

High Throughput Computing in the IceCube Neutrino Observatory

Condor Week





May 1, 2008

UW-Madison

Juan Carlos Díaz Vélez

juancarlos@icecube.wisc.edu

<http://www.icecube.wisc.edu/~juancarlos>



IceCube

IceCube Collaboration

Bartol Research Inst, Univ of Delaware, USA
University of Alaska Anchorage, USA
Pennsylvania State University, USA
University of Wisconsin-Madison, USA
University of Wisconsin-River Falls, USA
LBNL, Berkeley, USA
UC Berkeley, USA
UC Irvine, USA

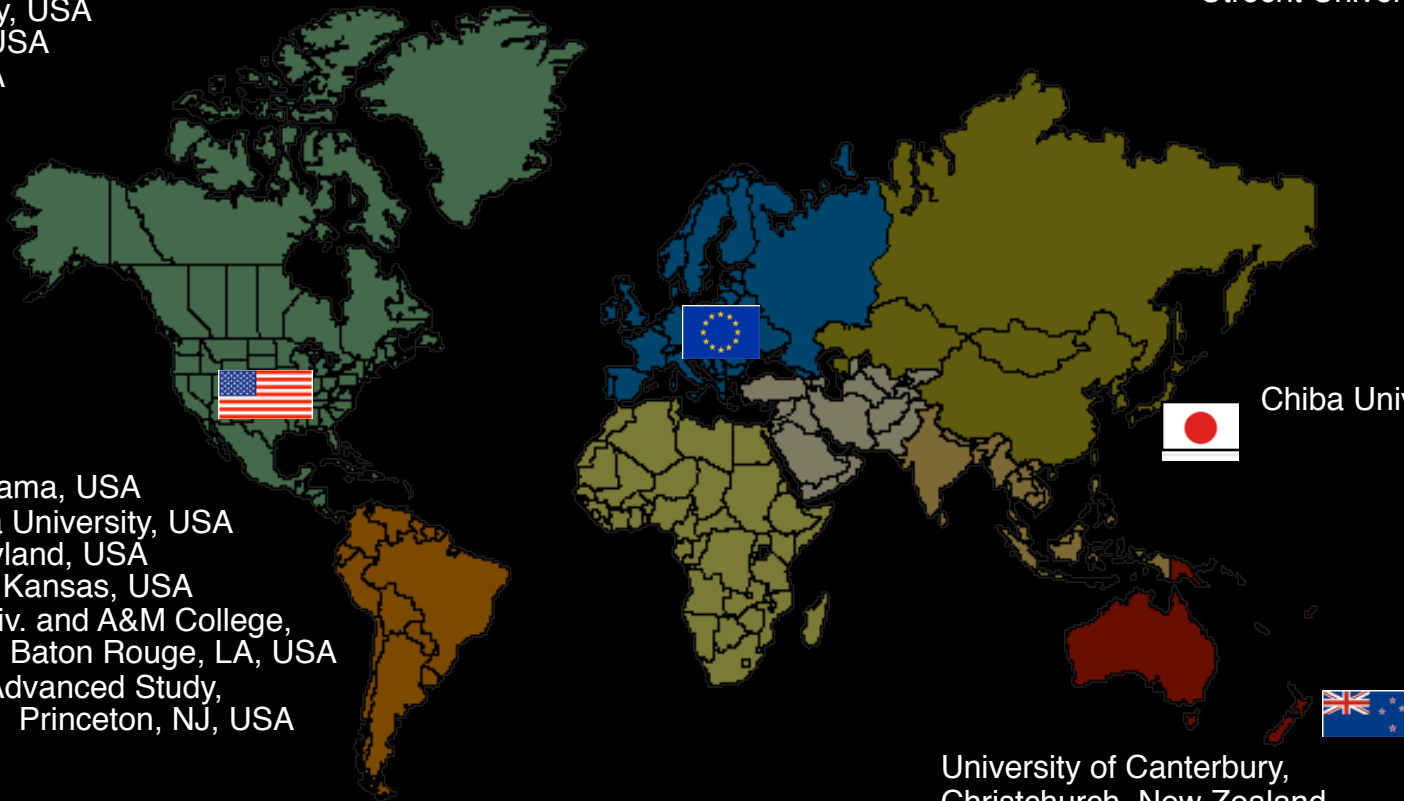
Université Libre de Bruxelles, Belgium
Vrije Universiteit Brussel, Belgium
Université de Mons-Hainaut, Belgium
Universiteit Gent, Belgium
Universität Mainz, Germany
DESY Zeuthen, Germany
Universität Wuppertal, Germany
Universität Dortmund, Germany

Humboldt Universität, Germany
Uppsala Universitet, Sweden
Stockholm Universitet, Sweden
Kalmar Universitet, Sweden
Imperial College, London, UK
University of Oxford, UK
Utrecht University, Netherlands


Univ. of Alabama, USA
Clark-Atlanta University, USA
Univ. of Maryland, USA
University of Kansas, USA
Southern Univ. and A&M College,
Baton Rouge, LA, USA
Institute for Advanced Study,
Princeton, NJ, USA

Chiba University, Japan

University of Canterbury,
Christchurch, New Zealand



what is a neutrino ?



ν

Condor Week
May 1, 2008
UW-Madison



- Proposed in 1930 by Pauli to explain missing energy in beta decay.
- Don't interact electro-magnetically.
- small mass ($m_e < 3 \text{ eV}$)
- Very small cross section for weak nuclear interactions
- Neutrinos come in three “flavors”
- mean free path (in lead) $\sim 1 \text{ ly}$

Standard Model: Leptons

e	μ	τ
ν_e	ν_μ	ν_τ

ν source candidates



ν source candidates



Photo credit: FORS Team, 8.2-meter VLT, ESO

supernova remnants




ν source candidates



Photo credit: FORS Team, 8.2-meter VLT, ESO

supernova remnants

AGNs



ν source candidates



Photo credit: FORS Team, 8.2-meter VLT, ESO

supernova remnants





Photo Credit: ESA/NASA/F. Mirabel (CEA)

microquasars


AGNs



unexpected sources



icecube



IceCube

AMANDA

South Pole

Dome


1500 m

2000 m

[not to scale]

Amundsen-Scott South Pole station

IceCube



IceTop: air shower array

- ✓ 80 Stations / 2 Tanks each
- ✓ 2 DOMs each per tank
- ✓ 125 m grid, 1 km² at 690 g/cm²
- ✓ $E_{\text{thres}} \sim 300 \text{ TeV}$ for ≥ 4 stations
- ✓ Useful rate up to $\sim \text{EeV}$




Digital Optical Module

AMANDA

IceCube: deep ice array

- 80 Strings / 60 DOMs each
- 17 m DOM spacing
- 125 m between strings
- 1 km³ instrumented

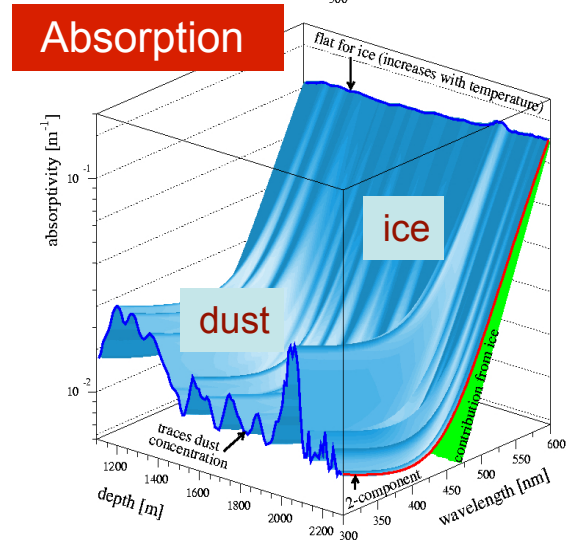
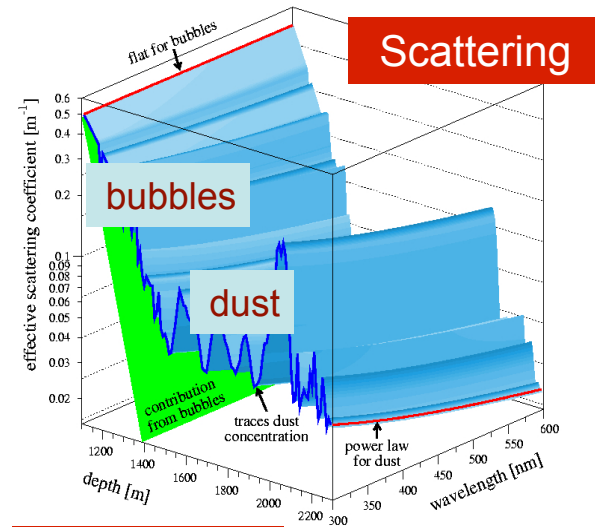






ice: the detector medium

Condor Week
 May 1, 2008
 UW-Madison

- Extremely low light level
 - Noise rate dominated by radioactive decay in glass
- Low absorption
- Significant scattering




optical sensor




Digital Optical Module

LED
flasher
board



main
board



HV board



- infrequently, a cosmic neutrino interacts with an atom in the medium and produces a nuclear reaction
- muon travels kilometers in ice/water

