Present and Future Computing Requirements for Theoretical Particle Physics

Michele Papucci (LBNL)
Stefan Hoeche (SLAC)

NERSC HEP Requirements Review

Theoretical Particle Physics & HPC

Theoretical Particle Physics & HPC

- Various areas where TH HEP needs large scale calculations:

Theoretical Particle Physics & HPC

- Various areas where TH HEP needs large scale calculations:
  - Improving accuracy of MonteCarlo event generators (development + testing + use) → necessary for experiments to estimate signal and background processes (where it is needed the greatest accuracy) at the LHC (e.g. W/Z+n jets, ttbar+jets, …)

Theoretical Particle Physics & HPC

- Various areas where TH HEP needs large scale calculations:
  - Improving accuracy of MonteCarlo event generators (development + testing + use) → necessary for experiments to estimate signal and background processes (where it is needed the greatest accuracy) at the LHC (e.g. W/Z+n jets, ttbar+jets, …)
  - Improving QCD calculations of various observables (mostly for the LHC). In general, higher accuracy reached than for event generators but more inclusive → kinematic distributions ("plots") instead of lists of collision events that can be used to simulate the detector response

Theoretical Particle Physics & HPC

- Various areas where TH HEP needs large scale calculations:
  - Improving accuracy of MonteCarlo event generators (development + testing + use) → necessary for experiments to estimate signal and background processes (where it is needed the greatest accuracy) at the LHC (e.g. W/Z+n jets, ttbar+jets, …)
  - Improving QCD calculations of various observables (mostly for the LHC). In general, higher accuracy reached than for event generators but more inclusive → kinematic distributions ("plots") instead of lists of collision events that can be used to simulate the detector response
    - fixed order: perturbative expansion in QCD coupling constant (α_s) → include more orders (e.g. NNLO top-antitop production)

Theoretical Particle Physics & HPC

- Various areas where TH HEP needs large scale calculations:
  - Improving accuracy of MonteCarlo event generators (development + testing + use) → necessary for experiments to estimate signal and background processes (where it is needed the greatest accuracy) at the LHC (e.g. W/Z+n jets, ttbar+jets, …)
  - Improving QCD calculations of various observables (mostly for the LHC). In general, higher accuracy reached than for event generators but more inclusive → kinematic distributions ("plots") instead of lists of collision events that can be used to simulate the detector response
    - fixed order: perturbative expansion in QCD coupling constant (α_s) → include more orders (e.g. NNLO top-antitop production)
    - resummation: observables depend not only on series in α_s, but on series of $\alpha_s^n \log^m$(<function of kinematic variables>) → need to resum the series of log's if log(...) is large (e.g. cross section of Higgs+0jet)

Theoretical Particle Physics & HPC

- Various areas where TH HEP needs large scale calculations:
  - Improving accuracy of MonteCarlo event generators (development + testing + use) → necessary for experiments to estimate signal and background processes (where it is needed the greatest accuracy) at the LHC (e.g. W/Z+n jets, ttbar+jets, …)
  - Improving QCD calculations of various observables (mostly for the LHC). In general, higher accuracy reached than for event generators but more inclusive → kinematic distributions ("plots") instead of lists of collision events that can be used to simulate the detector response
    - fixed order: perturbative expansion in QCD coupling constant (α_s) → include more orders (e.g. NNLO top-antitop production)
    - resummation: observables depend not only on series in α_s, but on series of $\alpha_s^n \log^m$(<function of kinematic variables>) → need to resum the series of log's if log(...) is large (e.g. cross section of Higgs+0jet)
  - Interpretation of LHC results for theoretical models of "new physics" → need to compare theory & data

Theoretical Particle Physics & HPC

- Various areas where TH HEP needs large scale calculations:
  - Improving accuracy of MonteCarlo event generators (development + testing + use) → necessary for experiments to estimate signal and background processes (where it is needed the greatest accuracy) at the LHC (e.g. W/Z+n jets, ttbar+jets, …)
  - Improving QCD calculations of various observables (mostly for the LHC). In general, higher accuracy reached than for event generators but more inclusive → kinematic distributions ("plots") instead of lists of collision events that can be used to simulate the detector response
    - fixed order: perturbative expansion in QCD coupling constant (α_s) → include more orders (e.g. NNLO top-antitop production)
    - resummation: observables depend not only on series in α_s, but on series of $\alpha_s^n \log^m$(<function of kinematic variables>) → need to resum the series of log's if log(...) is large (e.g. cross section of Higgs+0jet)
  - Interpretation of LHC results for theoretical models of "new physics" → need to compare theory & data
  - ...

Theoretical Particle Physics & HPC

- Various areas where TH HEP needs large scale calculations:
  - Improving accuracy of MonteCarlo event generators (development + testing + use) → necessary for experiments to estimate signal and background processes (where it is needed the greatest accuracy) at the LHC (e.g. W/Z+n jets, ttbar+jets, …)
  - Improving QCD calculations of various observables (mostly for the LHC). In general, higher accuracy reached than for event generators but more inclusive → kinematic distributions ("plots") instead of lists of collision events that can be used to simulate the detector response
    - fixed order: perturbative expansion in QCD coupling constant (α_s) → include more orders (e.g. NNLO top-antitop production)
    - resummation: observables depend not only on series in α_s, but on series of $\alpha_s^n \log^m$(<function of kinematic variables>) → need to resum the series of log's if log(...) is large (e.g. cross section of Higgs+0jet)
  - Interpretation of LHC results for theoretical models of "new physics" → need to compare theory & data

Important for DOE Energy Frontier Mission!

