How to access the SiD computing resources

Jan Strube
Tohoku University

Overview

- Overview over software
- How much CPU / storage does SiD have?
- Where are the SiD resources located?
- How do we use the resources to run our jobs?

SiD on the Grid

Open Science Grid (OSG) resources are new addition to the SiD resource pool

dedicated resources: PNNL, SLAC, CERN
temporary quota increase: FNAL, RAL Tier 1
opportunistic use: all others

Worldwide LHC Computing Grid (WLCG) resources have been established during LOI and CLIC CDR efforts

SiD takes advantage of the international computing grid infrastructure

DBD Production in Numbers

Production summary on SLAC confluence

50.7 million events at 1 TeV
(+ 4.7 million gghadrons)
6.55 million events at 500 GeV
(+ 4.4 million gghadrons)

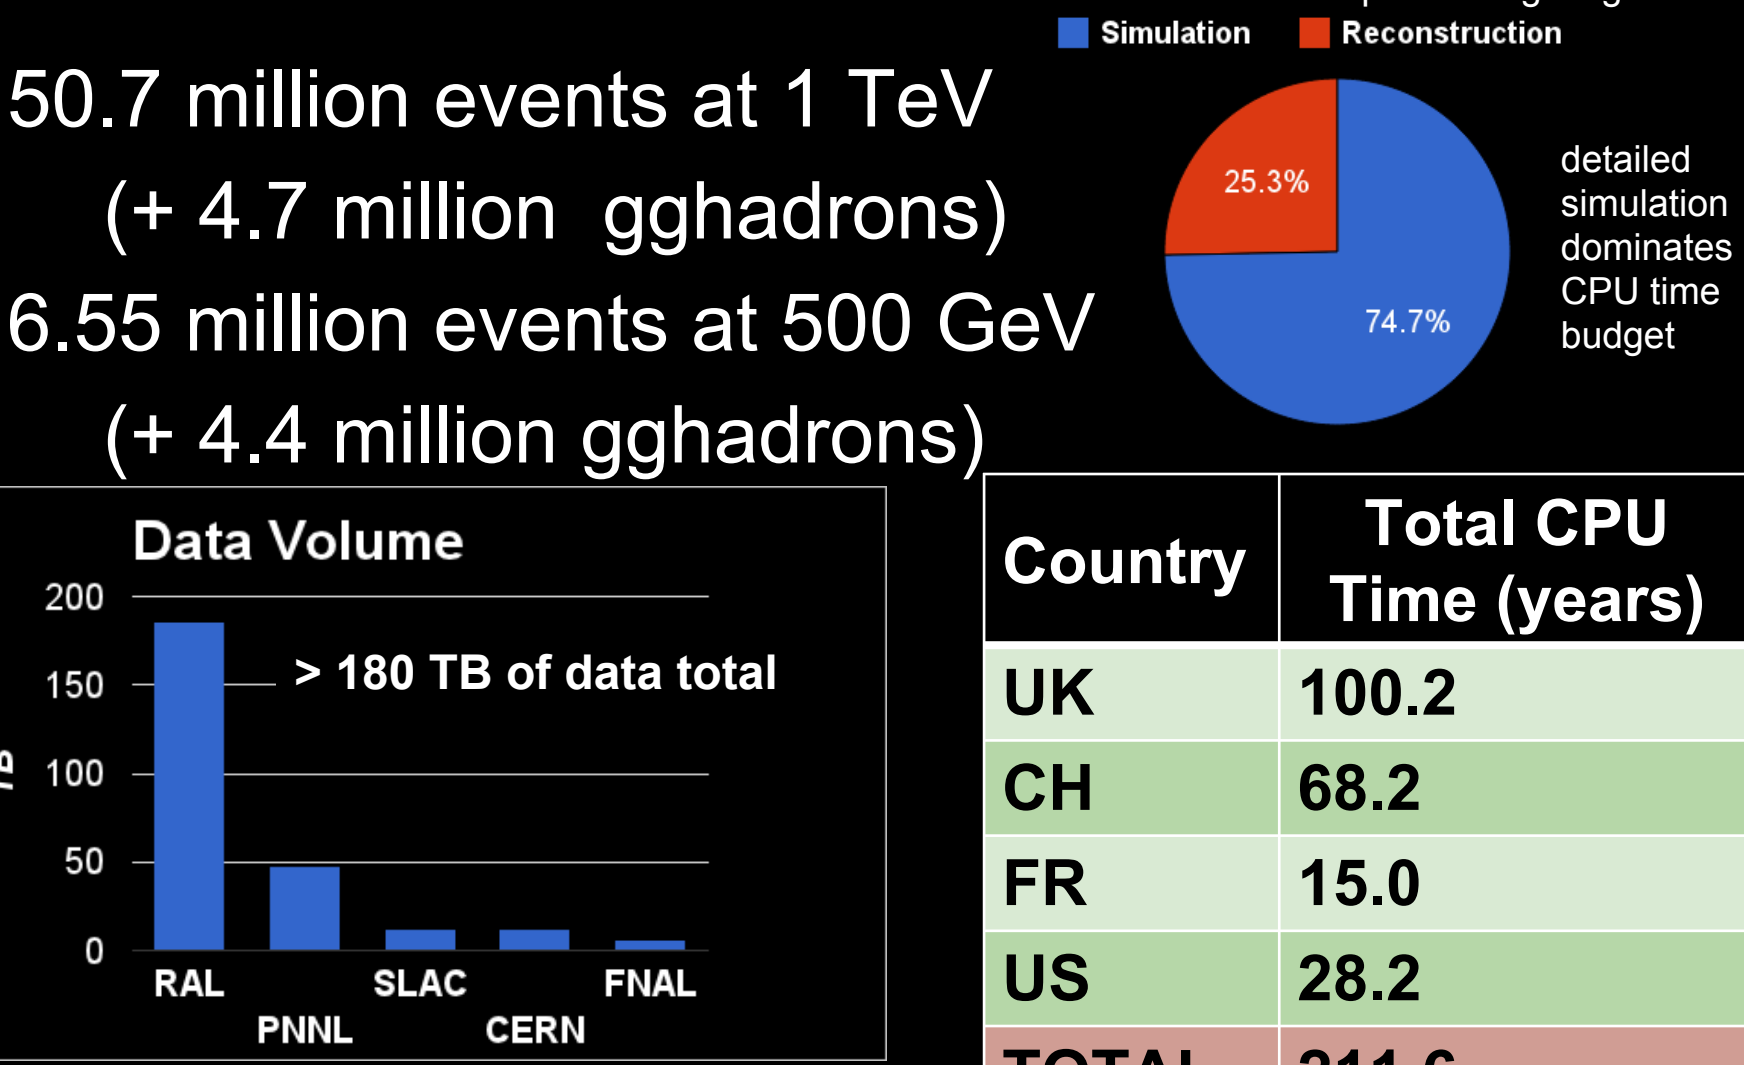