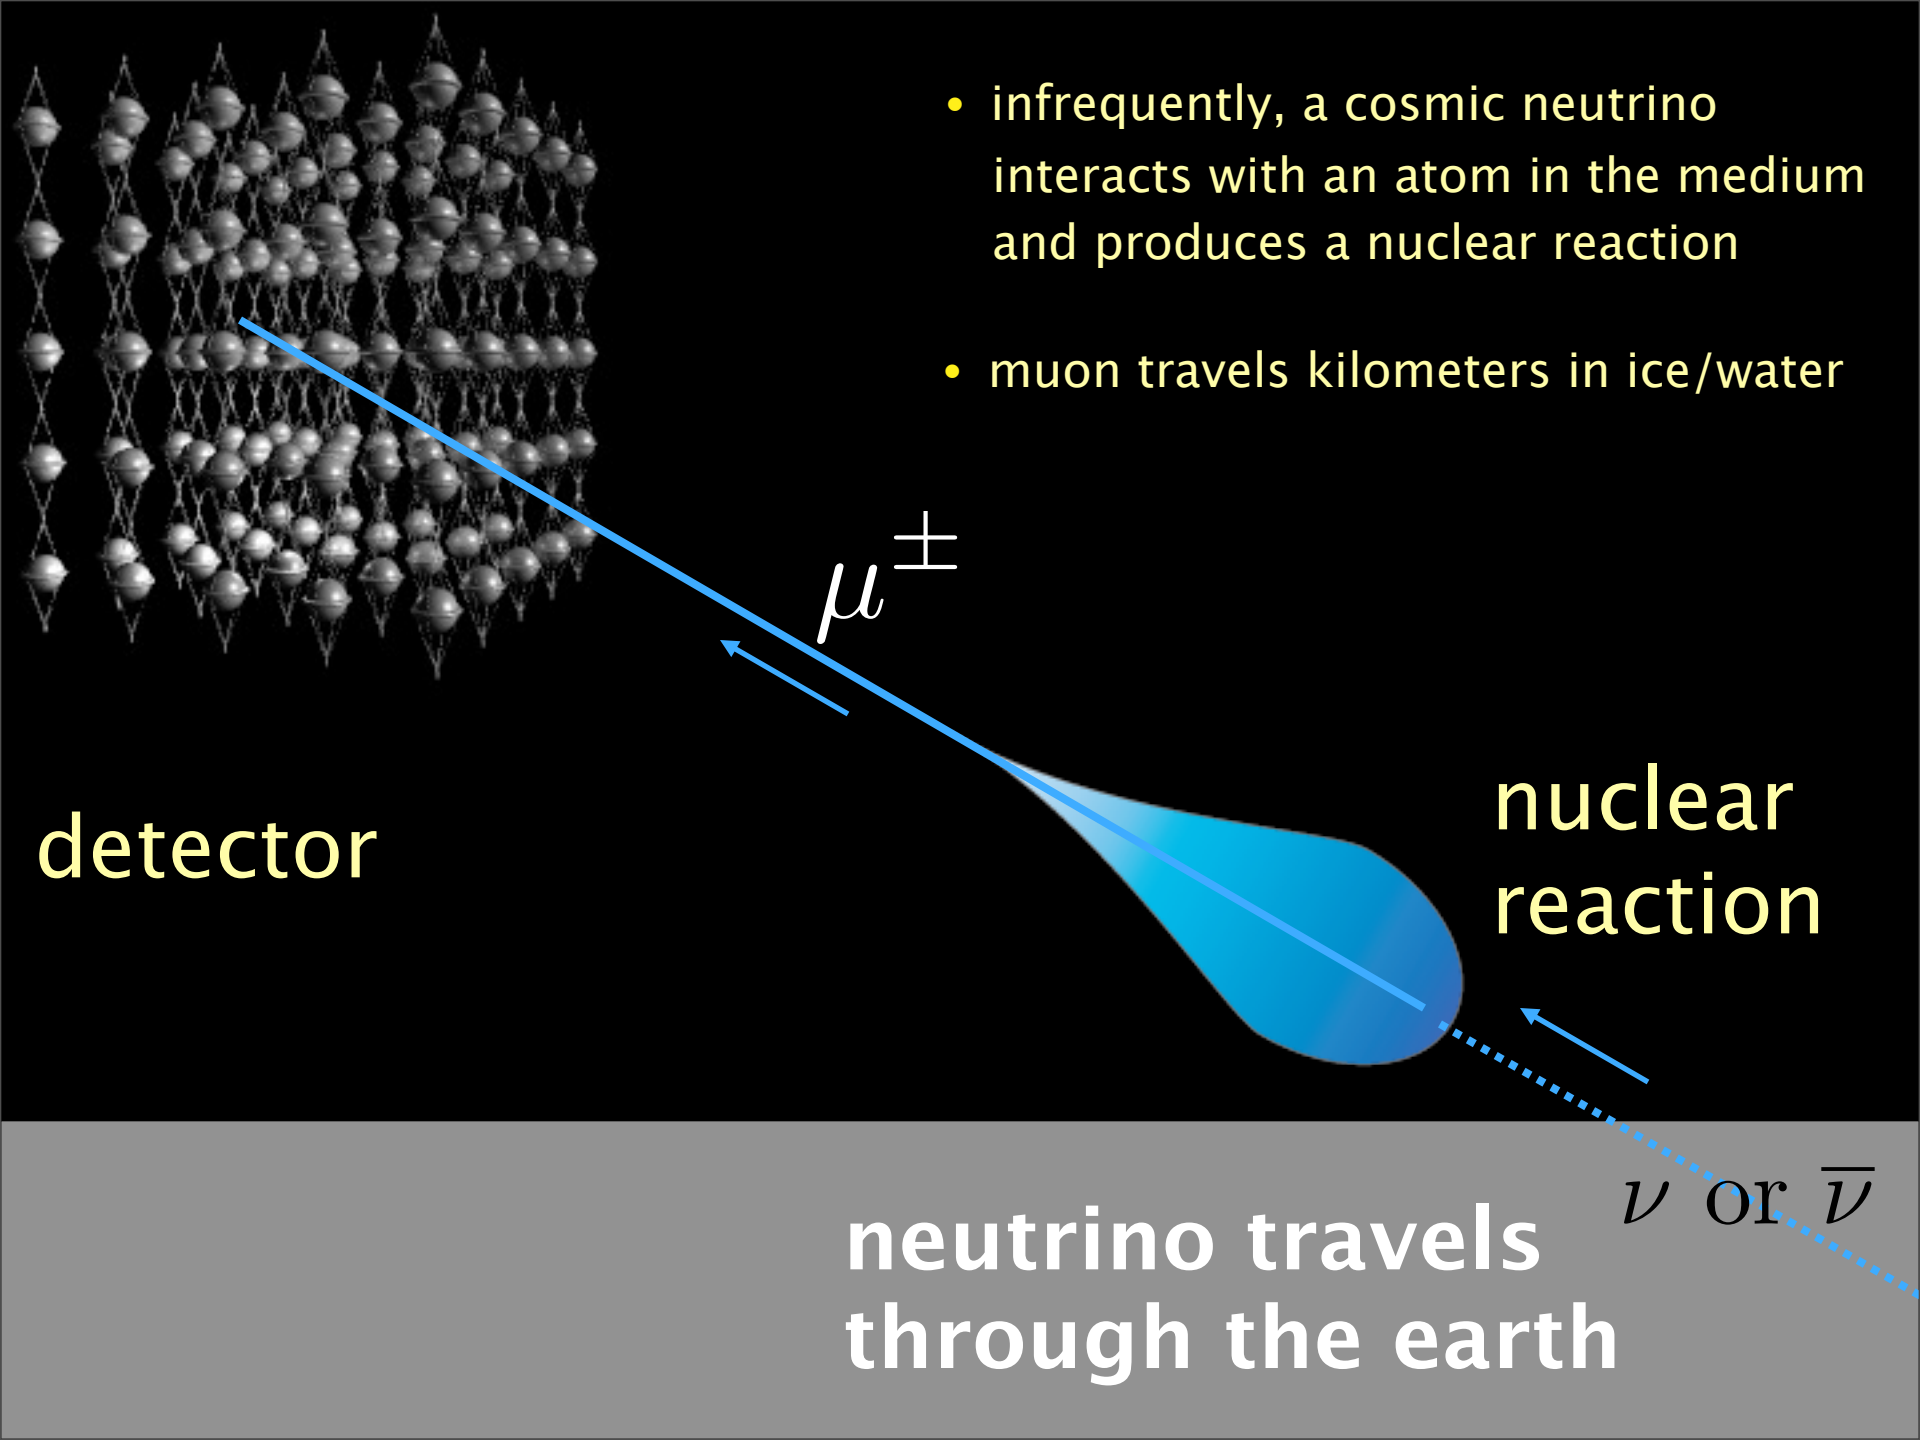
detector

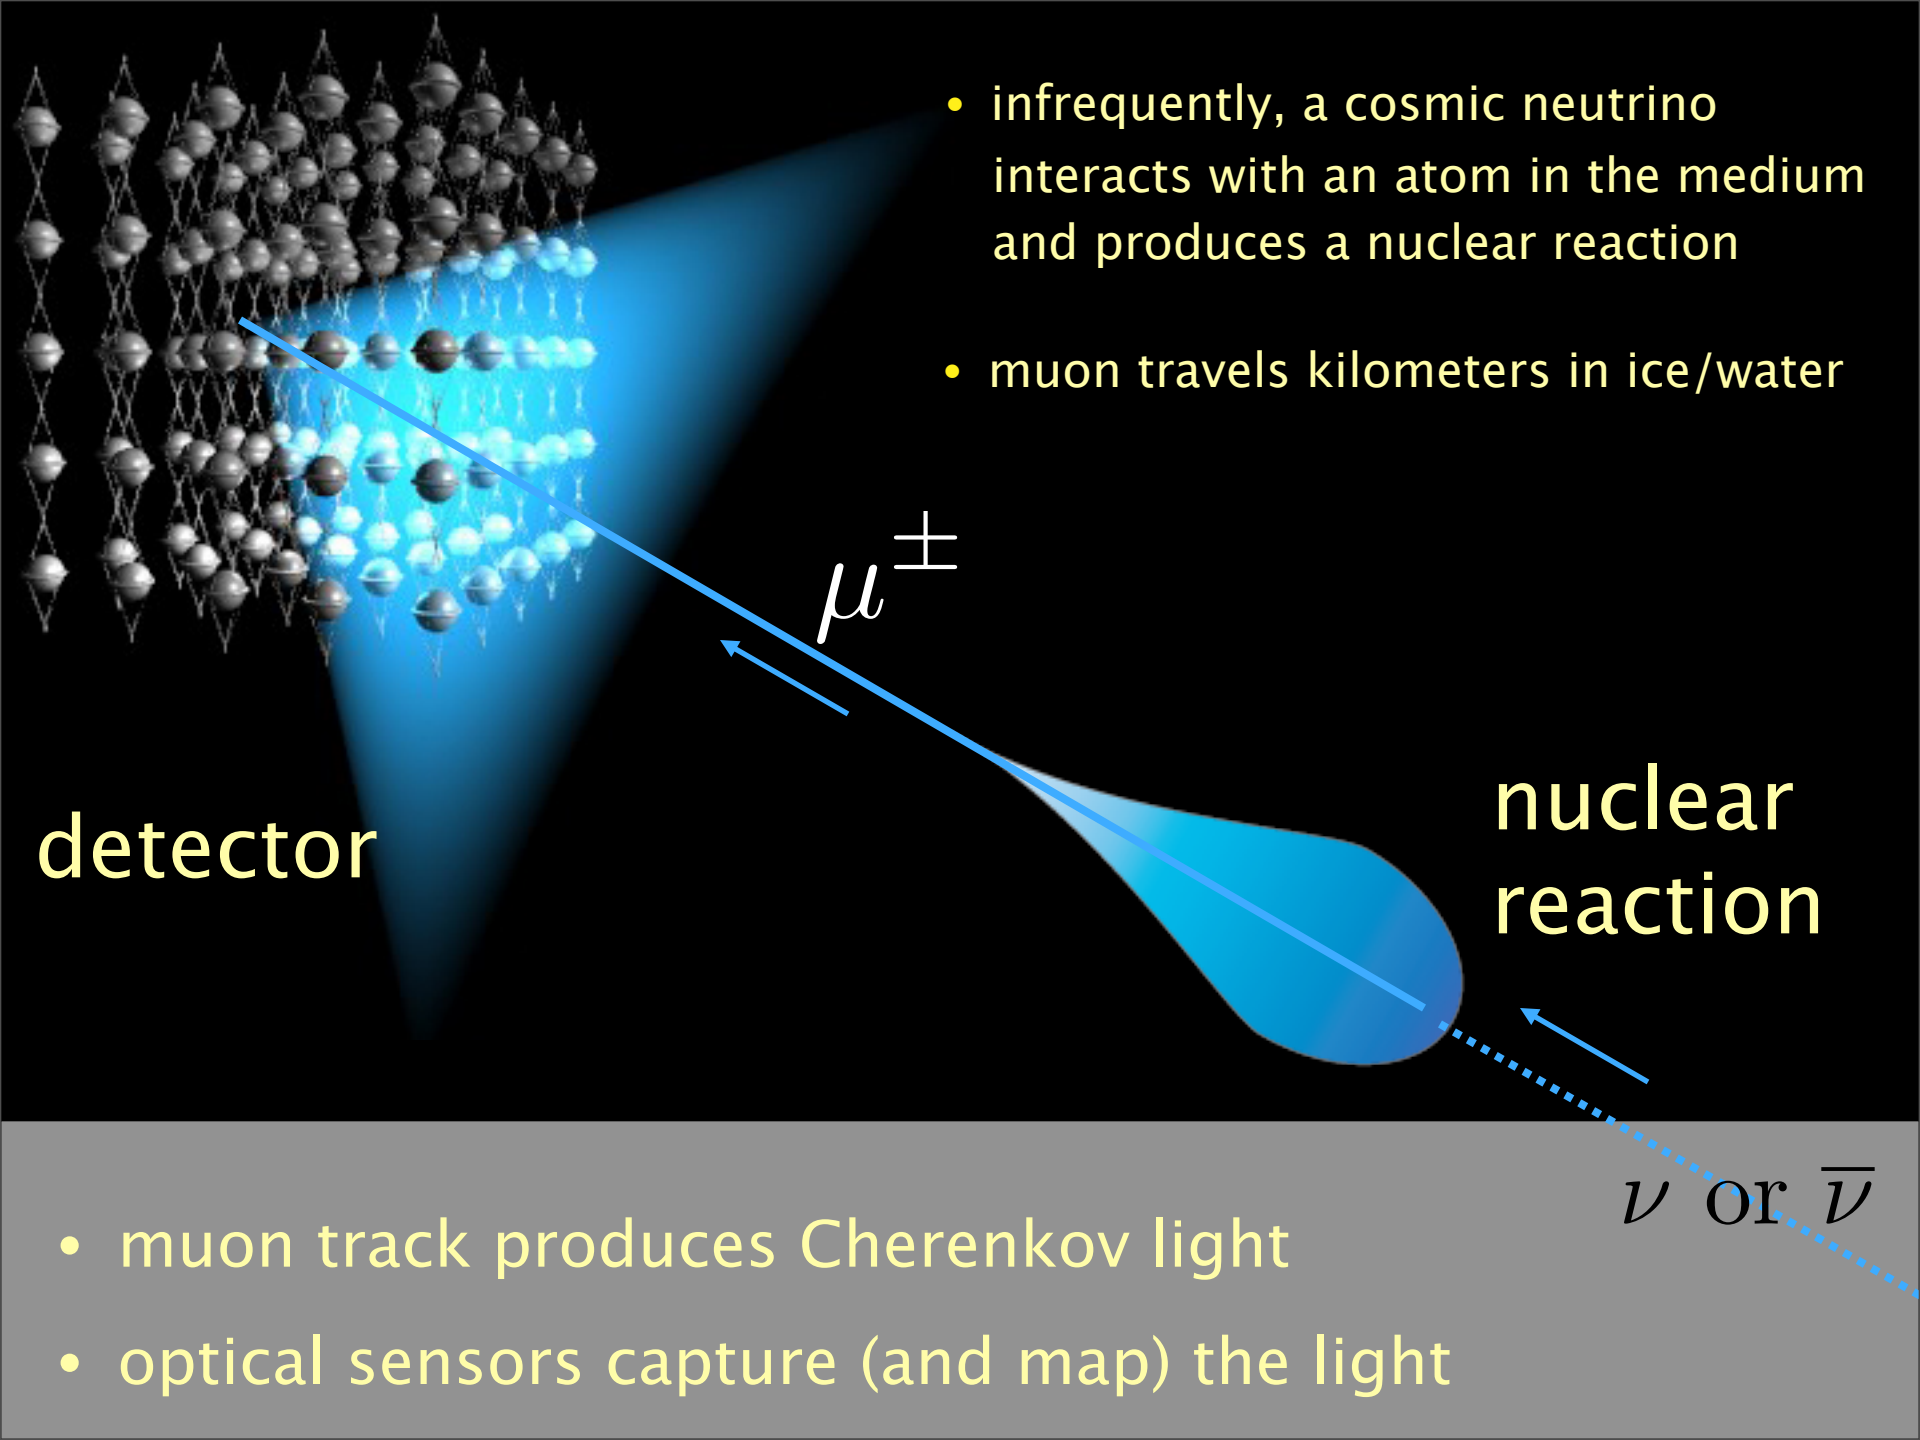
μ^{\pm}

nuclear
reaction

neutrino travels
through the earth

ν or $\bar{\nu}$





- infrequently, a cosmic neutrino interacts with an atom in the medium and produces a nuclear reaction
- muon travels kilometers in ice/water

detector

μ^\pm

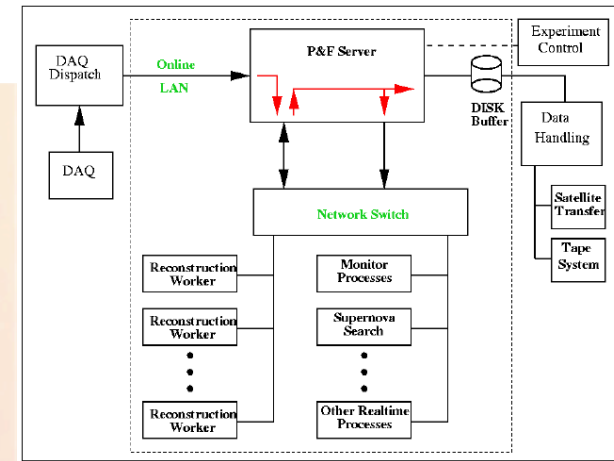
nuclear reaction

ν or $\bar{\nu}$

- muon track produces Cherenkov light
- optical sensors capture (and map) the light

icecube data rates (IC40)

- * $\sim 50 \nu$'s per day
- * $> 10^7 \mu$'s per day



- Based on simulations, the expected trigger rate for a 40 string detector is:

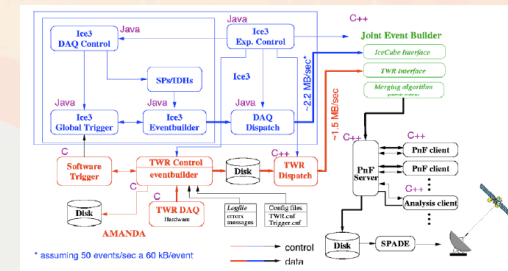
- ~ 950 Hz for an Simple Majority Trigger with ≥ 8 HLC hits.
- This translates into 6.5 MB/sec for the physics data or 548.4 GB/day .
- This trigger rate is expected to double for 80 strings.