Theoretical Particle Physics & HPC

- TH HEP is new to NERSC:
  - Part of the Case Studies presented based on other computing facilities:
    - SLAC clusters
    - CERN (LXPLUS)
    - Inst. for Advanced Study (Aurora)
    - ...
  - LBNL Theory Group started using NERSC in September 2012
    - Carver & PDSF (studies based on usage for end of Sep 2012 - Nov 2012)

Welcome to the NERSC Information Management (NIM) system. Use this interface to find information about NERSC users and repositories. Authorized managers can also modify information and create or review allocation requests. Staff-only options are in **gold**.

Last NIM login on 10/30/2012 16:02:05
Your password will expire on 03/24/2013

NERSC System Status:

Available Not Available MOTD

Carver Dirac Euclid
Hopper HPSS NGF
PDSF

MPP Available Repo Balance

HPSS Available Repo Balance

Repo 7-day Report

Project	Project Repos	PI	Last Active
LBL-HEP-Th	m1610 theory	Zoltan Ligeti	2012

Theoretical Particle Physics Simulations for LHC Processes

Format: Read-only <--> Edit user percentages

NOTE: all hours displayed below are user hours, not repo hours.

m1610 MPP Users, AY 2012 <--> **Show users for prior AY**

Login	Name	User Hrs Used	User Charged	Avg CF	% Used	% Allowed	User Balance	Repo User Status	Base Repo?	Dflt Now?
alioli	Alioli, Simone	14,380	14,380	1.0	10	50	60,620	Active	Y	Y
cdfsoft	Anderson, Jeffrey	0.0	0.0	1.0	0.0	100	82,256	Active	Y	Y
cwbauer	Bauer, Christian	0	0		0	10	15,000	Active	N	N
berggren	Berggren, Calvin	19,623	19,623	1.0	13	25	17,877	Active	Y	Y
chanowit	Chanowitz, Michael	7	7	1.0	0.0	10	14,993	Active	Y	Y
zligeti	Ligeti, Zoltan	0	0		0	100	82,256	Active	Y	Y
papucci	Papucci, Michele	0	0		0	25	37,500	Active	Y	Y
ruderman	Ruderman, Joshua	0	0		0	10	15,000	Active	N	N
frank	Tackmann, Frank	0	0		0	10	15,000	Active	Y	Y
avichi	Vichi, Allesandro	84	84	1.0	0.1	25	37,416	Active	Y	Y
jwalsh	Walsh, Jonathan	33,650	33,650	1.0	22	50	41,350	Active -- Processing	Y	Y
zuberi	Zuberi, Saba	0	0		0	10	15,000	Active	Y	Y
Total:		**67,744**	**67,744**							

12 records found

Case Studies

- 5 Studies submitted:
  - Monte Carlo event generator development
    - Precision QCD Calculations with BlackHat and Sherpa (L.Dixon & S.Hoeche)
    - Monte Carlo simulation of high-energy physics processes with Geneva and POWHEG (S.Alioli, C.Bauer)
  - Implications of LHC results for theoretical models
    - Searches for Supersymmetry in the pMSSM at the LHC (J.Hewett, T.Rizzo, A.Ismail)
    - Implications of LHC results on physics beyond the Standard Model (Natural SUSY and non-degenerate squarks) (M.Papucci)
  - Log Resummation in QCD observables
    - Resummation properties of jet vetoes at the LHC (J.Walsh, S.Zuberi)

Common features

Common features

- All case studies tend to use "embarrassingly parallel" computations → many thousands of independent smallish jobs in a batch

Common features

- All case studies tend to use "embarrassingly parallel" computations → many thousands of independent smallish jobs in a batch
  - notable exception: BlackHat+SHERPA recently parallelized (100-1000 multi-process calculations)

Common features

- All case studies tend to use “embarrassingly parallel” computations → many thousands of independent smallish jobs in a batch
  - notable exception: BlackHat+SHERPA recently parallelized (100-1000 multi-process calculations)

