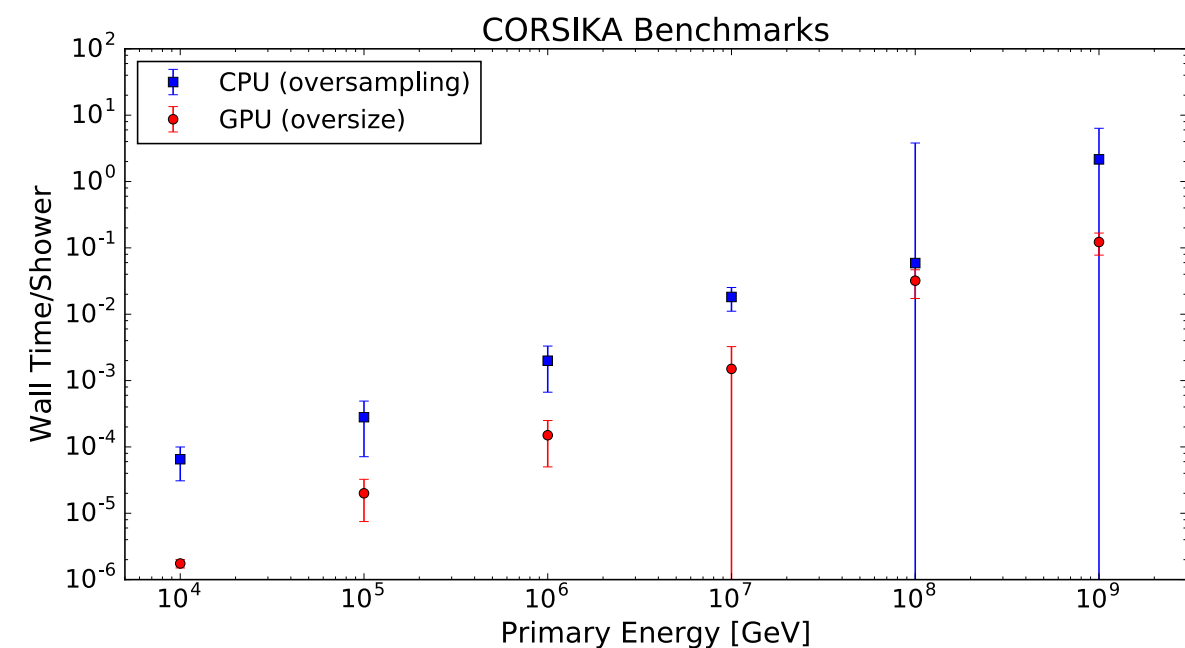
Resource Utilization

- Use of opportunistic resources has increased
- Combination of CPU-intensive datasets and GPU intensive datasets seems to utilize resources maximally.
 - Simpler workflows better than complex DAGs
 - Longer tasks (but not too long)
 - Reduce intermediate storage requirements, improved data compression (zstd)



CORSIKA Demands and Optimization Schemes


- Strategies such as event oversampling and DOM oversizing can speed up performance by large factors
- Systematic effects and GPU efficiency currently prevent us from taking full advantage of DOM oversizing.



| | Original | Optimization | Improvement Factor |
|----------|----------|--------------|--------------------|
| CPU/year | 14714 | 3516 | 4.18 |
| GPU/year | 4451 | 336 | 13.25 |

With optimizations, 1 yr/yr is achievable...


- If you exclude signal/systematics simulation
- No software bugs fixes/updates needed



Optimization Schemes

Oversizing speeds up photon propagation but with fast GPUs this actually leads to an underutilization of GPU since threads can't be fed data fast enough and large memory usage resulting from frame buffering.

Currently, most simulation use `oversize=1` (except at HE)



Possible to simulate direct hole ice propagation given that we are doing `oversize=1`.



GPU efficiency:

One way to improve GPU utilization is to share a single CISim server with multiple clients. One complication that arises in this model is a lack of determinism from random number generators for CISim (multi-client) determinism

Jakob's current solution is Random123 with an efficiency impact 40% but other possibilities include mall state non-cryptographic RNG.



Revisit DOM oversizing.

Introduces systematic errors in time distributions.




Benchmarks and Optimization of Datasets

- One of the issues with benchmarking and optimizing dataset configurations for GPU performance is the variance in GPU models.
- Datasets that perform well on some models have a much worse performance on others.
- Proposed solution for GPU utilization prediction:
 - Define a scaling factor for each model
 - Pyglidein runs benchmarks on GPUs to determine this ahead of time.
 - Check GPU model, refer to scaling factor.



Example: **SnowStorm** (Eric Ganster):


optimized configuration on a subset of GPUs but efficiency varies widely depending on the GPU model.



Predictive Resource Scheduling

Zeroth order optimization: Energy Binning

- The total GPU runtime depends on the total deposited charge
- The total deposited charge is *weakly* dependent on the primary energy
- A current approach is to roughly estimate the maximum runtime of each job if we break the spectrum into energy bins