- Spade/Satellite transmission

- The expected satellite bandwidth is 30 GB/day.
- A tape system is in place at the South Pole to store the full data.

- JEB/PnF

- JEB reads collected event streams from the I3DAQ and the TWR (AMANDA) DAQ.
- Merges events based on I3DAQ global trigger record in Joint Event Builder (JEB)
- PnF uses fast reconstruction algorithms to filter out events based on incident angle, nchannel, etc.
- Remaining data is written to tapes (which ideally won't be re-read).






simulation

Condor Week
May 1, 2008
UW-Madison



simulation	livetime	runtime
single shower	~2.62 sec	5 h/core
double shower	~62. sec	5 h/core
signal ν_{μ} (E^{-1})*		18 sec/event
signal ν_{μ} (E^{-2})		4 sec/event
signal ν_e (E^{-1})		20 sec/event
signal ν_e (E^{-2})		5 sec/event

*) estimated : runtime depends on configuration

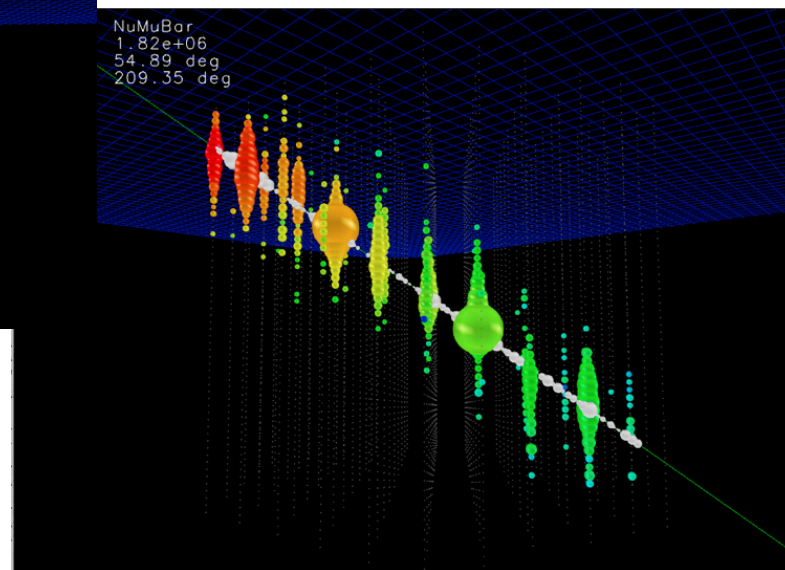
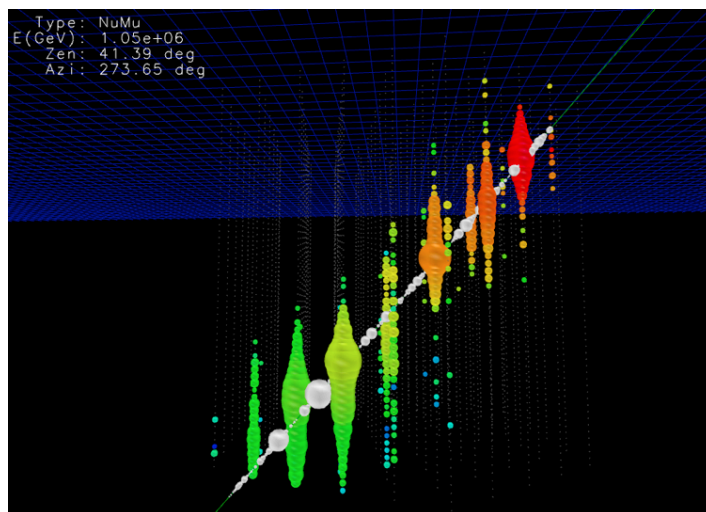