MC Event Generators

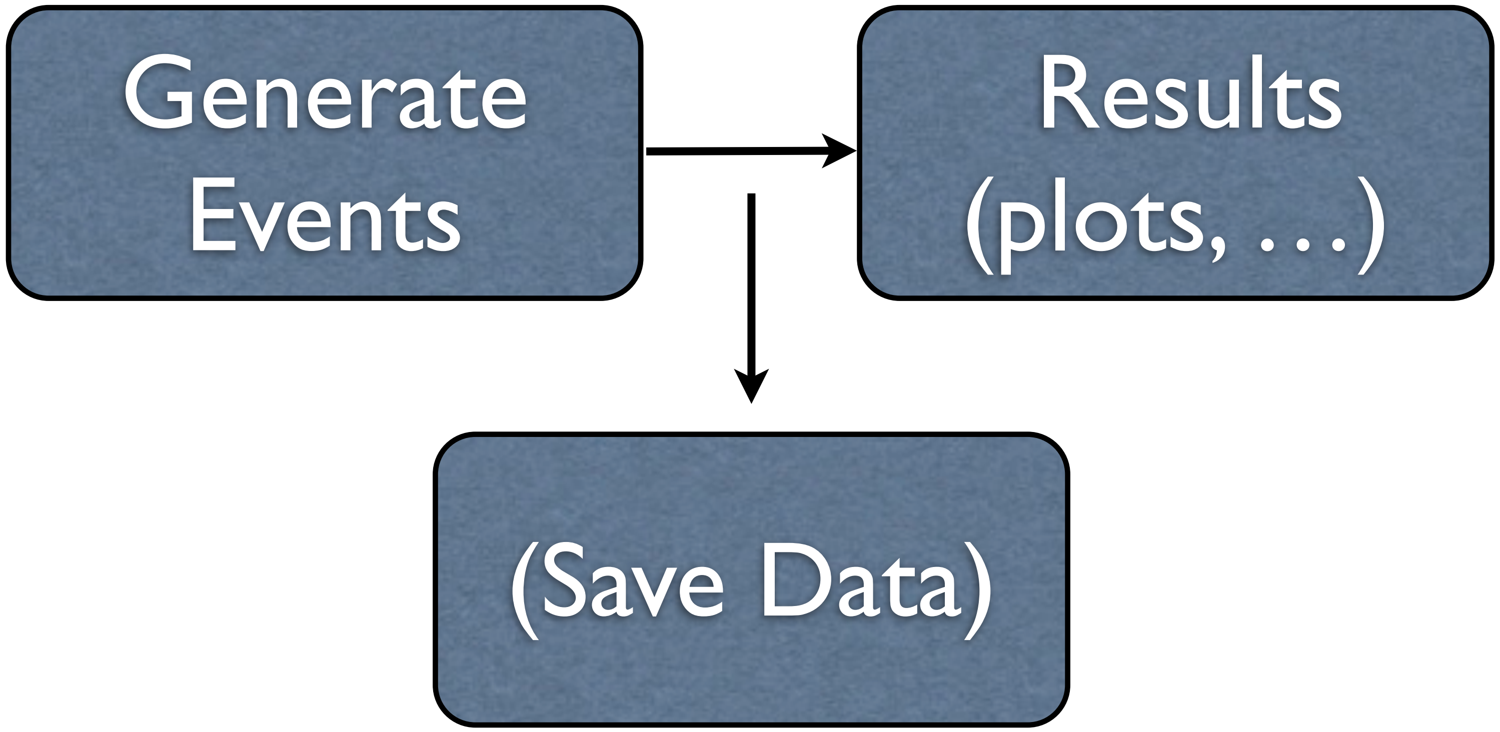

Common features

- All case studies tend to use "embarrassingly parallel" computations → many thousands of independent smallish jobs in a batch
  - notable exception: BlackHat+SHERPA recently parallelized (100-1000 multi-process calculations)