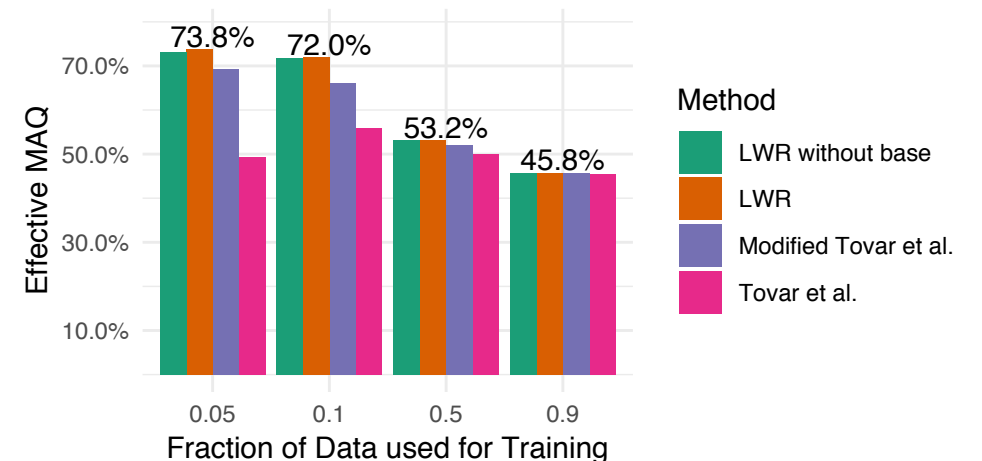


Predictive Resource Scheduling

Carl Witt <wittcarx@informatik.hu-berlin.de>

Jakob van Santen <jakob.van.santen@desy.de>

Ulf Leser <leser@informatik.hu-berlin.de>

- Low-wastage regression can improve memory allocation quality for IceProd jobs by nearly 50%.
- Largest improvement when memory requirement can be predicted from upstream tasks
- **Black-box, online method:** no knowledge of the task content or initial benchmarking needed
- Next steps:
 - Present at HPCS 2019
 - Implement requirement prediction in IceProd2 (who, when?)
 - Gather more log data from newer IceProd2 releases (memory use wasn't collected for nearly a year)
 - Investigate predictions based on dataset config (i.e. meta project version, generator, number of events, energy range, etc)



Simulation Coordination

- Role of Simulation Coordination shifting from away from a centrally managed production
- New aim is to coordinate WG-driven production:
 - Storage allocation
 - CPU/GPU utilization (adjust priorities)
 - Coordination between WG needs.
 - Large-scale common productions such as CORSIKA and MuonGun are still centrally managed.
- Simulation Requests (resource allocation)
 - Reviewed by tech leads.

Simulation Requests

Show completed requests

| | |
|---|-------------|
| LeptonInjector - SnowStorm | Not Started |
| GENIE simulations | In Progress |
| Muon Gun: single muons | In Progress |
| Complete old IceTop simulations | In Progress |
| Process AWS test to L2 | In Progress |
| Gen2 CORSIKA simulations | Testing |
| GENIE Systematics variations | Testing |
| Re-filter 2012 CORSIKA (2016) | Partial |
| LE CORSIKA Systematics | Complete |
| NuGen Systematics Production | Complete |

Datasets

IT81:2012

[CORISKA-ice-top](#) (8)

[CORISKA](#) (13)

IC86:2016

[neutrino-generator](#) (25)

[MuonGun](#) (5)

“generalized” set with static cylinder injection surface. This is very expensive computationally, but everyone can use it. Will take highest priority in simprod once NuGen sample is done. No plans to store photons since the memory requirement is too expensive.

- Working Group: Diffuse
- Status: In Progress
- Priority: 3
- Submitted: 2020-04-18, Completed: N/A
- Requester: Manuel Silva
- Analysis: ESTES/MESE/GENERAL
- Documentation: <https://wiki.icecube.wisc.edu/index.php/ESTES>

Comments

Stats

| Dataset | Status | Completed | Processing | Errors | Progress |
|---|--|-----------|------------|--------|---|
| 21315 | complete | 15000 | 0 | 0 | 100.00% |
| Level 2 (pass2) IC86.2016 MuonGun with weighted spectrum of E^{-2} , using spice_3.2.1 CLSim with scat=+0,abs=+0 with with angsens_unified p0=0.0 p1=0.0 with domeff=1.0. Energy range of $1e5 \text{ GeV} < E_{\mu} < 1e9 \text{ GeV}$. | | | | | |
| 21316 | errors | 39995 | 0 | 5 | 99.99% |
| Level 2 (pass2) IC86.2016 MuonGun with weighted spectrum of E^{-5} , using spice_3.2.1 CLSim with scat=+0,abs=+0 with with angsens_unified p0=0.0 p1=0.0 with domeff=1.0. Energy range of $1e4 \text{ GeV} < E_{\mu} < 1e5 \text{ GeV}$. | | | | | |
| 21317 | errors | 19994 | 0 | 6 | 99.97% |
| Level 2 (pass2) IC86.2016 MuonGun with weighted spectrum of $E^{-4.5}$, using spice_3.2.1 CLSim with scat=+0,abs=+0 with with angsens_unified p0=0.0 p1=0.0 with domeff=1.0. Energy range of $5e3 \text{ GeV} < E_{\mu} < 1e4 \text{ GeV}$. | | | | | |
| 21318 | processing | 99975 | 1 | 24 | 99.98% |
| Level 2 (pass2) IC86.2016 MuonGun with weighted spectrum of $E^{-4.5}$, using spice_3.2.1 CLSim with scat=+0,abs=+0 with with angsens_unified p0=0.0 p1=0.0 with domeff=1.0. Energy range of $1e3 \text{ GeV} < E_{\mu} < 5e3 \text{ GeV}$. | | | | | |
| 21319 | processing | 34157 | 17914 | 7929 | 56.93% |
| Level 2 (pass2) IC86.2016 MuonGun with weighted spectrum of $E^{-4.5}$, using spice_3.2.1 CLSim with scat=+0,abs=+0 with with angsens_unified p0=0.0 p1=0.0 with domeff=1.0. Energy range of $7e2 \text{ GeV} < E_{\mu} < 1e3 \text{ GeV}$. | | | | | |

Sim-Prod Requests



[Home](#)

[Simulation Request Form-](#)

[SmartSheets](#)

[Wiki](#)

[Mad-Dash](#)

[Grafana Monitor](#)

[Weighting](#)

Simulation Request Form

Simulation Dataset Request

Is simulation data you need for your analysis missing? Place your request and we'll see if data files already exist or if it needs to be produced.