IceCube : simulated μ track events

Condor Week

May 1, 2008



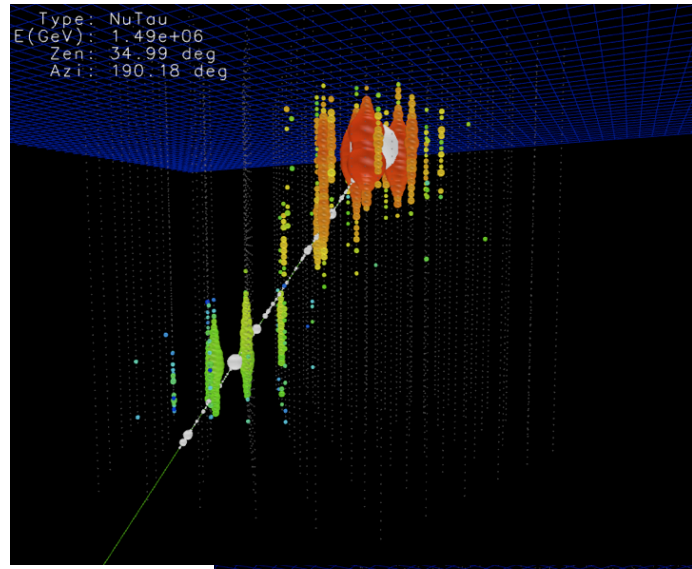
UW-Madison



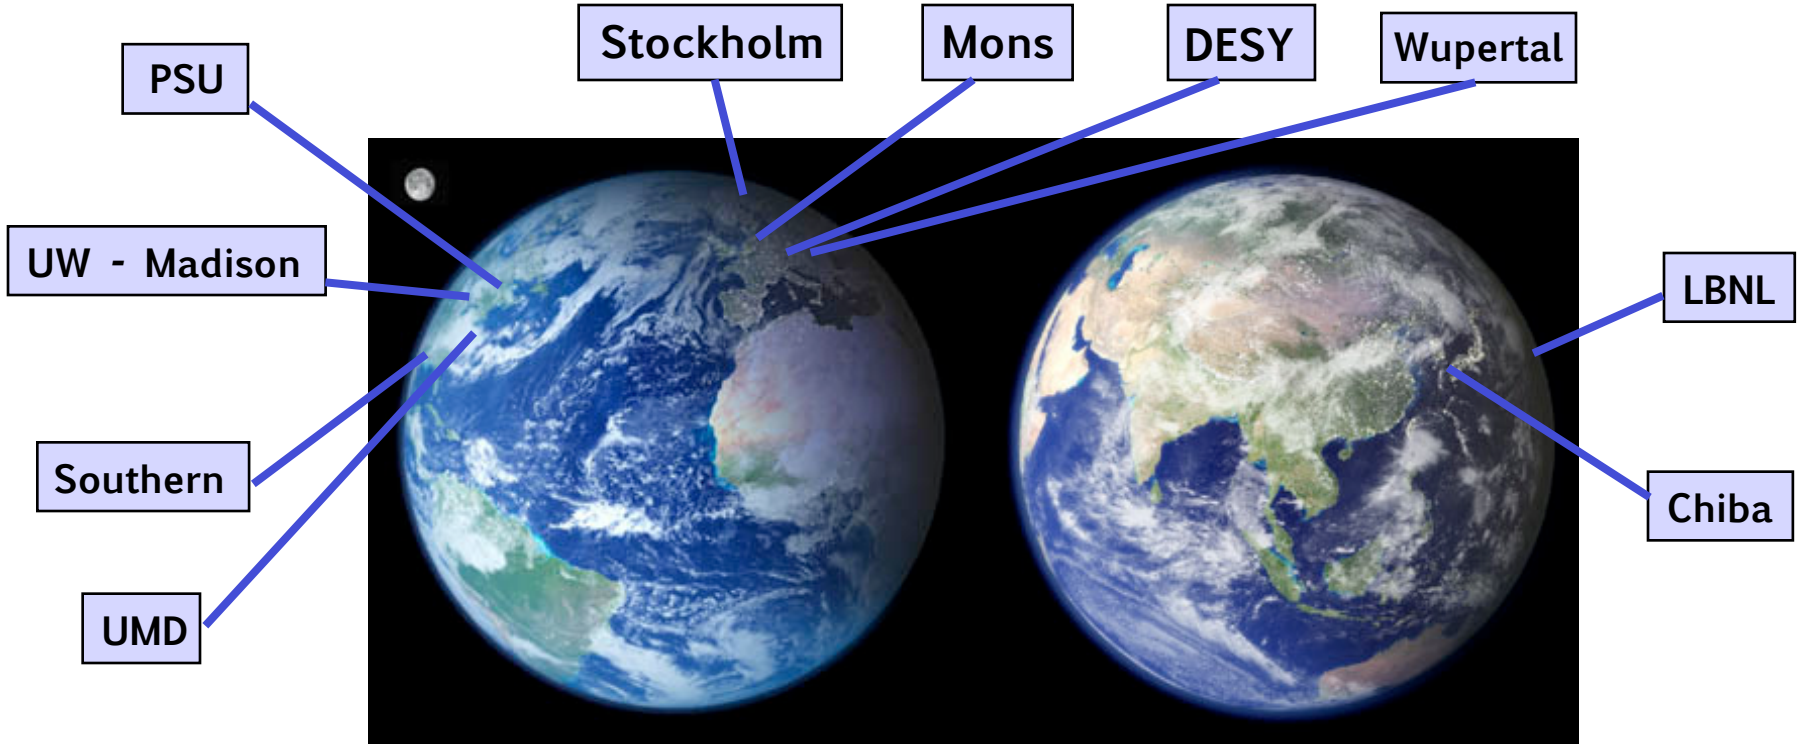




IceCube : simulated events

Condor Week
May 1, 2008
UW-Madison



icecube computing resources






IceCube

Condor Week
 May 1, 2008
 UW-Madison




- different architectures and OS and batch systems
- different policies
- each site provides a local contact person to...
 - work with local sys admin
 - maintain production
 - monitor runtime & completion
 - troubleshoot system &
 - check data integrity




Name	Location	Queueing System	Cores	Normalized
NPX	UW-Madison (Wisc.)	PBS	256	256
GLOW	UW-Madison (Wisc.)	Condor/OSG	120	120
GLOW	UW-Madison (Wisc.)	Condor/OSG	72	50
NPX-UWA (new!)	UW-Madison (Wisc.)	Condor	60	60
DESY	Desy (Germany)	SGE	150	190
FearTheTurtle	UMD (Maryland)	PBS/SGE	132	132
Swegrid	KTH Sweden	NorduGrid	111	70
MONS	Mons, BE	Condor	18	18
BRU	Brussels, BE	Condor/OpenPBS	58	58
KATRINA	Souther (Baton Rouge, LA)	OpenPBS	28	28
ALICEnext	Wuppertal, DE	PBS	512	256
PSU	PSU (PE)	PBS	10	10
CAU	CAU (GA)	PBS	40	40
AachenClust	Aachen, DE	Condor	12	12
CHIBA	Chiba U. (JP)	Condor	11	11

Total 1590 1311




iceprod

Condor Week
May 1, 2008
UW-Madison



- initially intended for documenting production parameters and to catalog simulation data.
- a job management system
 - originally intended for simulation production
 - extended to work for mass data filtering



iceprod

Condor Week
May 1, 2008
UW-Madison

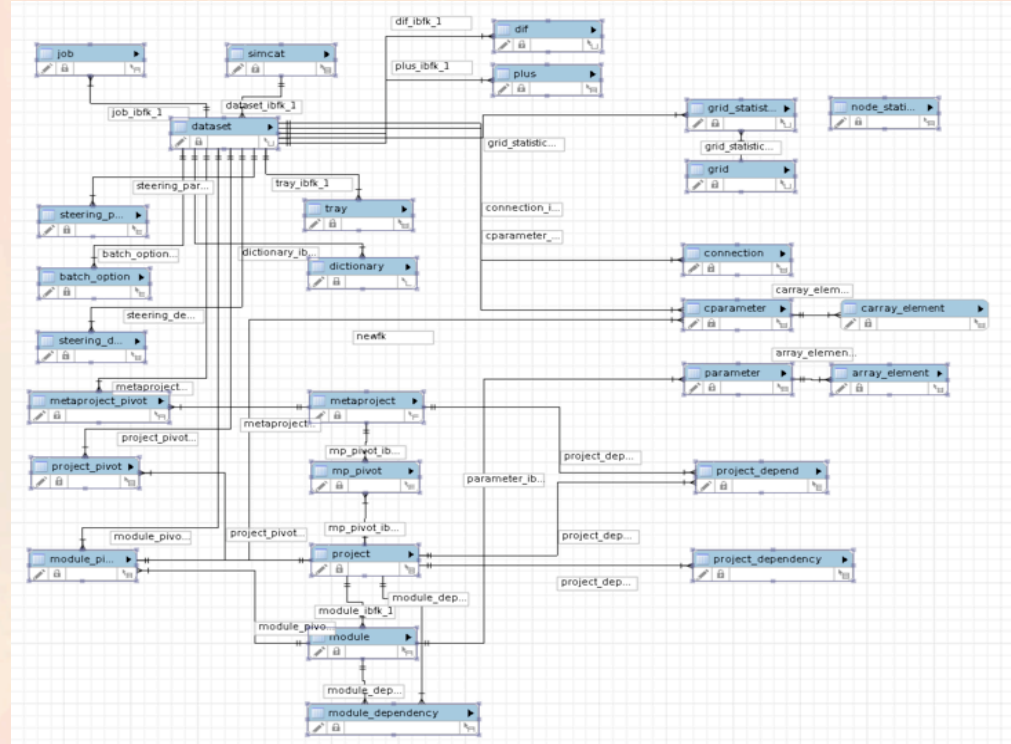


- a system for cataloging steering parameters for IceCube simulation datasets
- also a distributed job management and monitoring system.
- written in python
- daemons manage cluster job submission
- Jobs communicate to daemons via SOAP
- SOAP interface for submission
- central database tracks all production sets through heterogeneous collection of grids and clusters available to IceCube

iceprod

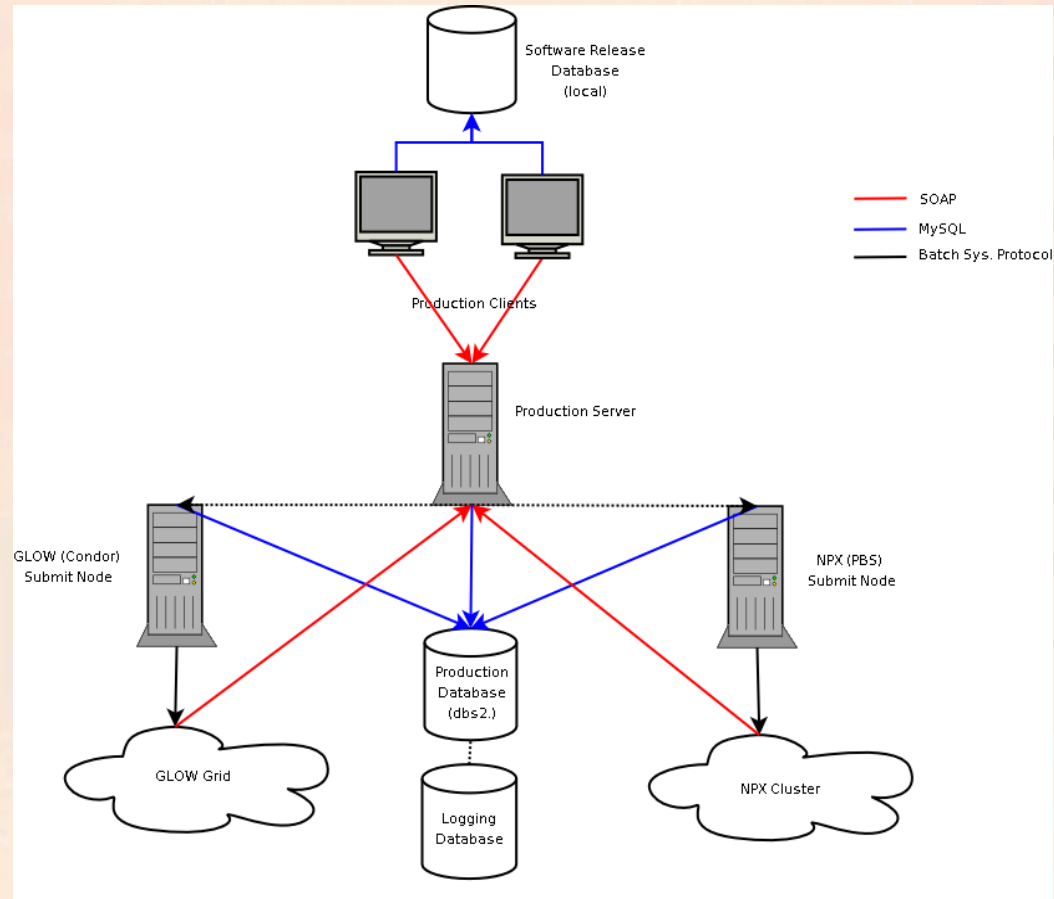
Production Database

- store production history including all configured module parameters
- provide information on configurable parameters for client