LHC results interpretation

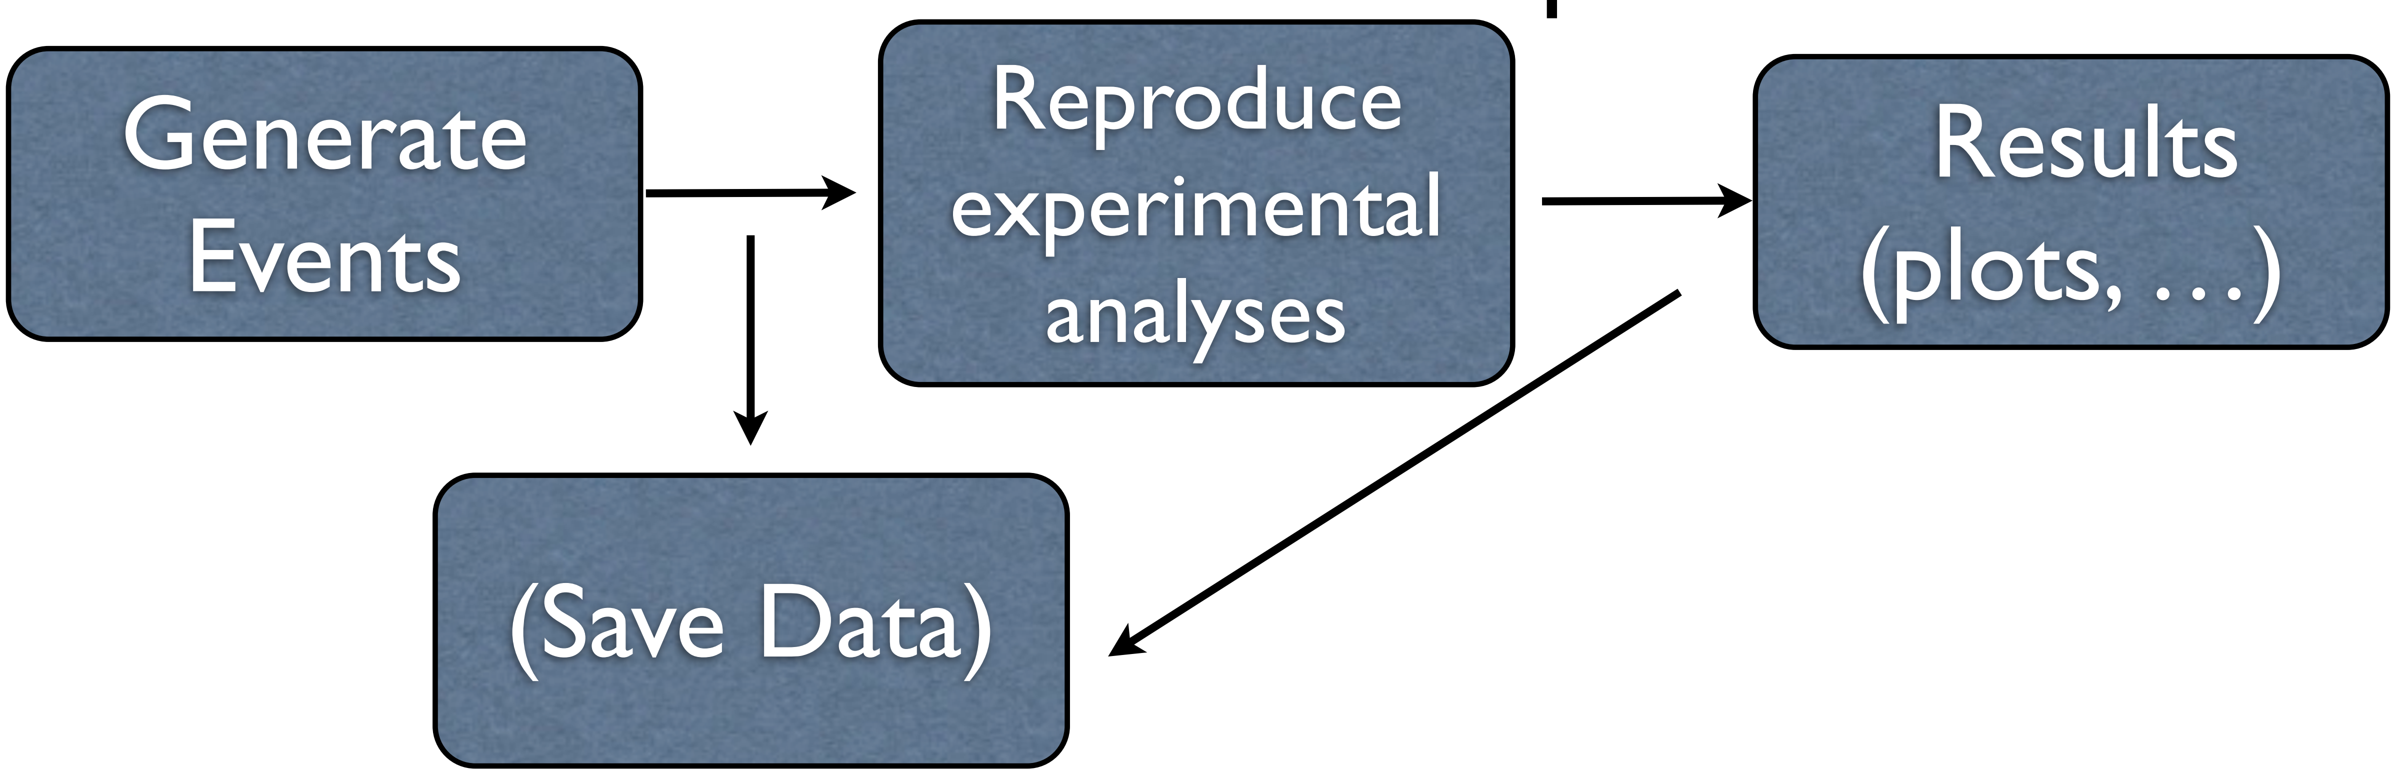

Common features

- All case studies tend to use "embarrassingly parallel" computations → many thousands of independent smallish jobs in a batch
  - notable exception: BlackHat+SHERPA recently parallelized (100-1000 multi-process calculations)