Country	Total CPU Time (years)
UK	100.2
CH	68.2
FR	15.0
US	28.2
TOTAL	211.6

Performance Summary

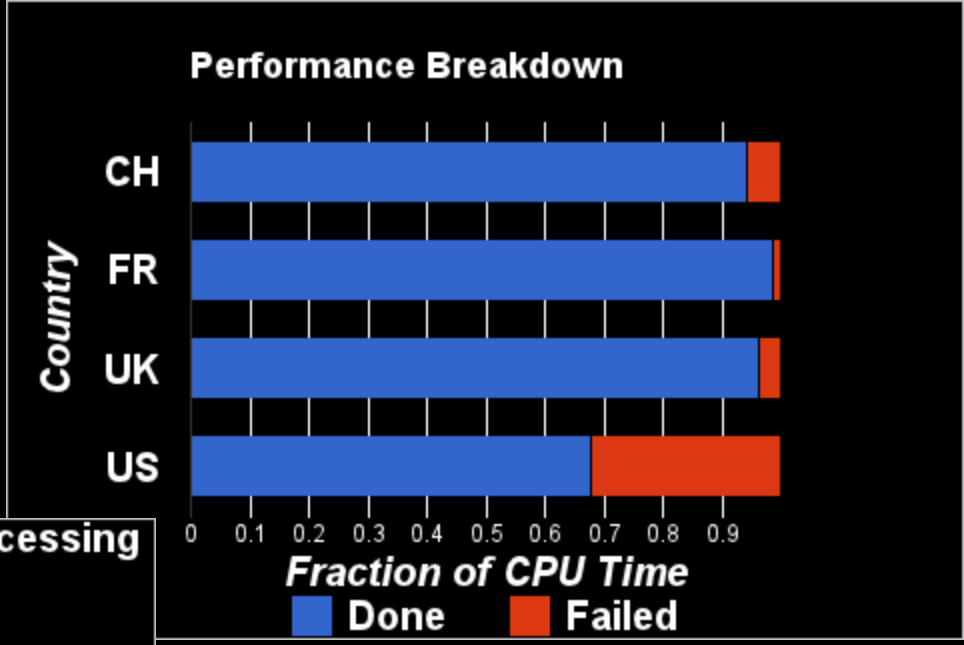

Number of Jobs

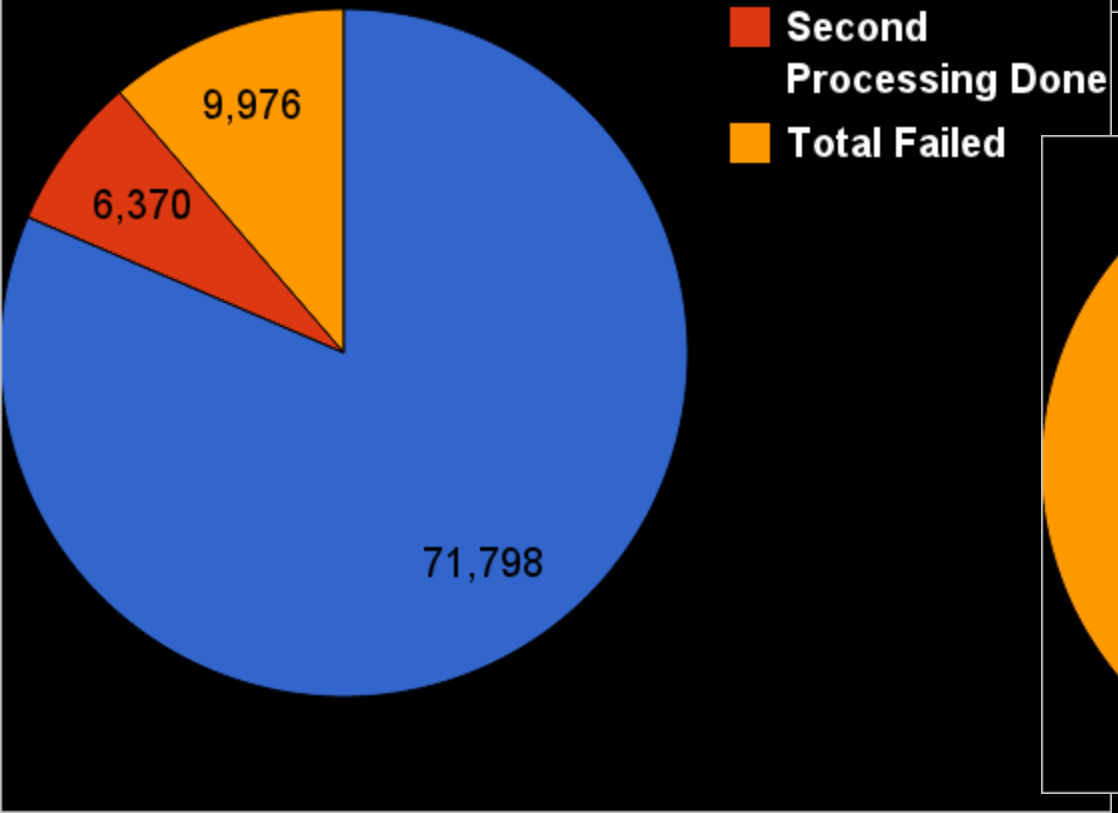

CPU Time

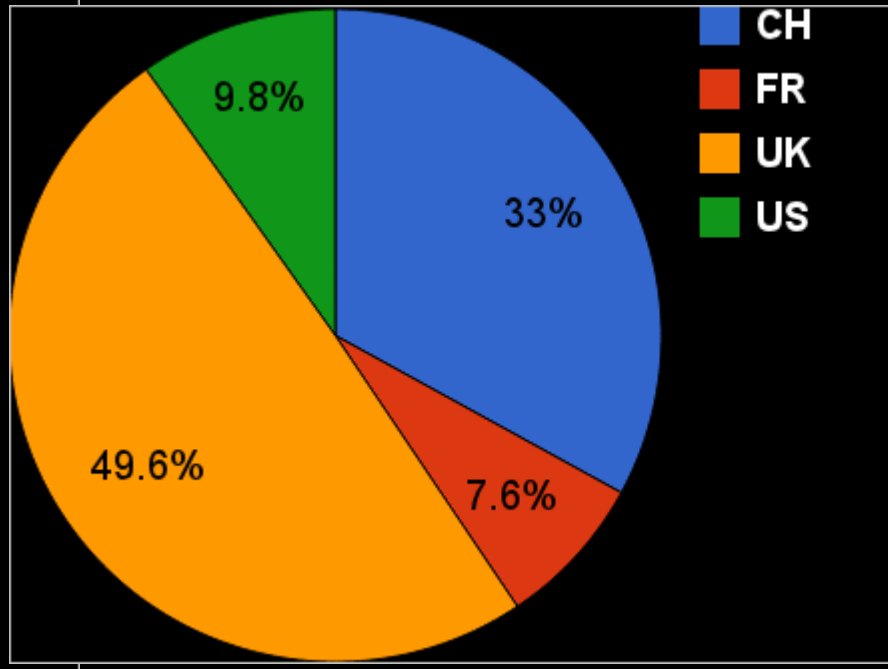

Reconstruction Workflow

1. Generate physics input (Tim Barklow)
2. Detector Simulation (SLIC)
3. “Overlay” of machine-induced background (org.lcsim)
4. Hit digitization / Track reconstruction (org.lcsim)
5. calorimetric reconstruction / PFA (slicPandora)
6. Vertex finding (LCFIPlus)
7. Data reduction / truth matching (org.lcsim)
8. Jet Finding / Flavor tagging (LCFIPlus)
9. Background reduction / Data analysis (MarLin / ROOT)

On the grid

at this point intermediate files could be dropped

On the grid or download DST files and run on local farm

Constraints

- Simulation takes the largest amount of time