iceprod

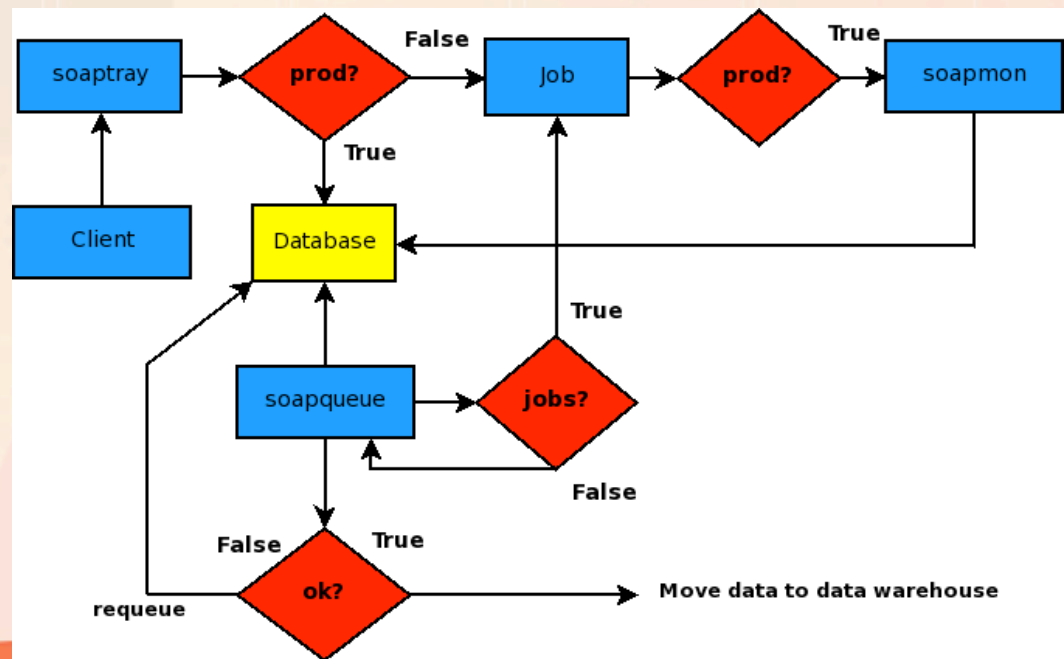
- **Production Database**
 - global database
 - store production history including all configured module parameters
 - provide information on configurable parameters for client
 - global database
- **Server daemons**
 - accept dataset requests from client
 - provides job management including error handling
 - separate daemons handle dataset submission, queue/job management & monitoring



iceprod

Server daemons:

- `soaptray` listens for incoming requests from client
- `soapqueue` checks queue for jobs to process
- `soapmon` receives monitoring updates from jobs
- `soapdh` handles data movement when jobs complete
- `soaphisto` merges and displays verification histograms



iceprod

- Production Database
- Production Server
 - accepts dataset requests from client
 - provides job management including error handling
 - separate daemons handle dataset submission, queue/job management & monitoring
- Queuing Plugins
 - adaptable to different sites and batch systems
 - inherit from base class I3Queue
 - overloads submission commands for specific batch system
 - people can write customized plugins that inherit from each batch system class



iceprod

- Production Database
- Server daemons
- Queuing Plugins
 - adaptable to different sites and batch systems
- Logging/monitoring Database
 - production status & troubleshooting
 - remote job management
 - unified monitoring for multiple clusters

Grid Monitor

ID	Name	Institution	System Type	Jobs Running	Version	soaptray	soapqueue	soapmon
20	AachenCLUST	RWTH-Aachen	condor_nfs	0	V00-05-04	RUNNING	RUNNING	RUNNING
24	ALICEnext	BU-Wuppertal	alicenext	0	V00-06-00	RUNNING	RUNNING	RUNNING
19	chiba	Chiba-U	condor_nfs	0	V00-05-04	STOPPED	STOPPED	STOPPED
9	desy	DESY	sge	20	V00-06-00	RUNNING	RUNNING	RUNNING
8	FearTheTurtle	UMD	pbs	184	V00-06-00	RUNNING	RUNNING	RUNNING
1	GLOW	UW-Madison	condor	499	V00-06-01	RUNNING	RUNNING	RUNNING
22	IIHE	IIHE-Brussels	condor_nfs	0	V00-06-00	RUNNING	RUNNING	RUNNING
4	Katrina	Southern University	pbs	166	V00-05-04	RUNNING	RUNNING	RUNNING
6	Mons	UMH	condor_nfs	0	V00-06-00	RUNNING	RUNNING	RUNNING
2	NPX	UW-Madison	pbs	0	V00-05-04	STOPPED	STOPPED	STOPPED
14	npv2	UW-Madison	pbs	96	V00-06-01	RUNNING	RUNNING	RUNNING
3	PDSF	LBNL	sge	0	V00-06-00	RUNNING	RUNNING	RUNNING
7	Super-K	Chiba University	Condor	0	V00-05-04	STOPPED	STOPPED	STOPPED
5	SWEGRID	StockholmUniversity	swegrid	1	V00-05-04	STOPPED	STOPPED	STOPPED

simulation 02-00-14 - Jobs

simulation 02-00-14 ▾ Any Generator ▾ Any Grid ▾ Any Dataset ▾ Any Job Status ▾ Apply Filters

Suspend Resume

Page: 1 2 3 > Last >

Results 1 - 100 of 585217.

	ID	Status	Metaproject	Generator	Grid	Host	Fails	Evicts	Events
<input type="checkbox"/>	796.0	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade00	0	0	1332
<input type="checkbox"/>	796.1	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade03	0	0	1350
<input type="checkbox"/>	796.2	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade04	0	0	1340
<input type="checkbox"/>	796.3	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade17	0	0	1298
<input type="checkbox"/>	796.4	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade25	0	0	1340
<input type="checkbox"/>	796.5	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade12	0	0	1286
<input type="checkbox"/>	796.6	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade13	0	0	1371
<input type="checkbox"/>	796.7	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade18	0	0	1290
<input type="checkbox"/>	796.8	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade10	0	0	1364
<input type="checkbox"/>	796.9	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade09	0	0	1355
<input type="checkbox"/>	796.10	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade14	0	0	1309
<input type="checkbox"/>	796.11	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade07	0	0	1288
<input type="checkbox"/>	796.12	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade06	0	0	1349
<input type="checkbox"/>	796.13	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade15	0	0	1324
<input type="checkbox"/>	796.14	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade11	0	0	1325
<input type="checkbox"/>	796.15	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade24	0	0	1336
<input type="checkbox"/>	796.16	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade16	0	0	1332
<input type="checkbox"/>	796.17	OK	simulation 02-00-14	CORSIKA-in-ice	desy	galaxy1	0	0	1358
<input type="checkbox"/>	796.18	OK	simulation 02-00-14	CORSIKA-in-ice	desy	galaxy4	0	0	1384
<input type="checkbox"/>	796.19	OK	simulation 02-00-14	CORSIKA-in-ice	desy	galaxy23	0	0	1297
<input type="checkbox"/>	796.20	OK	simulation 02-00-14	CORSIKA-in-ice	desy	galaxy11	0	0	1320
<input type="checkbox"/>	796.21	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade10	0	0	1318
<input type="checkbox"/>	796.22	OK	simulation 02-00-14	CORSIKA-in-ice	desy	galaxy21	0	0	1329
<input type="checkbox"/>	796.23	OK	simulation 02-00-14	CORSIKA-in-ice	desy	galaxy16	0	0	1377
<input type="checkbox"/>	796.24	OK	simulation 02-00-14	CORSIKA-in-ice	desy	galaxy29	0	0	1332
<input type="checkbox"/>	796.25	OK	simulation 02-00-14	CORSIKA-in-ice	desy	blade10	0	0	1333

iceprod

- Production Database
- Server daemons
- Queuing Plugins
- Logging/monitoring Database
 - production status & troubleshooting
 - remote job management
 - unified monitoring for multiple clusters
- Web Interface
 - cluster/dataset/job monitoring
 - search engine for production db
 - dataset statistics

Simulation Production
IceCube Internal

Juan Carlos Diaz-Velez
edit profile
check mail
log out

Dashboard Directory Internal Reports Masterpiece PQ Registration Simulation Time

Home Configuration Files Datasets Jobs Job Queues Graphs Grids Nodes Tickets

simulation 02-00-14 - neutrino-generator - Dataset

simulation 02-00-14 neutrino-generator Any Grid Any Dataset Status Apply Filters

Any Dataset Category

Page: 1 2 > Results 1 - 20 of 27.

Dataset 1051 simulation 02-00-14 neutrino-generator GLOW PHYSICS READYTOPUBLISH

Description

IC22 neutrino-generator NuE with E^{-1} neutrino spectrum, using AHA07v1 photon tables, $90\text{deg} < \theta < 180\text{deg}$, $10^{13} < E < 10^{19}$ GeV.

Jobs

Statistics

Actions

Finish Retire Nuke Clean Hide

Dataset 1045 simulation 02-00-14 neutrino-generator desy PHYSICS PROCESSING

Description

IC22+TWR neutrino-generator NuMu with E^{-1} neutrino spectrum, using AHA07v1 photon tables with AMASpan for TWR, $70\text{deg} < \theta < 180\text{deg}$, $10\text{GeV} < E < 10^{19}$ GeV. This dataset uses I3BasicHisto to generate histograms.