QCD Resummation

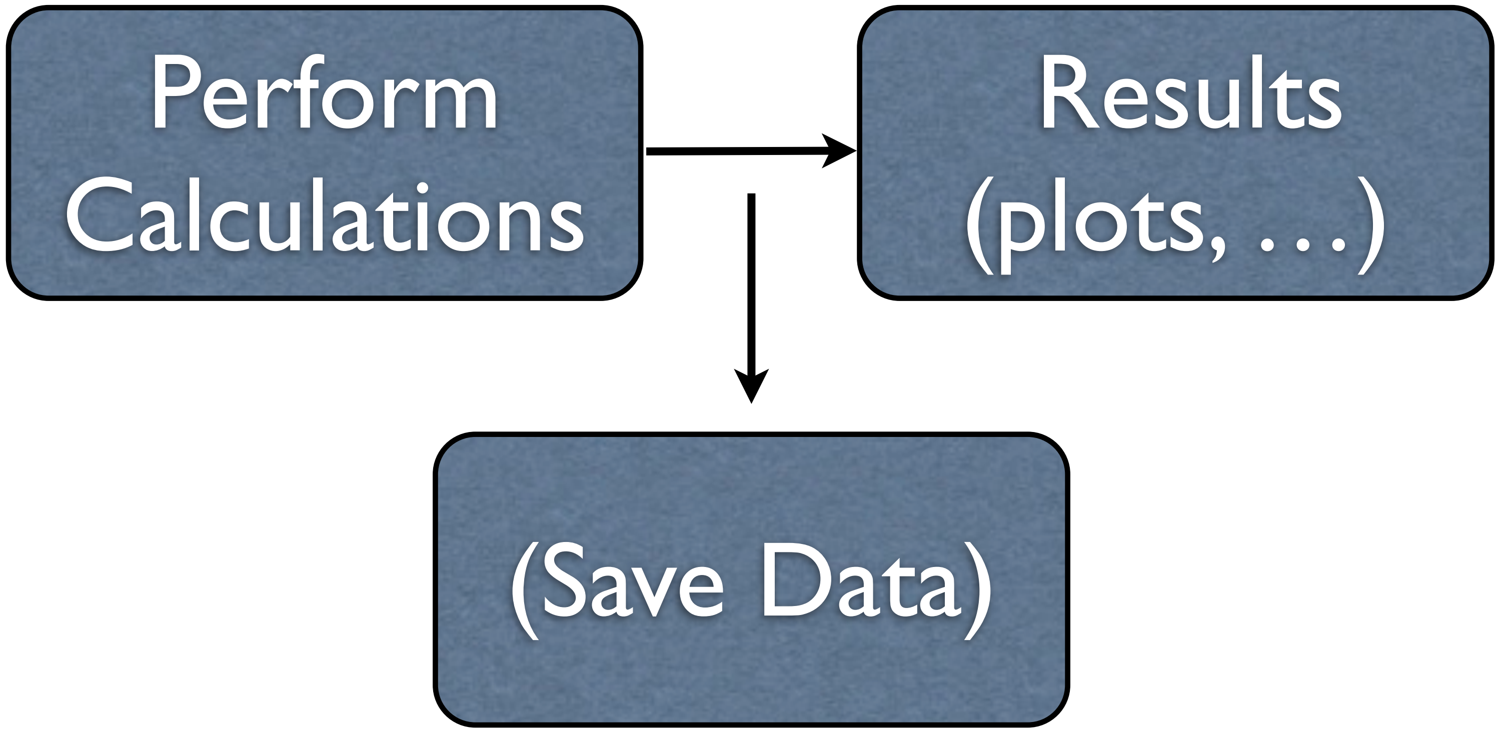

MC Event Generators

- Base: MonteCarlo algorithms using VEGAS and multi-channel phase space integration
- Tools: C++,fortran (POWHEG), python w/ interfaces to ROOT, LHAPDF, fastjet libraries. MPI (BlackHat), POSIX Threads (SHERPA) for parallelization
  - BlackHat+SHERPA: 2012 results based on non-NERSC cluster + Grid storage
  - POWHEG, Geneva: Sep 2012 - Nov 2012 Carver & PDSF usages, single user (Alioli). Geneva project x2-3

MC Event Generators

	in 2012	Sep 2012- Nov 2012 on Carver & PDSF	in 2017	@ NERSC
	BlackHat +Sherpa	Geneva, POWHEG (Alioli only)	BlackHat +Sherpa	Geneva, POWHEG (1 user)
Hours	3M	30k	9M	1M
#cores/run	128 (parallel)	150 (not-parallel, typical)	1024 (parallel)	200 (not-parallel, typical)
max # cores	128 x #jobs	200	1024 x #jobs	1k
I/O per run	<50GB	20GB	<100GB	100GB
I/O bandwidth	<1GB/sec	1-10GB/sec	<1GB/sec	1-10GB/sec
Runtime % for I/O	<1%	5-10%	<1%	5-10%
Shared space	30TB	40TB	100TB	100TB
Archival Data	10TB	0TB	300TB	20TB
Mem/node	2-8GB	1-3GB	4-16GB	4-8GB
Aggregate Memory		60TB		4PB

LHC result interpretation

- Base: MonteCarlo event generators + detector simulation software (pMSSM) + analysis software
- Tools: C++,fortran, python w/ interfaces to ROOT, LHAPDF, fastjet libraries. Pythia for event generation, PGS (pMSSM) for detector simulation, Mathematica (natSUSY) used in part of the batch processing
  - pMSSM: 2012 results based on SLAC cluster
  - natSUSY+…: results from Oct 2011-Oct 2012 based on LXPLUS (CERN)+Aurora(IAS) clusters. Moving on Carver+PDSF next month

LHC Result Interpetation

	in 2012		in 2017	@ NERSC
	pMSSM	NatSUSY + ...	pMSSM	NatSUSY + ...
Hours	800k	500k (aggregate)	1M	1-5M
typical #cores/run	3k	30-9k	5k	5k
max # cores	200k	10k/batch	200k	10k/batch
I/O per run	0.5TB	0.3TB/job	2.0TB	0.3TB/job
I/O bandwidth	0.1GB/sec	0.1GB/sec	0.1GB/sec	0.1GB/sec
Runtime % for I/O	variable	variable	variable	variable
Shared space	2TB	0.1TB (cluster limitations)	4TB	30TB
Archival Data	4.5TB	0TB	10TB	60TB
Mem/node	1-3GB	2-4GB	1-3GB	4-8GB

QCD Resummation

- Tools: C++,fortran w/ interfaces to LHAPDF, fastjet libraries. Use of Mathematica in batch jobs to perform part of the calculations
  - Sep 2012 - Nov 2012 results based on Carver+ PDSF
  - Current Mathematica bottleneck on Carver: 4 licenses rarely available in batch jobs → fall back on Euler cluster (LBNL theory group)

QCD Resummation

	Sep 2012- Nov 2012 on Carver & PDSF	2017 @ NERSC
Hours	30k + 10k	50-250k (single user)
typical #cores/run	150	200
max # cores	200	~1k
I/O per run	1GB	1-4GB
I/O bandwidth	small	small
Runtime % for I/O	negligible	negligible
Shared space	1TB	1TB
Archival Data	0TB	1TB
Memory	1-3GB/core	4GB/core

Other TH HEP usages

- no Case Study submitted for fixed order multiloop but typical usage:

Computing usage summary:

Project: Top-pair total inclusive cross-section in NNLO QCD

Authors: Michael Czakon and Alex Mitov

- Time frame: 1 year. Includes:
  - Development
  - Actual computation
- Nature of computing problem:
  - Parallelizable.
  - Average load: of the order of 200 jobs running over 1 year. But it <u>fluctuates</u>.
  - Typical length of individual jobs ~ hours - day(s).

Future projects:

- Differential NNLO production of tops and other particles.
- Estimate of computational needs: same order of magnitude.
- Time frame: years.

Computing Frontier: pQCD Working Group   Alexander Mitov   Fermilab, 12 Oct 2012

Summary

- HEP Theory needs large scale computations:
  - mostly LHC-related work:
    - perturbative QCD calculations:
      - fixed order & resummation
      - Monte Carlo event generators
    - Study of LHC results & their implications for new physics
  - mostly batches of large number of serial jobs, but few cases of multiprocess calculations (BlackHat) (0.1-1k cores)
- By 2017 typical project/paper will need of the order of 1M hours, with order of 100GB-1TB of I/O per job and up to 100 TB of space, with 4-16GB of memory/node

Backup

Case Study Title: Monte Carlo simulation of high-energy physics processes with Geneva and POWHEG

Principal Investigator: Simone Alioli, Christian Bauer (Zoltan Ligeti for NERSC usage)
Worksheet Author (if not PI): Simone Alioli
NERSC Repositories: m1610

1. Project Description

1.1 Overview and Context

Monte Carlo simulation of high-energy physics processes for the LHC. Development and testing of Monte Carlo event generators Geneva and POWHEG. Requires massive serial batch jobs system (1k - 10k cores) and facilities to store large amount of data ~100TB

1.2 Scientific Objectives for 2017

Improve the theoretical accuracy of the description of both inclusive and exclusive observables, obtained via a Monte Carlo event simulation. Implement new processes and perform sound phenomenological studies and data comparisons.