***Obligatorio**

Dirección de correo electrónico *

Tu dirección de correo electrónico

Your name *

Tu respuesta

Your status *

Undergraduate student

Sim-Prod Requests



- Home
- Simulation Request Form
- SmartSheets-
- Wiki
- Mad-Dash
- Grafana Monitor
- Weighting

Simulation Request Smartsheets

Simulation Requests (View Only) : **smartsheet** Report Abuse Help


Card View Filter All Levels View by Status

| Testing (2) | In Progress (7) | Partial (1) | Complete (7) | Withdrawn (0) |
|---|---|--|---|---------------|
| <p>Gen2 CORSIKA simulations L Ithulasidhar@wisc.edu Gen2 https://wiki.icecube.wisc.edu...</p> | <p>Complete old IceTop simulati... D Dennis Soldin Cosmic Rays https://wiki.icecube.wisc.edu...</p> | <p>Re-filter 2012 CORSIKA (2016) Juan Carlos Díaz Vélez GENERAL</p> | <p>NuGen Systematics Producti... MS Manuel Silva Diffuse https://wiki.icecube.wisc.edu...</p> | |
| <p>CORSIKA for Moon analysis thorsten.gluesenkamp@fau.de Systematics/Reconstruction</p> | <p>GENIE Systematics variations G grenzi@icecube.wisc.edu BSM https://wiki.icecube.wisc.edu...</p> | | <p>LE CORSIKA Systematics D david.gillcrist@icecube.w... Calibration</p> | |
| | <p>IC79 Pass2 NuGen Simulations ML Michael Larson Neutrino Sources</p> | | <p>GENIE simulations G BSM https://wiki.icecube.wisc.edu...</p> | |
| | <p><i>Reprocess Detector+Filter with correct...</i> Process AWS test to L2 2020 MS Manuel Silva GENERAL</p> | | <p>Muon Gun: single muons MS Manuel Silva Diffuse https://wiki.icecube.wisc.edu...</p> | |
| | <p><i>Reprocess Detector+Filter wi</i></p> | | | |

Backup



Photon Propagation

- μ energy lost + cascades \rightarrow photons \rightarrow p.e.
- Photon propagation : ice properties + PMT response + DOM glass/gel



- Direct photon tracking (CLSim, PPC)
- Pre-generated lookup splined table :
 - I3PhotonicsHitMaker
 - Amplitude and time distribution
- Hybrid photon tracking
 - HitMaker + CLSim

- PPC originally written for CUDA libraries
- CLSim written for OpenCL
- PPC ported to OpenCL
- CLSim ported to CUDA
- New emerging technologies:
 - HIP, SyCL, OneAPI, C++20/23