Jobs

Statistics

Actions

Dataset 1044 simulation 02-00-14 neutrino-generator desy PHYSICS PROCESSING

Description

IC22+TWR neutrino-generator NuMu with E^{-1} neutrino spectrum, using AHA07v1 photon tables with AMASpan for TWR, $70\text{deg} < \theta < 180\text{deg}$, $10\text{GeV} < E < 10^{19}$ GeV. This dataset uses I3BasicHisto to

simprod

- Server daemons
- Queuing Plugins
- Logging/monitoring Database
 - production status & troubleshooting
 - remote job management
 - unified monitoring for multiple clusters
- Web Interface
 - cluster/dataset/job monitoring
 - search engine for production db
 - dataset statistics
- Usage reports
 - daily email reports are sent to subscriber list

[Simprod] SimProd Usage Summary for Thu Apr 26 10:00:0... ice3simusr@icecube.wisc.edu 9:59 PM

Asunto: [Simprod] SimProd Usage Summary for Thu Apr 26 10:00: De: ice3simusr@icecube.wisc.edu 9:59 PM

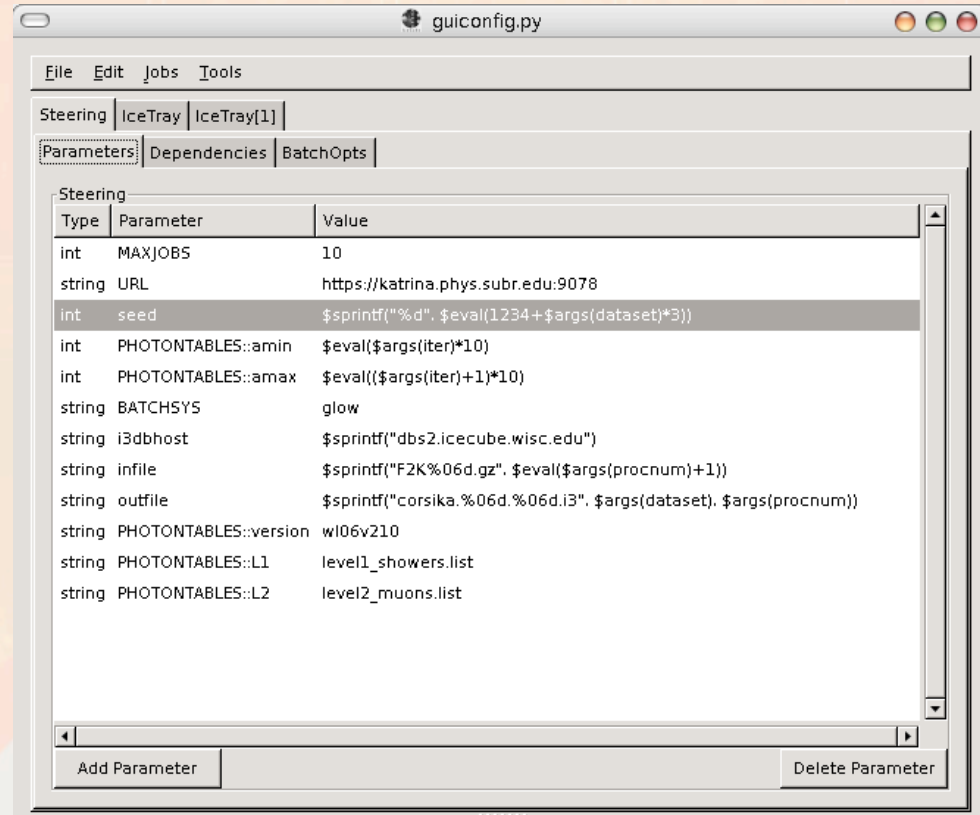
```
----- monthly summary -----
sys_t = 637934.749814
ok = 7045.0
usr_t = 72976855.6935
real_t = 119045455.388
suspended = 430.0
error = 0.0
events = 1210641.0
```

grid	sys_t	ok	usr_t	real_t	grid_id	suspend	error	events
GLOW	5.6e+04	5.2e+02	2.6e+06	2.9e+06	1	1.7e+02	0	8.4e+04
PDSF	1.2e+04	93	1.9e+06	2.3e+06	3	0	0	3.1e+05
Katrina	2.7e+05	2e+03	5.7e+07	1e+08	4	52	0	2e+05
FearThe	2.7e+05	3.9e+03	1.1e+07	1.3e+07	8	1.7e+02	0	4.8e+05
desy	0	0	0	0	9	20	0	0
npv2	3.6e+04	5.2e+02	7.6e+05	1e+06	14	20	0	1.4e+05

Simprod mailing list
Simprod@icecube.wisc.edu
<http://www.icecube.wisc.edu/mailman/listinfo/simprod>

simprod

- Queuing Plugins
- Logging/monitoring Database
- Web Interface
 - cluster/dataset/job monitoring
 - search engine for production db
 - dataset statistics
- Usage reports
 - daily email reports are sent to subscriber list
- GUI Production Client (cmd line interface also available)
 - editor of XML IceTray steering files
 - interface to SOAP daemons for job submission



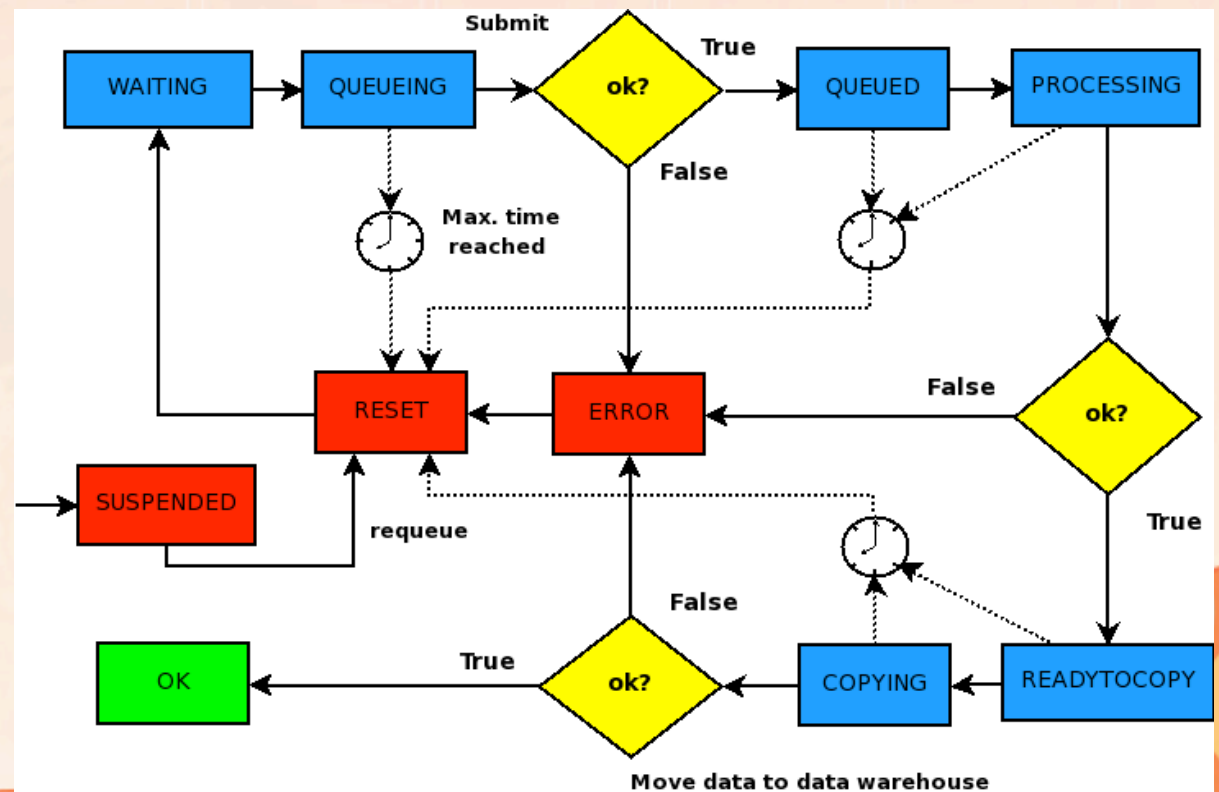
simulation ticketing system

RT Queue

Number	Subject	Status	RT Link
6211	IC80 MC production	open	view in RT
6176	increase statistics of E-2 production	open	view in RT
6148	Re: Simulation bug?	resolved	view in RT
6147	Re: Simulation bug?	resolved	view in RT
6115	Low energy muons for trigger studies	new	view in RT
6108	MC request from point source working group	open	view in RT
6097	Re: [ice3cascade] Lower energy nu e's centered on AMANDA	resolved	view in RT
6096	Re: [ice3cascade] Lower energy nu e's centered on AMANDA	resolved	view in RT
6094	Re: [ice3cascade] Lower energy nu e's centered on AMANDA	open	view in RT
6090	CORSIKA request (320 again)	new	view in RT
6089	IC80 coincident muon data set	open	view in RT
6079	Re: AutoReply: Two requests	resolved	view in RT
6078	Re: AutoReply: Two requests	open	view in RT
6067	Nugen IceCube-80+AMANDA Nue simulation request	resolved	view in RT
6056	atmospheric neutrinos for the unfolding analys	resolved	view in RT
6034	IC80 nugen dataset	resolved	view in RT
5907	benchmark dataset request : simplegenerator	resolved	view in RT
5832	"Standard" nu_e simulation production	resolved	view in RT
5823	[Fwd: Sim-prod simple generator cascades]	open	view in RT
5627	request benchmark dataset	resolved	view in RT

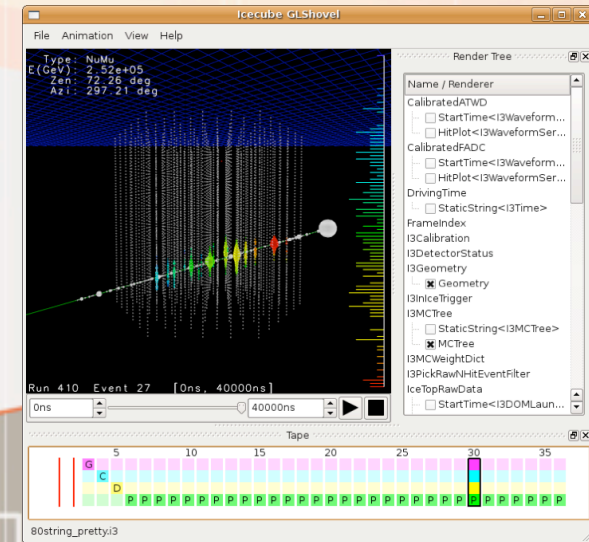
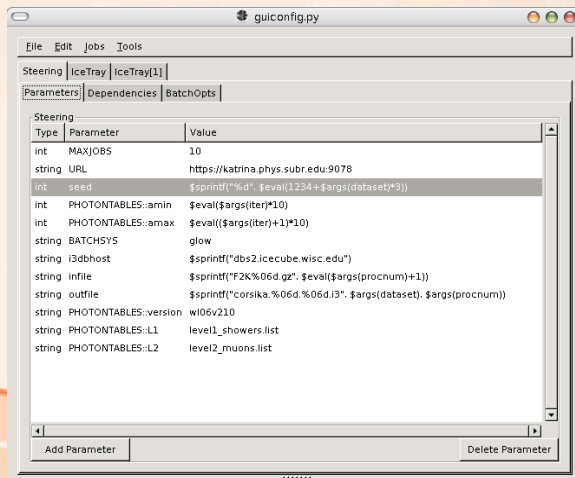
job management

- job goes through series of states with configurable timeouts in case communication is dropped.