2. Computational Strategies (now and in 2017)

2.1 Approach

"Embarassingly parallel" runs. Thousands of jobs submitted in batch, each job lasting few hours up to a day.

2.2 Codes and Algorithms

The softwares ran are Geneva and POWEG. They are based on Fortran and C++ codes, currently non-parallelized. They are run as multiple, independent batch jobs.

3. HPC Resources Used Today (2012)

3.1 Computational Hours

Used ~30k hours from Sep 2012 to today on Carver and PDSF

3.2 Compute Cores

Conventional # of cores for production run is 100-200. 100-200 is the maximum on Carver, while I used 1k cores on PDSF when available

3.4 Shared Data

Shared data are used to access information and write output from both Carver and PDSF clusters

3.5 Archival Data Storage

Not yet used, since new to the system. Eventually publication-related data will be archived there

4. HPC Requirements in 2017

4.1 Computational Hours Needed

Current usage will be around 250k hours per year, may increase by a factor of 2-4 by 2017.

4.2 Number of Compute Cores

I will mainly use serial queues, better with ~1k cores available

4.3 Data and I/O

Combined throughtput of data written/read on/from the filesystem will easily reach 500TB - 1 PB

4.4 Shared Data

Estimated NERSC project directory space will be 100TB in 2017

4.5 Archival Data Storage

Estimated archival storage will be 20TB in 2017

4.6 Memory Required

Estimated memory per node will be 4GB in 2017.

Requirements Summary Worksheet

Please try to fill out this worksheet based on your answers above to be best of your ability prior to the workshop.

	Used in 2012 (Sep-Nov only)	Needed at NERSC in 2017
Computational Hours	30k	1M
Typical number of cores* used for production runs	150	200
Maximum number of cores* that can be used for production runs	200	~1k
Data read and written per run	0.02TB	0.1TB
Maximum I/O bandwidth	1-10GB/sec	1-10GB/sec
Percent of runtime for I/O	5-10%	5-10%
Shared filesystem space	40TB	100TB
Archival data	0TB	20TB
Memory per node	1-3GB	4-8GB
Aggregate memory	60TB	4PB

* "Conventional cores." For GPUs and accelerators, please fill out section 4.7.

Case Study Title: Precision QCD Calculations with BlackHat and Sherpa

Principal Investigator: Lance Dixon
Worksheet Author (if not PI): Stefan Hoeche
NERSC Repositories:

1. Project Description

1.1 Overview and Context

The aim of our project is to construct a Monte-Carlo event generator for the accurate theoretical prediction of exclusive reactions involving jets at hadron colliders. This event generator is available to HEP theorists and experimentalists. It is used by the ATLAS, CMS, D0, CDF and H1 collaborations to maximize the discovery potential of Energy Frontier experiments.

The overall project is composed of two parts, which can be characterized as a general event generation framework (Sherpa), and a library for the computation of scattering amplitudes at the next-to-leading order (NLO) in perturbative QCD (BlackHat). Predictions can be generated at the parton level or at the particle level.

BlackHat and Sherpa have been used extensively in the first years of LHC running to analyse data and to predict signal and background levels for future measurements.

1.2 Scientific Objectives for 2017

Measurements of Standard Model parameters and searches for new physics at hadron colliders call for precise theoretical predictions in reactions involving many jets. State of the art particle level calculations are mostly based on leading order perturbative QCD. Within the next few years they will be superseded by calculations based on next-to-leading perturbative QCD. Methods based on next-to-next-to leading order perturbative QCD are under investigation. These improvements are mandatory to satisfy requirements of current Energy Frontier experiments like those carried out by ATLAS and CMS and of potential future experiments.

BlackHat and Sherpa play a leading role in this endeavor. The two programs have been used to compute the first NLO predictions for the production of W/Z+3 and W/Z+4-jet final states. The first working technique to combine NLO calculations of varying jet multiplicity was implemented using BlackHat and Sherpa. The programs will be extended in the following years to tackle more complicated processes and to provide predictions at even higher accuracy.

2. Computational Strategies (now and in 2017)

2.1 Approach

BlackHat and Sherpa are stand-alone programs written entirely in C++. They include Python API's and provide interfaces to external libraries like Root and LHAPDF. The calculation is performed as a Monte-Carlo integration using adaptive sampling algorithms like VEGAS and the multi-channel method.

2.2 Codes and Algorithms

BlackHat makes use of arbitrary precision numbers in the implementation of GMP and of the QD library for quadruple precision arithmetics. It implements the OpenMP standard to parallelize the computation and will be extended in the near future to outsource parts of the calculation to GPU's.

Sherpa uses the POSIX thread library to parallelize the calculation of scattering matrix elements. The adaptive Monte-Carlo integration is parallelized using the MPI standard, which allows scalable parallel computing.

3. HPC Resources Used Today (2012)

3.1 Computational Hours

We use about 150k CPU hours per project.
There are typically one or two such projects per month.

3.2 Compute Cores

We use of the order of 128 compute cores per job and multiple jobs per production run. The maximum number of cores is limited only by network bandwidth. We have used only up to 256 cores so far, due to limited availability of large scale parallel computing facilities.

3.4 Shared Data

Our project has mechanisms for sharing data among collaborators. Sherpa uses dynamically linked libraries for runtime extensions, which are user-independent and relocatable. BlackHat consists of a core program and data files, which are user-independent as well.