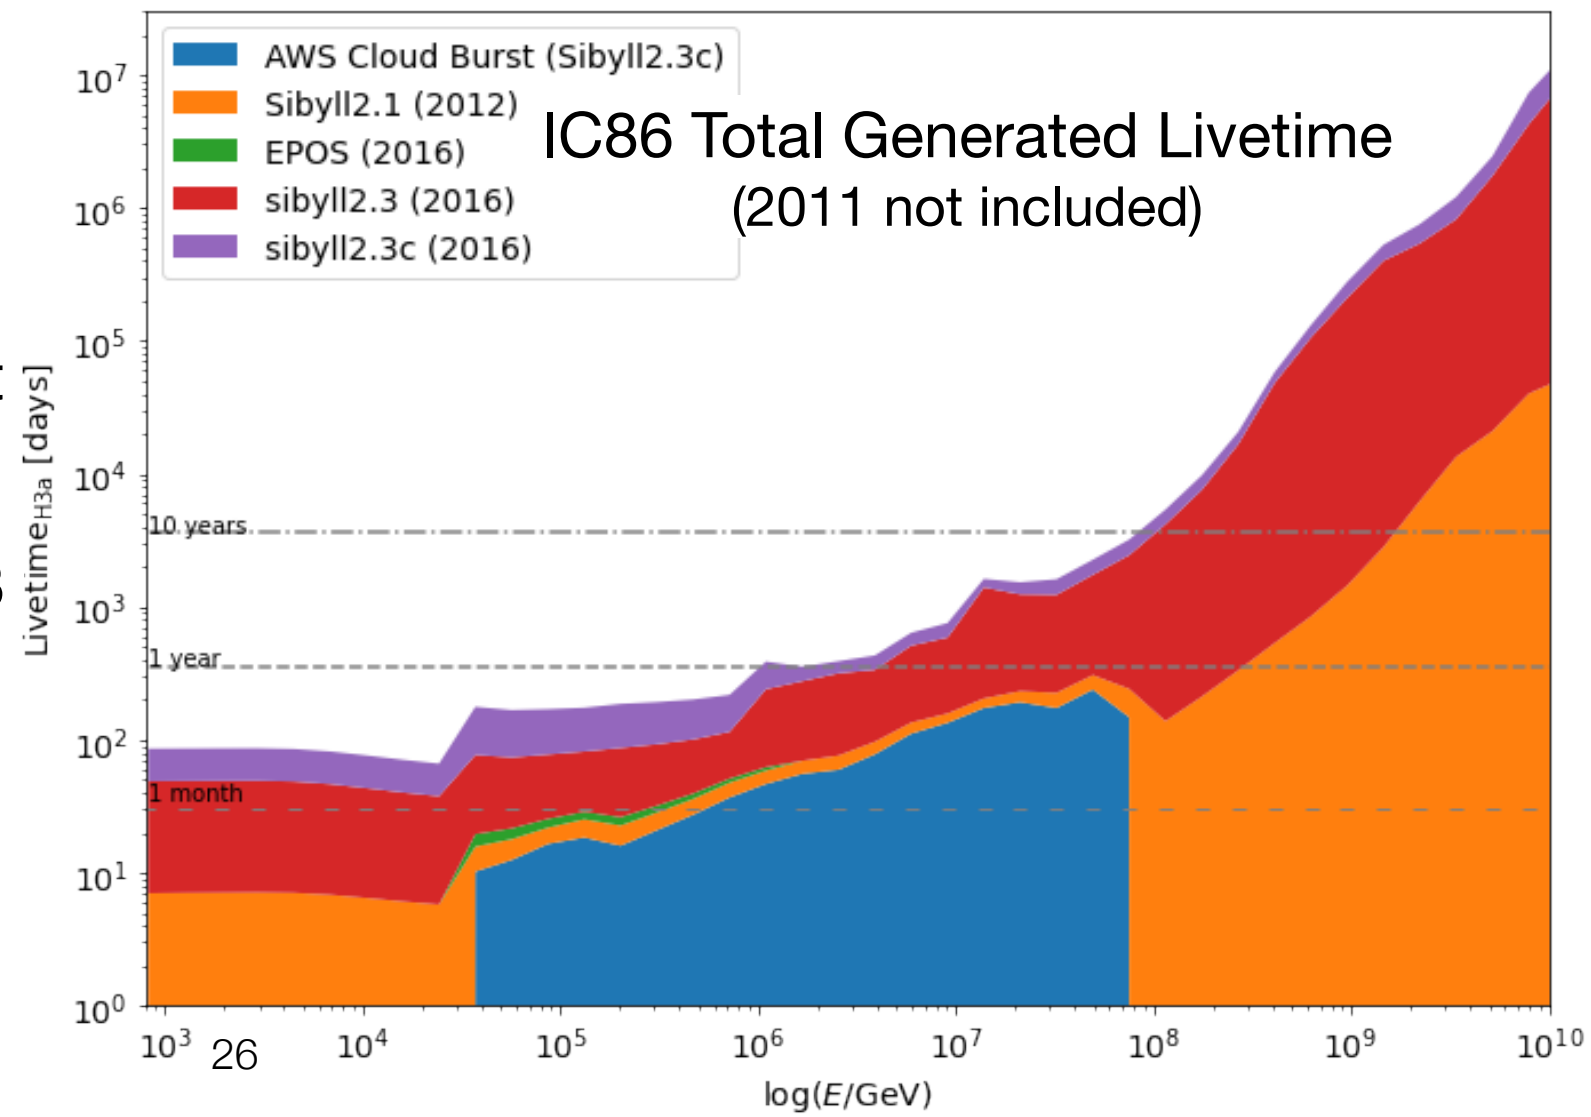


CORSIKA Livetime

- A number of older datasets have been replaced as a result of previously reported issues.
- 2012 datasets are currently being reprocessed with pass2
- 2011 datasets will also be reprocessed with pass2.
 - Much of this is L3 CSD.
- Datasets are distributed on different sites (DESY, CEPH, etc.)
- Pass2 & L3 processing (output to I3 data warehouse)
- New MuonGun production

Oversampling in CORSIKA

- Can accumulate statistics much faster.
- But events are not statistically independent




Simulation Survey

Juan Carlos Díaz Vélez

Science and Computing Advisory Panel


Which working groups do you belong to?

10 respuestas




Which years of data are you working on? (check all that apply)

9 respuestas



Is simulation data crucial for your analysis?

10 respuestas




- Yes, definitely
- No, not at all
- Would be good to have some simulation

Simulation Production data: current usage

Do you use data from simulation production?


10 respuestas



- Yes, I use simulation production data only
- Yes, I use simulation production data AND my own simulation
- No, I use my own simulation
- I don't use any simulation


Does simulation production data satisfy your needs?

9 respuestas




If you (your working group for you) ever requested datasets from Simulation Production, please check all that apply

2 respuestas




Which types of SIGNAL datasets do you need?

8 respuestas




Which years do you need?

6 respuestas




If you are working with coincident background, do you think it would be sufficient to use MuonGun (single muons) instead of CORSIKA?

5 respuestas




Which type of BACKGROUND datasets do you need for your analysis?

6 respuestas




Which ice model would you like to use?

6 respuestas



Which types of SYSTEMATIC datasets do you need for your analysis?

7 respuestas




Please specify any other datasets that would be useful for you that were not covered above

0 respuestas

Todavía no hay respuestas para esta pregunta.

Would you like SimProd to take care of the production of any of the above datasets? (we will contact you if you check yes, so please make sure you stated your name in the beginning of this form)

4 respuestas



Please specify the properties of your dataset(s) (signal/background, flavor, energy, etc.)

4 respuestas

Nutau Analysis: Muongun background (medium energy)

EHE Analysis (work in progress): Very high energy NuE, NuMu, NuTau (+ eventually E^{-1} CORSIKA)

Dark matter signal from the center of the Earth, all flavors, energy depending on the dark matter mass (<10 TeV in any case)

I produce the low energy muongun targeted for deepcore. (This is the background for our analysis/WG. Josh Hignight previously and Wing Ma currently produce the genie simulation for our analysis/WG.)

signal simulation for realtime analyses

Do you use any simplifications/tricks to speed up your simulation? Please specify.