icetray

- Software frame work for IceCube data processing
- C++/Boost/Python
- used for both DAQ data and Monte Carlo
- Steering files are written in Python
- XML interface for production system



iceprod

- **statistics**

- I3XMLSummaryService

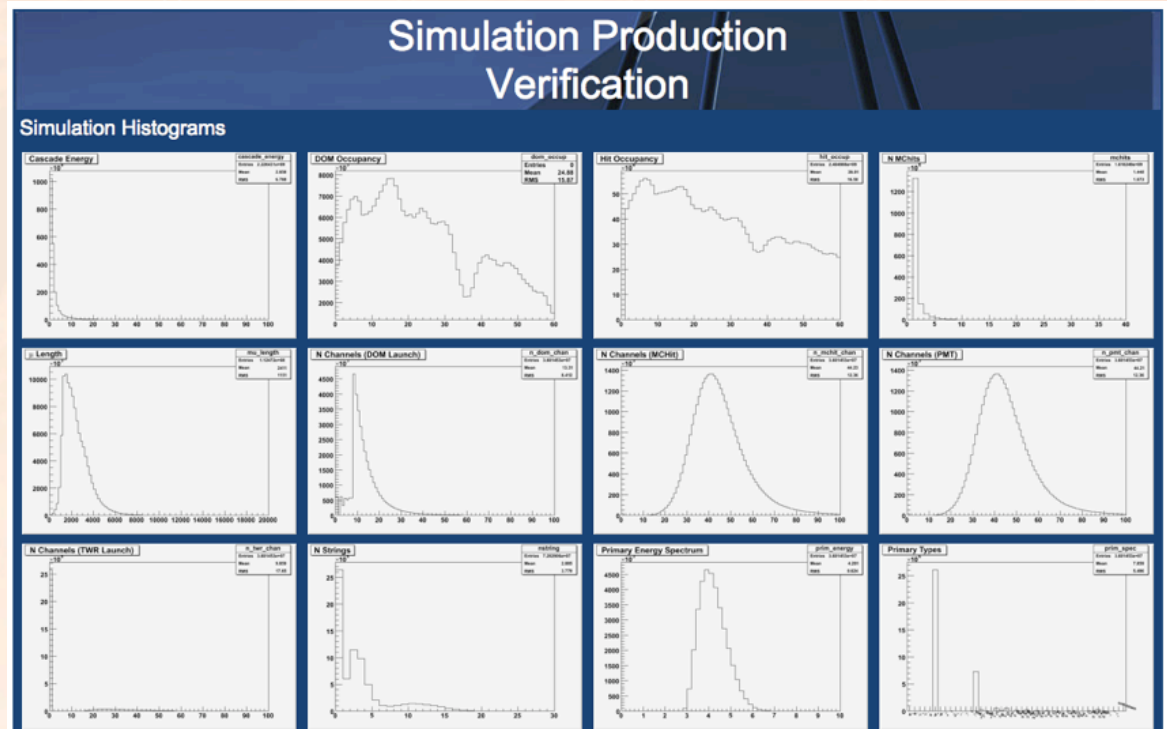
- Installs an I3SummaryService
 - modules can write floats values to service
 - values get written to XML summary file
 - interfaces to iceprod so database can collect statistics


- I3EventCounter

- counts events at various stages of simulation chain
 - writes statistics to I3SummaryService

iceprod

- simulation verification
 - automatically generate plots from simulation
 - basic sanity checks of data





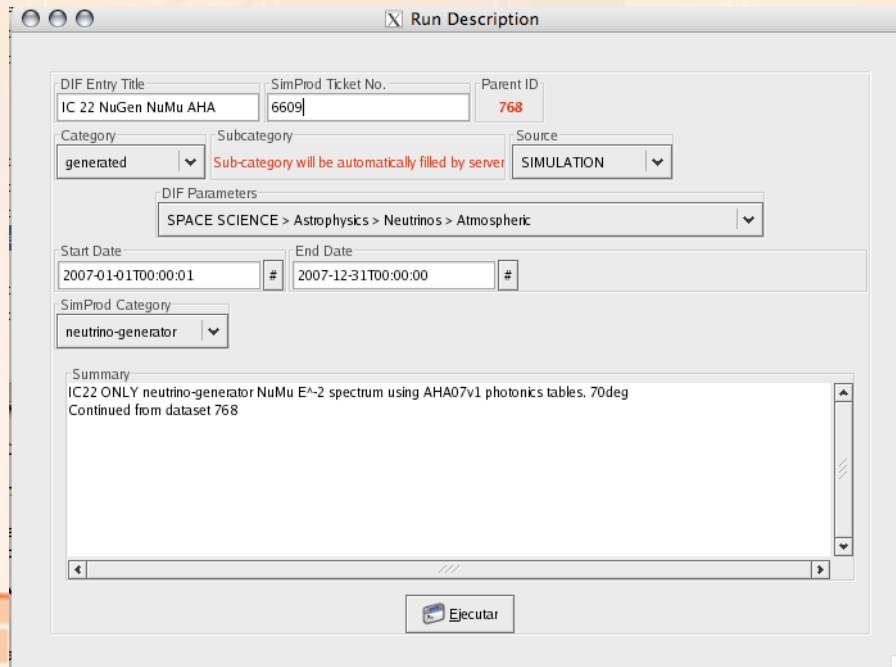
data storage and transfer

Condor Week
May 1, 2008
UW-Madison

- server initiated
 - soapdh daemon handles movement of data after job completes. Various protocols supported
- client initiated
 - job handle direct upload to target url
 - implemented in 0.7.3 but requires more testing
 - not possible for some clusters
 - gridFTP client (globus toolkit) needs to be added to i3tools in order to use gsiftp protocol.
 - requires care to insure security

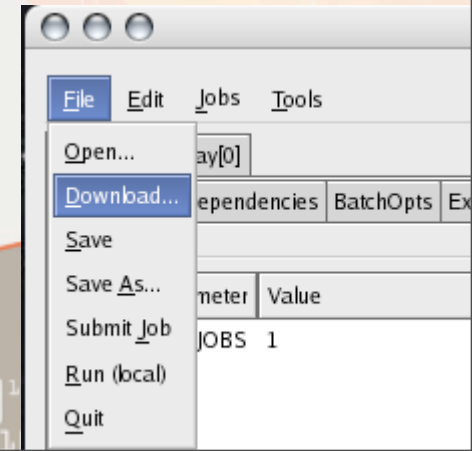
iceprod


- dataset Inheritance
 - dataset parent id is automatically set when you download a configuration from the database
 - you can diff configurations from related datasets
 - you can track changes in configurations



Ancestry

- Dataset 648 (diff)
 - Dataset 651 (diff)
 - Dataset 741 (diff)
 - Dataset 747 (diff)
 - Dataset 753 (diff)
 - Dataset 759 (diff)
- Dataset 768

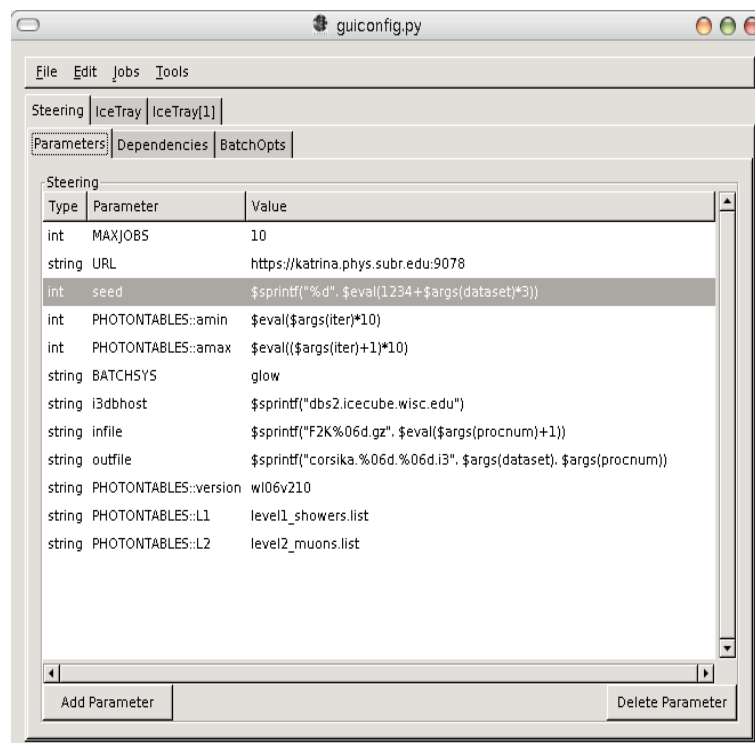





Submitting Jobs

Condor Week
May 1, 2008
UW-Madison

- Use the GUI to build a simulation from scratch, edit an XML file or download and edit a configuration file from a previous run.





submitting jobs

Condor Week
May 1, 2008
UW-Madison

- Select and configure services and modules


The screenshot displays the xsimprod application interface. The main window has tabs for Steering, IceTray, IceTray[1], IceTray[2], and IceTray[3]. Below these are input fields for Events (1000000) and Iterations (1). The Services tab is active, showing a list of available services and modules.

DbID	Default Name	Class	Project
73	particledatase	I3ParticleDataServiceFactory	phys-services
76	mediumservicef	I3MediumServiceFactory	phys-services
77	fileomkey2mbid	I3FileOMKey2MBIDFactory	phys-services
78	trandomservice	I3TRandomServiceFactory	phys-services
79	gsrlrandomservi	I3GSLRandomServiceFactory	phys-services
80	springrandomser	I3SPRNGRandomServiceFactory	phys-services
30	javavmfactory	I3JavaVMFactory	c2j-icetray
32	wrappedclassfa	I3WrappedClassFactory	c2j-icetray
71	servicefactory	I3ServiceFactory	icetray
28	rootboxesfacto	RootI3BoxesFactory	root-icetray
29	rootexecutionf	RootI3ExecutionFactory	root-icetray
31	rootmodulesche	RootModuleSchedulerFactory	root-icetray
43	rootconfigfact	RootConfigFactory	root-icetray
34	omdbfactory	I3OMDbFactory	I3Db
39	dbomkey2mbidfa	I3DbOMKey2MBIDFactory	I3Db
37	psi_photonicsf	I3PSI_PhotonicsFactory	hit-constructor
72	f2kfilefactory	I3F2kFileFactory	amanda-core
74	monolithdatadi	I3MonolithDataDirectoryFactory	monolith-reader
75	stmsflowfactor	I3StmFlowFactory	stms-flow

The Parameter Table window shows the following parameters:

Type	Name	Value
int	NStreams	$\$args(nproc)$
int	Seed	20
int	StreamNum	$\$args(procnun)$

Buttons: Add Parameter, Delete Parameter



parameter expressions

Condor Week
May 1, 2008
UW-Madison

Comand line options:

`$args(procnum)`, `$args(nproc)`, etc - is replaced by value passed via options `--procnum=` , `--nproc=` , etc.

Steering parameters as variables:

`$steering(<varname>)` - is replaced by steering parameter `<varname>`


Python expressions:

`$eval(<expr>)` - restricted python expressions

String formatting:

`$sprintf(" %s %d %f", <expr1>, <expr2>, <expr3>)`

Recursive expresion evaluation.




parameter expressions

Condor Week
May 1, 2008
UW-Madison

Example 1: SPRNGRandomService

```
<parameters>  
  <int>  
    <name>Stream</name>  
    <value>$args(procnum)</value>  
  </int>  
  <int>  
    <name>Nstreams</name>  
    <value>$args(nproc)</value>  
  </int>  
</parameters>
```

```
l3sim.py --nproc=10 --procnum=2 config.xml
```



parameter expressions

Condor Week
May 1, 2008
UW-Madison

Example 2: dependency files

```
<dependency>  
    $format("corsika_%06d.f2k", $eval($args(nproc) + 101))  
</dependency>
```

With arguments...

```
l3sim.py --nproc=10 --procnum=2 config.xml
```

Is evaluated as..

```
corsika_000103.f2k
```



Submitting a Cluster of Jobs

- Select 'submit' from the File menu
- Enter a description of what you are simulating.