3.5 Archival Data Storage

Our project archives data on Grid storage and on local storage systems at various universities and national laboratories. We currently use about 30TB of disk space.

4. HPC Requirements in 2017

4.1 Computational Hours Needed

Our computing requirements are expected to increase in 2017 due to the use of next-to-next-to leading order perturbative QCD.

4.2 Number of Compute Cores

Our codes will likely use about 1024 compute cores in parallel. The maximum will be limited by network bandwidth and by the fault tolerance of MPI. We will run multiple jobs concurrently and we expect the additional use of GPU's.

4.3 Data and I/O

We expect to write less than 100GB of data per run in a mostly continuous stream.
We do not have specific requirements regarding the speed of I/O.

4.4 Shared Data

We will require of the order of 100TB shared disk space.

4.5 Archival Storage

We expect to use up to 300TB of archival storage.

4.6 Memory Required

We expect shared memory requirements of 4-16GB per compute core.

4.7 Many-Core and/or GPU Architectures

We are currently working on an extension of BlackHat to make use of GPU's.
Sherpa will enable the use of GPU's as needed.

4.8 Software Applications and Tools

We expect to use C++ and FORTRAN compilers, OpenMP, Python, Root and MPI.
Our build systems will need the autotools and possibly scons.

4.9 HPC Services

We will need account support and usage statistics.

4.10 Time to Solution and Throughput

No special requirements.

4.11 Data Intensive Needs

No special requirements.

4.13 What Else?

No other requirements.

Requirements Summary Worksheet

Please try to fill out this worksheet based on your answers above to be best of your ability prior to the workshop.

	Used in 2012	**Needed at NERSC in 2017**
Computational Hours	3000000	~ 90000000
Typical number of cores* used for production runs	128	1024
Maximum number of cores* that can be used for production runs	bandwidth limited	bandwidth limited
Data read and written per run	<0.05TB	<0.1TB
Maximum I/O bandwidth	<1GB/sec	<1GB/sec
Percent of runtime for I/O	<1	<1
Shared filesystem space	30TB	100TB
Archival data	10TB	300TB
Memory per node	2-8GB	4-16GB
Aggregate memory	#nodes x mem per node	#nodes x mem per node

* “Conventional cores.” For GPUs and accelerators, please fill out section 4.7.

Case Study Title: Searches for Supersymmetry in the pMSSM at the LHC

Principal Investigator: JoAnne Hewett, Tom Rizzo
Worksheet Author (if not PI): Ahmed Ismail
NERSC Repositories:

1. Project Description

We look at signals of models of new physics, focusing on supersymmetry at the Large Hadron Collider (LHC). Even in its most minimal form, supersymmetry has over 100 free parameters, and supersymmetric models in different regions of parameter space can have very different experimental signatures. Experimental limits on supersymmetry are often presented in terms of very constrained versions of this large parameter space.

2. Computational Strategies

For our study, we generate sets of ~200-250 K models, and simulate particle physics events at the LHC for each one. For a single set of models, the simulation data takes up about 0.5-1 TB, and this is expected to increase by a factor of ~4 as we will need to simulate more events to keep up with the LHC over the next five years. We have several different sets of models, and it is useful to archive data from previous sets for comparison purposes. Right now, we make use of all the cores that are available to us, and are able to simulate 3000-5000 models at a time. The simulation software that we use is not optimized for multi-core processors or GPUs, and takes a few hours to run on average for a single model. In principle, we could make use of as many conventional cores as we have models in a given set, i.e. several hundred thousand.

4. HPC Requirements in 2017

In summary, we expect to need approximately 1 million conventional CPU-hours and 4 TB of storage, with another several TB for archival, by 2017. With these resources, we can study the ability of the LHC to discover or exclude general supersymmetric models, beyond the simplified limits that are normally presented by experimental collaborations. In addition, we can investigate the complementarity of collider and astrophysical experiments in searching for new physics, an important endeavor that is currently not generally performed by collaborations.

Requirements Summary Worksheet

Please try to fill out this worksheet based on your answers above to be best of your ability prior to the workshop.

	Used in 2012	**Needed at NERSC in 2017**
Computational Hours	800000	1000000
Typical number of cores* used for production runs	3000	5000
Maximum number of cores* that can be used for production runs	200000	200000
Data read and written per run	0.5 TB	2.0 TB
Maximum I/O bandwidth	0.1 GB/sec	0.1 GB/sec
Percent of runtime for I/O	variable	variable
Shared filesystem space	2.0 TB	4.0 TB
Archival data	4.5 TB	10 TB
Memory per node	GB	GB
Aggregate memory	TB	TB

* “Conventional cores.” For GPUs and accelerators, please fill out section 4.7.

Case Study Title: Implications of LHC results on physics beyond the Standard Model (Natural SUSY and non-degenerate squarks)

Principal Investigator: Michele Papucci
Worksheet Author (if not PI):
NERSC Repositories:

1. Project Description

1.1 Overview and Context

Study of the consequences of the LHC results for theoretical models of new physics at the TeV scale. Currently studied: the status of Natural SUSY and the current LHC limits on non-flavor degenerate squarks.

1.2 Scientific Objectives for 2017

Perform detailed studies of LHC (and Dark Matter indirect detection in particle astrophysics experiments) results and draw implications for various theoretical models of physics beyond the Standard Model.

2. Computational Strategies (now and in 2017)

2.1 Approach

"Embarrassingly parallel" runs on serial queues. Thousands of jobs submitted in batch, each job lasting 1-8 hours. Each job correspond to the analysis of a particular LHC process (or a set of those), for a specific theoretical model and specific values of its parameters. A parameter scan (1D, 2D or 3D) is usually performed. Size of the parameter grids vary, but usually O(100-1k) points for 2D grids

2.2 Codes and Algorithms

C++/Fortran non-parallelized standard HEP software packages. Standard Monte Carlo event generators for generating events LHC processes. C++/python analysis code currently under development (named ATOM). Usage of Mathematica for symbolic calculation or data processing sometimes needed.