3 respuestas

Nutau Analysis: Throw away a fraction of events after muon propagation based on energy losses

EHE Analysis: Some approximations + table based photon propagation

I use the star pattern to simulate the event instead of the detector physical location

Yes, we use the inner target cylinder feature in MuonGun to target the DeepCore volume. We also use a KDE I developed to target the portions of energy/coszen phase space that are most likely to survive our event selection and be present in our final level sample.
http://code.icecube.wisc.edu/svn/sandbox/kleonard/kde_filter/

If you checked "I had special requirements that couldn't be handled by SimProd." above, please specify.

3 respuestas

See the speedup section above, let me know if you have any further questions about this

LeptonInjector

There is lot of optimization involved in choosing the desired target cylinder settings and energy/coszen KDE settings for DeepCore, so we (Andrii Terliuk previously, but now myself) had to learn how to produce MuonGun so we could create small test sets to explore these parameters and determine what was appropriate for our analysis. Then once we had picked our settings, it was easy enough to just plug them back in and launch the production of a full set. My understanding is that SimProd could help with the last part (making the final set once we know the desired settings), but by that point most of the groundwork was already complete so it was easier to just finish it ourselves.

Let us know what you think!

4 respuestas

The wikipage https://wiki.icecube.wisc.edu/index.php/Simulation_Production#Datasets can be more detailed and updated. For example, dataset 11820 can't be located on cobalt.

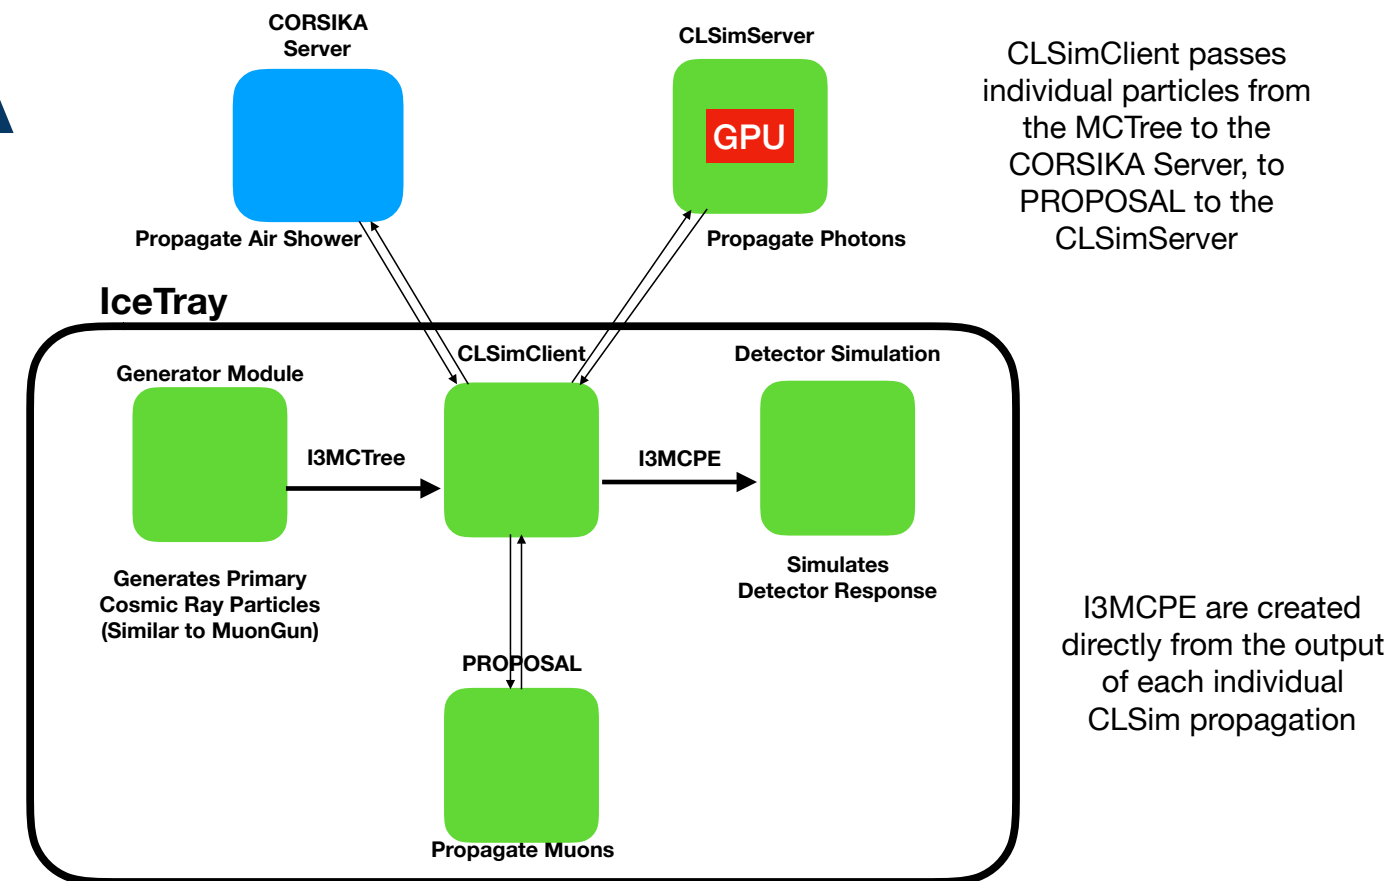
Thanks!