  - drives CPU requirements and file sizes
  - Could insert a file-merging step? (sandbox size limit)
- Limited number of beam-induced “overlay” events
  - Many different jobs access the same event at the same time
  - Heavy burden on storage elements. Relieved by replicating / duplicating / shuffling events
- Small size of DST files (< 100 MB).
  - Good for user analysis. Bad for file transfers / tape storage

How much CPU / storage does SiD have?

- Resources on the grid are shared within the ILC virtual organization (VO)
- VO members are
  - ILD -- DESY and KEK
  - Clicdp -- CERN / (IN2P3)
  - SiD -- Everything else !
- All three are currently running detector optimization campaigns
- Storage: 150 TB at RAL (full, some used by Clicdp), 150 TB PNNL, 20 TB at SLAC
- CPU: ILC VO Total: ~10k CPU + PNNL

Takeaway message I

- The production is a complicated beast. Don't submit jobs willy-nilly.
  - We don't have a fair-share mechanism within the ILC VO.
  - Storage elements / data distribution / file access need planning
- UK resources were nearly exclusively SiD
  - Liaison now John Marshall.
- Resources for the newly formed Clicdp group have not yet been negotiated
  - Currently overlap with SiD for historical reasons
  - To be negotiated within LCC Computing group

How to access the grid

1. Obtain grid certificate from your local authority
2. Register with the ILC VO (Do not use the same certificate for ILC and LHC)
3. Register with ILCDirac and follow tutorials
   a. https://twiki.cern.ch/twiki/bin/view/CLIC/DiracForUsers
   b. https://confluence.slac.stanford.edu/display/ilc/Running+LCSim+Analysis+Jobs+on+the+Grid+with+DIRAC

Where are the SiD samples located?