The screenshot shows two windows from the IceCube GUI. The 'guiconfig.py' window has the 'File' menu open, with 'Run (local)' selected. Below the menu is a table of parameters:

Parameter	Value
OBS	3
PROJECT	simulation-V01-02-01
string BATCHSYS	condor
string LIBDIR	/data/juancarlos/simprod-v1/archives
string PHOTONTABLES	/home/icecube/tables
string OUTDIR	/tmp/results


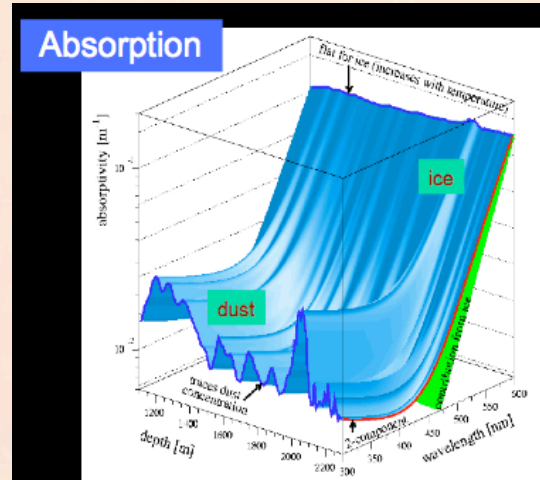
The 'Run Description' dialog box is open, showing the following fields:


- DIF Entry Title: corsika 23 string dataset
- Category: generated
- Subcategory: Sub-category will be automatically filled by server
- Source: SIMULATION
- DIF Parameters: SPACE SCIENCE > Astrophysics > Cosmic Rays > Cosmic Ray Muons
- Start Date: 2007-01-01T00:00:00
- End Date: 2007-12-31T00:00:00
- SimProd Category: CORSIKA
- Summary: CORSIKA production with 23 strings of IceCube, simple multiplicity trigger threshold = 8 with HLC. Lifetime per file is 7.136 sec. LC Time Window was corrected to +/- 1000 nsec. In the BDB it was set to +/- 500 nsec.
- Ejecutar button

grid computing challenges

photonics

- photon interaction probability tables are produced with detailed module of ice properties
- full set of tables is >14 GB (too large to load in memory on most nodes)
- we sort events in zenith bins and process process each bin separately.
- most of our current production clusters have tables pre-installed on nodes
- This limits our ability to add new clusters or large grids for simulation production.







event sorting

Condor Week
May 1, 2008
UW-Madison



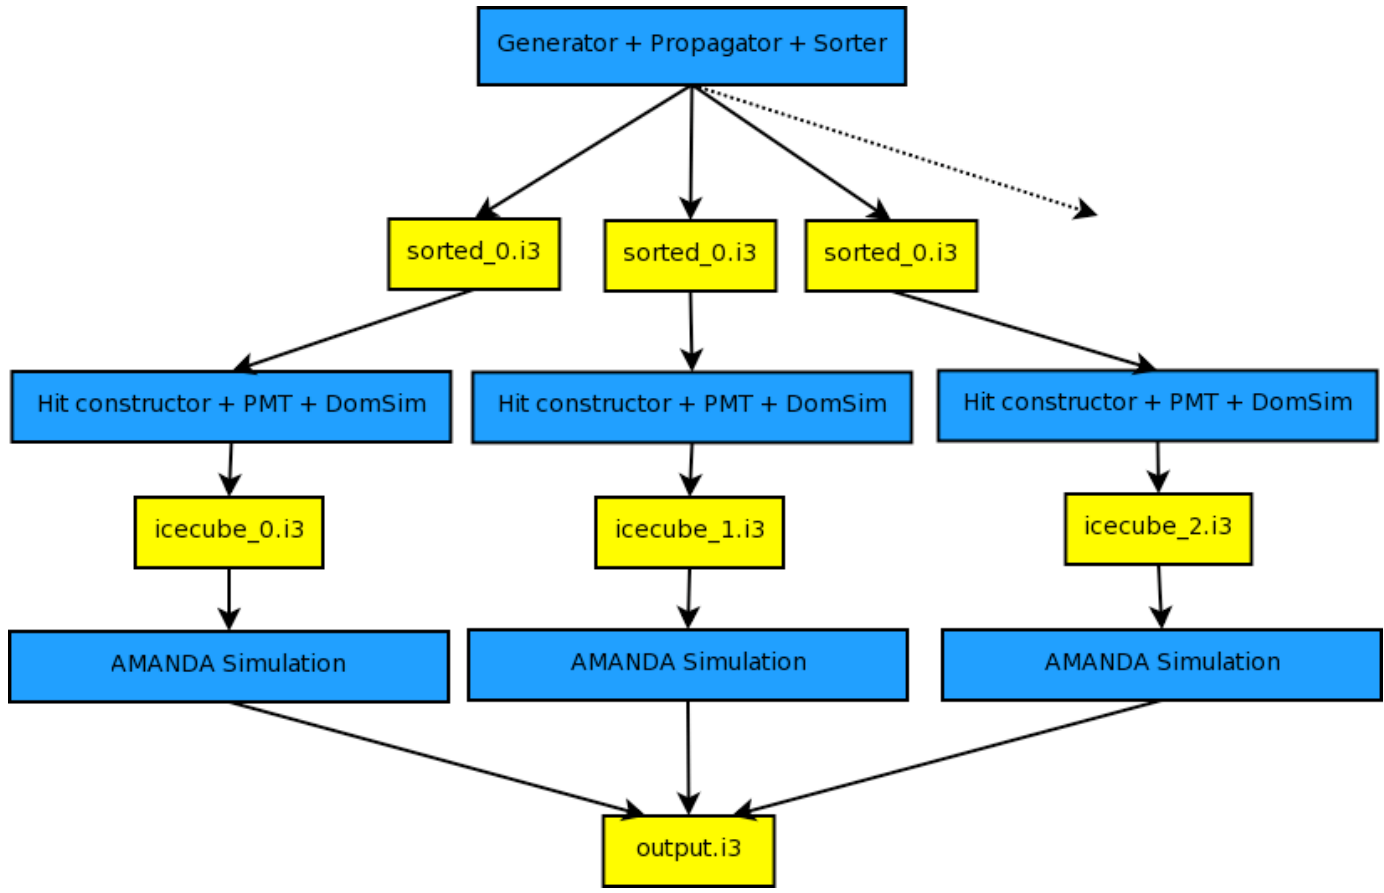
- custom event writer iterates through particle tree to find θ_{\min} .
- Events are then assigned to a bin (file) on the basis of θ_{\min} .
- Load adjacent bin to account for secondaries which spill over.






“DAG” of joint AMANDA IceCube Simulation with Fine Photonics Binning

Condor Week
May 1, 2008
UW-Madison






high throughput computing

Condor Week
May 1, 2008
UW-Madison






FLOPY \neq (60*60*24*7*52)*FLOPS






icecube on GLOW





Condor Week
 May 1, 2008
 UW-Madison






HTC

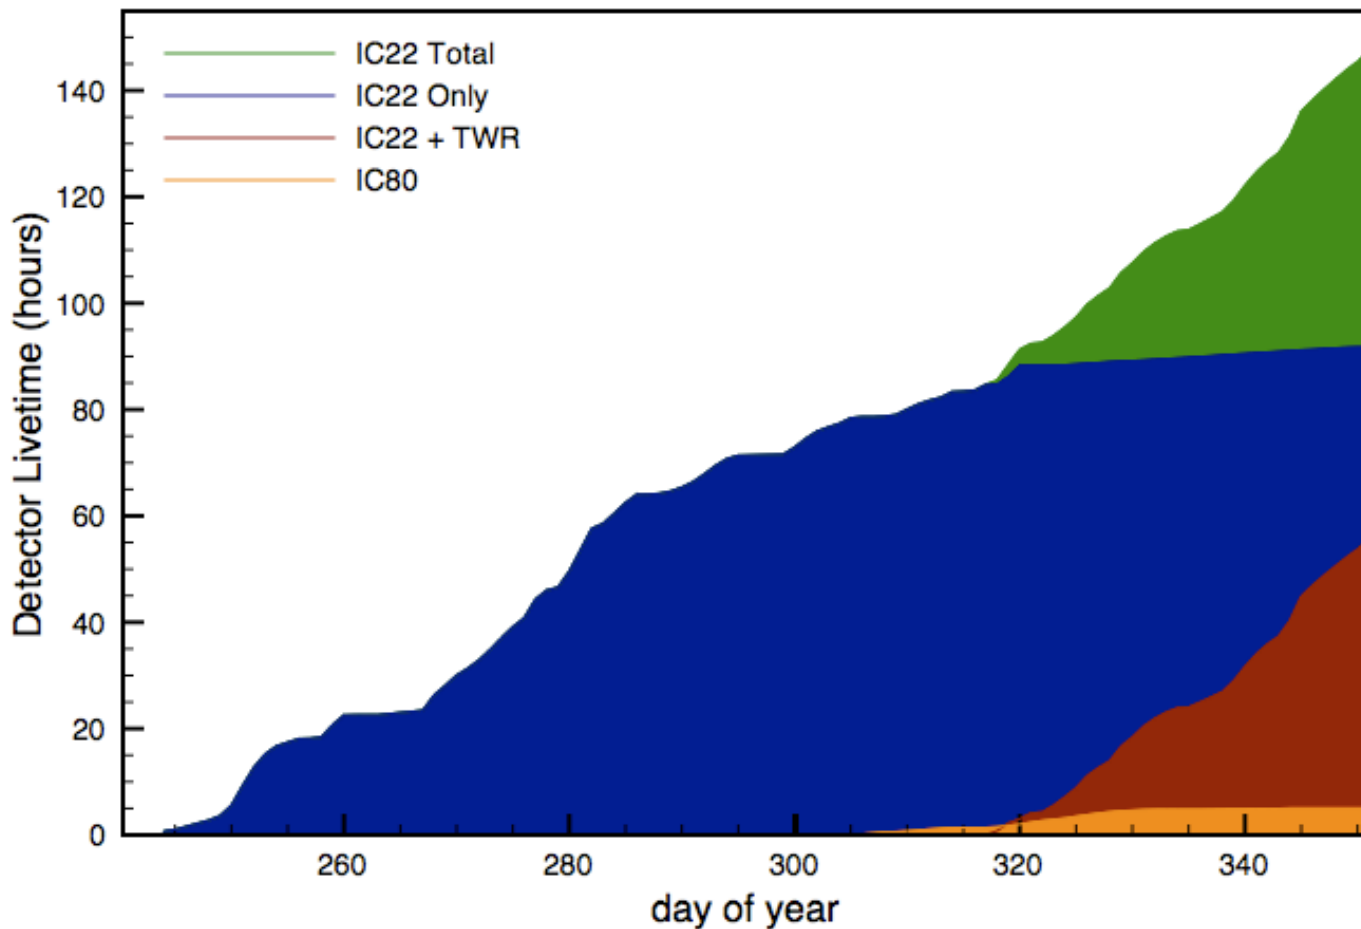
Condor Week
 May 1, 2008
 UW-Madison






simulation production statistics

Condor Week
May 1, 2008
UW-Madison






I/O

Condor Week
May 1, 2008
UW-Madison



- Unfortunately, CPU power is not everything
 - you have to be able to move data as fast as you produce it
 - at least over the long term
- Currently, all of our data is housed at UW-Madison
 - we are exploring options such as distributing data via GridFTP
 - also assigning more CPU bound simulations to sites with narrower network bandwidth



next...

Condor Week
May 1, 2008
UW-Madison

- Currently working towards using OSG
 - we now have a test submit node for OSG and will begin testing soon
 - need to resolve photonics tables problem
- LONI (Louisiana Optical Network Infrastructure)
- DAG implementation in IceProd



antarctica







new South Pole station

IceCube construction



1 gigaton particle detector,
or 1 kilometer cube or 1 teraliter

one of 21 drill modules arrive in antarctica



IceCube construction