I am not sure if I qualify for this feedback.

can we have filters on iceprod2 website to select a specific generator or generation year? with old ice-rod it was easy, now it's hard to find what I want!

DYNSTACK CORSIKA

Kevin Meagher & Jakob van Santen



- Replaces CORSIKA's post-reaction particle stack with a C++11 plugin
- General API for doing things like the neutrino kill threshold, plus helpful extras (take configuration from the steering card, manipulate event headers/trailers, etc)
- In mainline CORSIKA since 7.56 (modulo typos)
- Write plugins in C++11 without touching corsika.F, depend only on the standard library
- Build a better CORSIKA for in-ice background simulation.
- Reduce memory and disk requirements for high energy simulations.

Analysis-specific, targeted background simulation

MuonGun (IceCube implementation of MUPAGE)

arXiv:0907.5563v1 [astro-ph.IM] 31 Jul 2009

PROCEEDINGS OF THE 31st ICRC, ŁÓDŹ 2009

Atmospheric MUons from Parametric formulas: a fast GENERator for neutrino telescopes (MUPAGE)

M. Bazzotti*, S. Biagi*[†], G. Carminati*[†], S. Cecchini*[‡],
T. Chiarusi[†], A. Margiotta*[†], M. Sioli*[†] and M. Spurio*[†]

*Dipartimento di Fisica dell'Università di Bologna, Viale Bertini Pichat 6/2, 40127 Bologna, Italy

[†] INFN, Sezione di Bologna, Viale Bertini Pichat 6/2, 40127 Bologna, Italy

[‡] INAF-IASF, Via Gobetti 101, 40129 Bologna, Italy





Fig. 1: Sketch of some input parameters. The cylinder surrounding the instrumental volume is the *can*, with radius R_{can} and height H_{can} . The events are generated on an extended can with R_{ext} . The origin of the coordinate system does not have to be located at the center of the detector. The lower disk is at a depth H_{max} with respect to the sea/ice surface.

MuonGun - Software is stable but parameterizations need to be updated. Re-parametrization with SIBYLL 2.3c. This requires a large volume of CORSIKA simulation (no photon propagation). In the past this was accumulated over many datasets in f2k format.

- Current parametrization based on SYBILL 2.1
- MuonGun parametrization needs to be updated with new hadronic interaction models.
- This requires quite a bit of CORSIKA which in the past was produced and saved in f2k format as a side effect of standard production.
- We are in the process of producing CORSIKA showers in I3 format as a backfill on CPU nodes.

Single Photo-Electron Templates

Spencer N. Axani

saxani@mit.edu

SPE Templates:

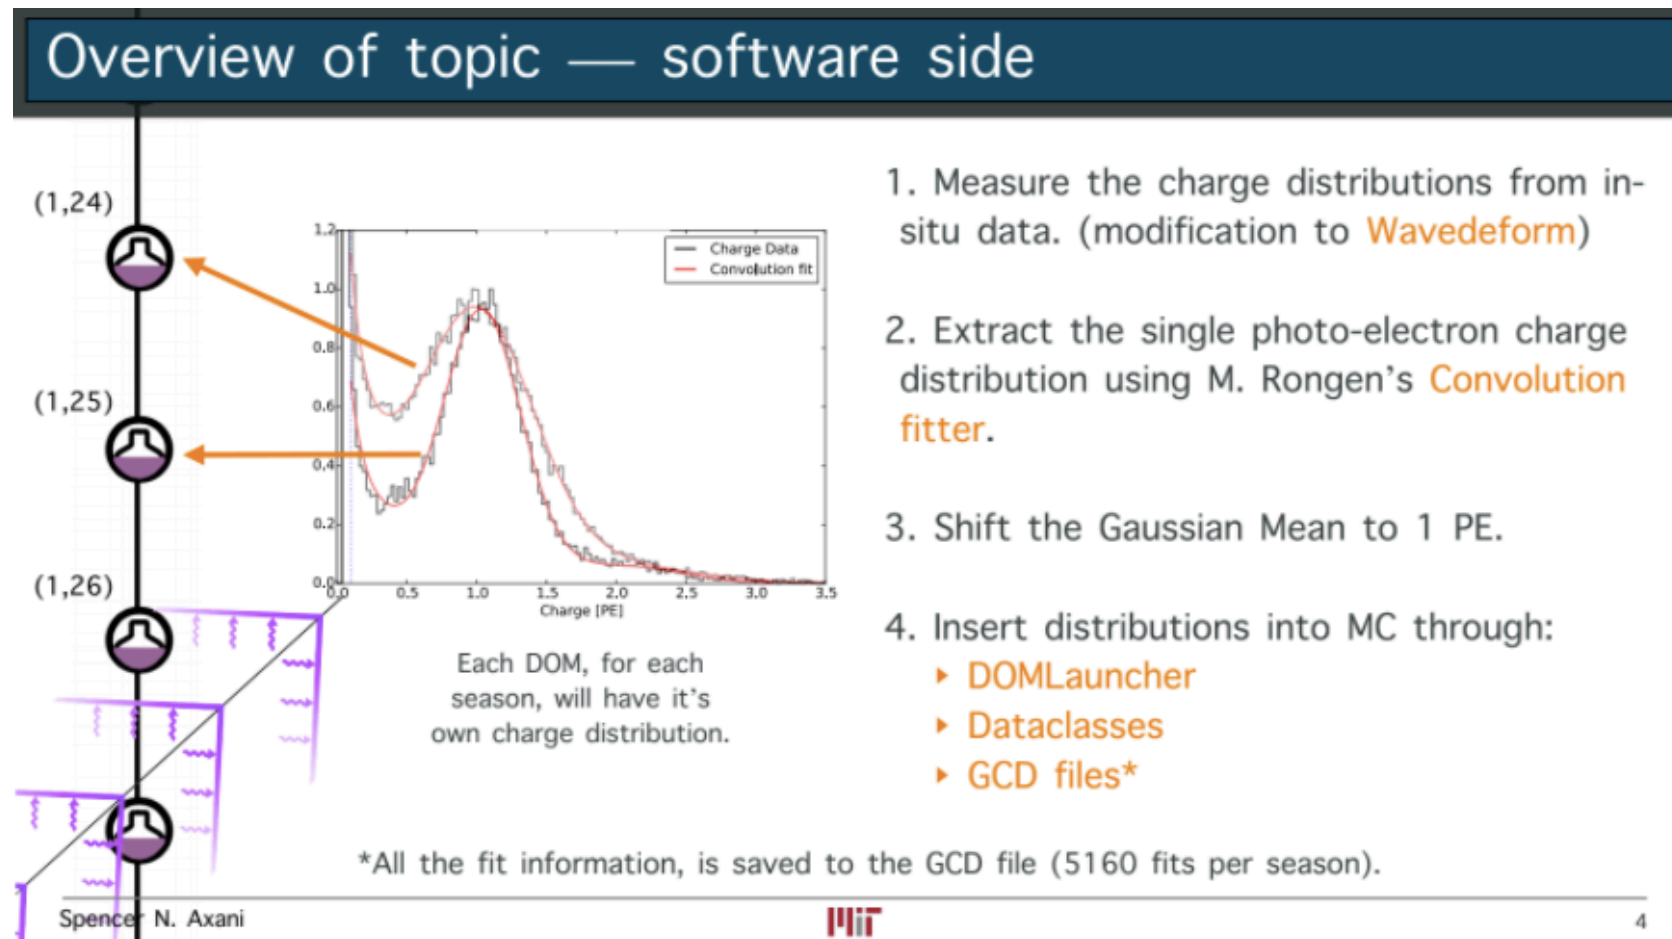
The individual in-ice DOM single photon charge distributions.