3. HPC Resources Used Today (2012)

3.1 Computational Hours

Projects performed so far did not use NERSC clusters yet. Used LXPLUS cluster at CERN, the Euler cluster of the LBNL Theory Group in the Physics Division, the Aurora cluster at the Institute for Advanced Study. Overall about 150k-300k hours per project, two projects completed in the last year. Running time distributed highly unevenly during the year. Will start using Carver and PDSF in the near future

3.2 Compute Cores

As many as they are available, mainly limited by cluster usage: 100-200 concurrently used cores on Aurora, 30 to 9000 concurrently used cores on LXPLUS.

3.4 Shared Data

Each jobs reads/writes up to 100-500MB of data on temporary disk space available on the node. When available FIFO pipes or ramdisks have been used to speed up job runs. Small amount of final output (up to 10MB/job), mostly due to cluster policy constraints (LXPLUS scratch limited to 20GB/user)

4. HPC Requirements in 2017

4.1 Computational Hours Needed

Given the increase on number of projects and usage of more accurate MC event generators which are more time-consuming, a factor of 5-10 increase in the number of hours is expected already from next year. Therefore I expect to use 1-5M hours per year by 2017.

4.2 Number of Compute Cores

I will mainly use serial queues, at least 1k, as many as are available. Batch runs will likely consist of 5-10k jobs, each lasting from 1h up to a day.

4.3 Data and I/O

Similar to present runs, about 0.5GB per job

4.4 Shared Data

Possibility of saving 100-300 MB per job during the course of a project. Assuming realistic numbers of 10k jobs/run and 10 runs for a project, a minimum of 30TB would be required.

4.5 Archival Data Storage

30-60TB.

4.6 Memory Required

Estimated memory per core will be 1-2GB in 2017, possibly larger by another 1-2GB if ramdisks are supported to speed up I/O.

4.8 Software Applications and Tools

Besides currently available software, a sufficiently large number of Mathematica licenses is needed. Batches of few thousands jobs using Mathematica will be present, each session lasting up to an hour.

	Used in Sep 2011-Oct 2012	**Needed at NERSC in 2017**
Computational Hours	About 500k, distributed on various clusters	1-5M
Typical number of cores* used for production runs	30-9000, depending on availability	5k
Maximum number of cores* that can be used for production runs	Varies, 1 batch is about 10000 jobs	10k
Data read and written per run	0.3TB/job	0.3TB/job
Maximum I/O bandwidth	small	small
Percent of runtime for I/O	10-20%	10-20%
Shared filesystem space	0.1TB	30TB
Archival data	0TB	60TB
Memory per node	1-3GB/core	4GB/core
Aggregate memory	N/A	N/A

* "Conventional cores." For GPUs and accelerators, please fill out section 4.7.

Case Study Title: Understanding resummation of jet vetoes at the LHC

Principal Investigator: Jonathan Walsh, Saba Zuberi
Worksheet Author (if not PI): Michele Papucci (based information provided by J.Walsh)
NERSC Repositories: m1610

1. Project Description

1.1 Overview and Context

QCD calculation of LHC processes: effects of resummation of large logarithms on observables containing explicit requirements on jets (e.g. veto on the number of jets)

1.2 Scientific Objectives for 2017

Improve the theoretical accuracy of QCD predictions for LHC processes by resumming large logarithms of various kinematic variables. Effects of vetoes on the number of jets. Study of non-global logarithms.

2. Computational Strategies (now and in 2017)

2.1 Approach

"Embarrassingly parallel" runs on serial queues. Thousands of jobs submitted in batch, each job lasting few hours up to a day.

2.2 Codes and Algorithms

C++/Fortran non-parallelized HEP software packages. They are run as multiple, independent batch jobs. The use of Mathematica for data processing in each of the jobs is also necessary in some of the projects, but is currently constrained by the presence of only 4 licences for the whole Carver cluster, which are often in use.

3. HPC Resources Used Today (2012)

3.1 Computational Hours

Used ~33k hours from Sep 2012 to today on Carver and ~10k on PDSF. I estimate a yearly usage between 50k to 250k hours, depending on the number projects both now and in 2017.

3.2 Compute Cores

As many as they are available. Conventional # of cores for production run is 100-200.

3.4 Shared Data

Small amount of output data (about 1MB per job).

4. HPC Requirements in 2017

4.1 Computational Hours Needed

50k-250k hours per year.

4.2 Number of Compute Cores

I will mainly use serial queues, better with ~1k cores available.

4.3 Data and I/O

Similar to present runs (few MBs per job)

4.4 Shared Data

Negligible on NERSC scales.

4.6 Memory Required

Estimated memory per core will be 1-3GB in 2017.

4.8 Software Applications and Tools

Besides currently available software, a larger number of Mathematica licenses is needed. At least 50 available parallel licenses per batch run will be needed.

Requirements Summary Worksheet

Please try to fill out this worksheet based on your answers above to be best of your ability prior to the workshop.

	Used in 2012 (Sep-Nov only)	Needed at NERSC in 2017
Computational Hours	30k+10k	50-250k
Typical number of cores* used for production runs	150	200
Maximum number of cores* that can be used for production runs	200	~1k
Data read and written per run	1GB	1-4GB
Maximum I/O bandwidth	small	small
Percent of runtime for I/O	negligible	negligible
Shared filesystem space	1TB	1TB
Archival data	0TB	1TB
Memory per node	1-3GB/core	4GB/core
Aggregate memory	N/A	N/A

* "Conventional cores." For GPUs and accelerators, please fill out section 4.7.