Stdhep (generator level):

https://confluence.slac.stanford.edu/display/ilc/Standard+Model+Data+Samples

On the grid:

https://confluence.slac.stanford.edu/pages/viewpage.action?pageId=138785074

https://confluence.slac.stanford.edu/display/ilc/DBD+Data+Samples

Currently: ~ 10 TB DST @ 1 TeV ⇒ 51 MEvents

~ 250 GB @ 500 GeV ⇒ 6.5 MEvents

~ 200 GB @ 250 GeV ⇒ 12 MEvents

Introduction to ILCDirac

- The Grid is a heterogeneous set of computing sites
  - Different architectures, configurations, limitations
- Any tool that claims it can hide this heterogeneity from you is lying
- Dirac is a service to submit computing jobs to grid sites. Similar to your local batch farm
- ILCDirac wraps several ILC applications and executes them with your credentials on grid sites

Example ILCDirac Script (snippet)

See https://confluence.slac.stanford.edu/display/~jstrube/RecoChain.py for complete example

```
from DIRAC.Core.Base import Script
Script.parseCommandLine()
from ILCDIRAC.Interfaces.API.DiracILC import DiracILC
dirac = DiracILC(True,"some_job_repository.rep")

from ILCDIRAC.Interfaces.API.NewInterface.UserJob import UserJob
job = UserJob()
from ILCDIRAC.Interfaces.API.NewInterface.Applications import SLIC
slic = SLIC()
slic.setVersion('v2r9p8')
slic.setInputFile("LFN:/ilc/prod/ilc/some/file.stdhep")
slic.setSteeringFile('MyMacro.mac')
slic.setDetectorModel('sidloi3')
slic.setOutputFile("out.slcio")
res = job.append(slic)

job.setName("MyJobName")
job.setJobGroup("Agroup")
job.setCPUTime(86400)
job.setInputSandbox(["file1","file2"])
job.setDestination("LCG.CERN.ch")
job.setBannedSites(['LCG.DESY-HH.de', 'LCG.DESYZN.de', 'LCG.KEK.jp'])
(50000)
job.setSystemConfig('x86_64-slc5-gcc43-opt')
job.setOutputData("out.slcio", "sidloi3/analysis", "PNNL-SRM")
job.setOutputSandbox(['*.log', '*.xml', '*.lcsim', '*.steer'])

job.submit()
```

Initialization

Set up the application (SLIC)

from input sandbox

from file catalog

Deal with Grid specifics:
- Don't interfere with ILD
- block broken sites
- Make sure binaries run
- require CPU time

log files (web frontend)

grid output data

ILCDirac Features

- Submit scripts written in Python
  - Choose input files
  - reco steps mix and match
  - (semi-optionally) select sites for running
  - (semi-optionally) select site for output
- Web interface http://ilcdirac.cern.ch/DIRAC/
  - bookkeeping, restart failed jobs
- File Catalog
  - meta data search
  - find physical location of files
  - upload / download files

Takeaway message II

- The grid is advertised as fire-and-forget
  - Believe this at your own peril
- Our tools are getting better but you have to know a few details of where your jobs will go
  - Select only sites where we have negotiated resources
  - Several sites advertise they accept ILC jobs to the dirac system, but they are mis-configured, which leads to failed jobs and angry mails from admis
  - Only use storage at the main sites
    - FNAL has “volatile” storage...

Summary

- We've been using the grid very successfully.
  - Running in LHC overheads and negotiating with different admins
  - This has been pretty informal, but needs to be formalized
- Our existing configurations save you the pain of installing the correct version of all of the different software packages
  - But you still need to help us manage resources
  - And report bugs
- Go get your grid certificate and get started