The SPE Templates are now used in our simulation.

We expect that including the SPE Templates will introduce a shift in the efficiency of $\sim 130\%$.

Code has been reviewed and approved and is currently part of the standard simulation.

Updated noise parameters in Vuvuzela due to change.



1. Measure the charge distributions from in-situ data. (modification to **Wavedeform**)
2. Extract the single photo-electron charge distribution using M. Rongen's **Convolution fitter**.
3. Shift the Gaussian Mean to 1 PE.
4. Insert distributions into MC through:
 - ▶ **DOMLauncher**
 - ▶ **Dataclasses**
 - ▶ **GCD files***

Upgrade/Gen2

DOM Simulation:

PMTResponse uses PTM base class.

Specific PTM Classes implemented:

- HamamatsuR7081_02PMT (IceCube IceTop PDOM)
- HamamatsuR5912_100PMT (DEgg)
- HamamatsuR15458_02PMT (mDOM)

DOM Launcher: electronics to be implemented in the future. Separate module. Readout electronics have been under development. Currently very simple placeholder in simulation for all new sensors:

Merge pulses within 10ns.

10 ns pulse width for all pulses.

Realistic DOM frontend simulation and pulse reconstruction under development (Leander Fischer).

There has also been some development of GEANT4 OM models:

mDOM (Münster), DEgg (Chiba), and a generic system for many OM types (Nahee Park).

Tasks: Lots of work to do. See task list in backup slides

We are incorporating upgrade simulation/software into general software calls. Alternating on topics.

Parallelization

- Our current model is already optimally parallel for CPU performance.
- We basically have multiple instances of IceTray running (1 per core) on a node (i.e. *embarrassingly parallel*)
- Any amount of parallelization involving inter-process communication and locking will reduce performance.
- There are some non-CPU performance reasons to parallelize, namely RAM limitations.
 - You could reduce memory utilization by sharing tables, splines, etc.
- One of the limitations in using threads in IceTray is the Global Interpreter Lock (GIL) in boost-python.
- This is called whenever entering Python and can lead to deadlock when used in combination with thread locks.
- Another issue is determinism from calls to random number generators. This can be addressed by using parametrized RNGs that can produce very large number of independent threads and assign a separate thread to each atomic computing task.

Parallelization

- **Event-level parallelism**

- This is essentially what we are doing now.
- You could improve things by sharing memory between threads (or processes prior to fork()).
- Driver module reads input file and dispatches events to workers round-robin.
- Nathan wrote an implementation of this a few years ago:
 - <http://code.icecube.wisc.edu/svn/sandbox/nwhitehorn/i3mpi>


- **Module-level parallelism**

- One module per thread/process. Requires some work in organizing break up for load balancing.
- Assigns modules (or group of modules) to different threads/processes
- Load balance could be improved by using thread pools. But this requires being able to move modules from one thread to another.

- **Intra-module parallelism**

- Requires direct changes to each module that utilizes non-negligible CPU
- Otherwise, performance would drop if threads are waiting idle for work.

flow of experimental and simulation data



Upgrade Task List

IceCube Upgrade simulation, reconstruction and physics task list.

Contents [\[hide\]](#)

- 1 [Simulation](#)
 - 1.1 [Detector simulation](#)
 - 1.2 [Calibration device simulation](#)
- 2 [Data processing](#)
- 3 [Reconstruction](#)
- 4 [Physics](#)

Simulation

Detector simulation

Maintainer: Tom Stuttard

| Category | Task name | Task description | Status | Task owner(s) |
|----------|----------------------|--|--|-------------------------------|
| mDOM | GEANT4 model | Detailed GEANT4 model of OM, including geometry, materials (including glass, gel). Supports photon or decay simulations. | Preliminary model in place. | Alexander Kappes, Nahee Park? |
| mDOM | GEANT4 decays | Implement radioactive decays in OM glass (based on lab measurements) in GEANT4 models. | Preliminary model in place | Martin Unland |
| mDOM | Production sim model | Derive OM photon acceptance curves (or some other model) at PMT, vs wavelength, incidence position, incidence angle). Parameterised from GEANT4 models & lab tests. Not required if directly use GEANT4 model in production simulations. | Preliminary geometric angular acceptance model + wavelength dependence implemented. Out-of-date (not derived from latest GEANT4 models). | TBC |
| mDOM | PMT model | Model response of PMT to incident photons. Aim to have a single model (based on existing PMTResponseSimulator) that supports all PMTs, steered using datasheet values and lab measurements. | PMTResponseSimulator updated to support IceCube, DEgg and mDOM PMTs (in combo trunk). | Matti Jansson |
| mDOM | SPE template | Distribution of charge observed for a single photoelectron. Expected that IceCube SPE template code can be re-used, tuned to lab measurements. | Not started | TBC |

Upgrade Task List (continued)

| | | | | |
|------------|-----------------------------|--|--|-------------------------------|
| mDOM | Noise model (parameterised) | Tune vuvuzela parameters to match lab measurements and GEANT4 simulations. Derive new dt parameterisation to account for correlated noise between different PMTs on multi-PMT OMs. | Parameterisation tuned to GEANT4 model, but does NOT include correlations between PMTs. Currently choosing pre-generated noise model as baseline. | Martin Unland, Michael Larson |
| mDOM | Noise model (pre-generated) | Produce files containing noise hits using GEANT4 model with radioactive decays | Done | Martin Unland |
| mDOM | mDOM readout sim / reco | Simulate readout electronics. Signal shaping -> slow (100 MSPS) ADC sampling -> discriminator -> leading/trailing edge timestamp using fast (1 GSPS) sampler -> simple vs complex waveform selection -> onboard charge extraction (simple waveform only) -> data packet formation & transmission | Preliminary work started, but requires ~final readout design before simulation/reconstruction work can be finalized (decision on readout expected summer 2020). Code is here . | Leander Fischer |
| DEgg | GEANT4 model | See entry for mDOM. | Preliminary model in place. | Ken'ichi, Lu Lu |
| DEgg | GEANT4 decays | See entry for mDOM. | Not started | Ken-ichi |
| DEgg | Production sim model | See entry for mDOM. | Very simple placeholder (two IceCube OMs). | TBC |
| DEgg | PMT model | See entry for mDOM. | Initial discussions between experts started | Wing Ma? |
| DEgg | SPE template | See entry for mDOM. | Not started | TBC |
| DEgg | Noise model (parameterised) | See entry for mDOM. | Parameterisation tuned to GEANT4 model, but does NOT include correlations between PMTs. Prioritising pre-generated noise model as baseline. | TBC |
| DEgg | Noise model (pre-generated) | See entry for mDOM. | Not started | Ken'ichi |
| DEgg | DEgg readout sim | Simulate readout electronics. Signal shaping -> ADC sampling (250 MSPS) -> TDC firmware triggering/discriminating -> simple vs complex waveform selection -> onboard charge extraction (simple waveform only) -> data packet formation & transmission | Not started | Ryo? |
| Triggering | DAQ/Trigger sim | Simulate data aggregation from all OMs at the surface (include IC86 OMs) and triggering. Include handling the fact that there is no hardware HLC in the Upgrade OMs. | TBC | Alex Olivas |

Upgrade Task List (continued)

| | | | | |
|-----------|---------------------------------|--|--|-----------------------|
| Framework | Correlated noise in vuvuzela | Add functionality to handle correlated noise between neighbouring PMTs/OMs (e.g. from a single decay). | Complete (although do not yet have parameterisations that utilise this functionality) | Michael Larson |
| Framework | Pre-generated noise in vuvuzela | Generate noise in vuvuzela by sampling from pre-generated noise files (produced from GEANT4 simulations). | Complete | Michael Larson |
| Framework | GEANT4 models in production sim | Implement use of GEANT4 OM models directly in production simulation. Likely by propagating photons (CLSim/ppc) to a bounding surface and storing, then separately running GEANT4. Result would be photoelectrons at photocathode surface, which would then be passed to the PMT simulator. | Not started | Nahee Park? Lu Lu? |
| Framework | OM geometry in CLSim/ppc | Update CLSim/ppc to handle non-spherical and variable radius OM shapes | Not started. Note though that a CLSim branch developed by Sebastian Fiedlschuster for hole ice and cable simulations has gone some way down this path, but is not synchronised with combo trunk. Need to coordinate with Sebastian here. | TBC |
| Framework | Snowstorm | Implement Snowstorm for efficiency production of systematics sets | Not started | TBC |
| Framework | GENIE reader | Switch from the deprecated genie-icetray to the new GENIE reader software for neutrino signal production. | Not started | TBC |
| Framework | simprod | Integrate the Upgrade MC submission scripts with simprod. | Not started | TBC |
| Framework | Update MC production scripts | Update the MC production scripts in line with the latest development to the DeepCore (oscNext) scripts. | Not started | TBC |