# EXECUTING WORKFLOWS

Thomas Uram
Data Science group, Argonne Leadership Computing Facility

WHY SHOULD YOU CARE ABOUT WORKFLOWS?

▪ You came to the SDL workshop to learn to use large-scale parallel systems for simulation, data, and learning:

- architectures
- programming models
- communication
- solvers

- visualization
- profiling
- optimizing for node-level performance
- scaling parallel applications to tens of thousands of nodes

▪ Once you've written your application and achieved near-peak node-level performance and scalability, what comes next?

Argonne
NATIONAL LABORATORY

WORKFLOWS

We developed a new intersection control paradigm called AIM.

ALLOCATION EXAMPLES

- Small number of large jobs running on Mira over course of the year followed by some analysis
  - 40 jobs x 16K nodes * 16 cores/node * 12 hours ~= 100M core-hours
  - ~Reasonable to manage manually
- Large number of simulation+analysis jobs running across multiple facilities, requires coordination of jobs submitted to multiple schedulers, large/long data transfers, and interaction with project storage and archival storage
  - Programmatic management would be a clear benefit

  Workflows can save you!
  …but you might have to code them yourself!

WHY USE WORKFLOWS?

- Automate job submission

- Simplify computational campaigns

- Increase concurrency by disentangling data dependencies

- Robustness: Improve error handling and recovery (retries)

- Coscheduling of multiple resources

- Systematize data management

- Provenance/Metadata tracking
  - Validation
  - Reuse

WHAT IS A WORKFLOW?

▪ It depends who you ask

▪ Basically a collection of jobs to be run
  – Could be a sequence of individual jobs
  – Could be a sequence of varying numbers of jobs

▪ Available means of describing and running workflows
  – Script jobs
  – Job dependencies
  – Ensemble jobs
 – In situ
  – Custom workfows
  – Workflow software

Argonne
NATIONAL LABORATORY

WORKFLOW VIA SCRIPT JOB

j1

j2

j3

one long job

▪ You can submit a job that executes your application
  – qsub -q cache-quad -n 512 -t 60 -A yourproject application.exe

▪ Alternatively, you can submit a script job that executes
  multiple applications sequentially
  – qsub -q cache-quad -n 512 -t 60 -A yourproject script.sh

script.sh

```
aprun -n 512 -N 1 application.exe
aprun -n 512 -N 1 application.exe
```

▪ With this approach, you wait in the queue one time to run your application multiple times
▪ However, small long jobs tend to stay in the queue longer than small short jobs
  – It may be better to submit individual jobs with dependencies

Argonne
NATIONAL LABORATORY

WORKFLOW VIA COBALT JOB DEPENDENCIES

- A simple way to achieve a linear workflow is simply to set dependencies between your jobs
  - qsub -q cache-quad -n 512 -t 60 -A yourproject a.out

    Job 12345 submitted
  - qsub -q cache-quad -n 512 -t 60 -A yourproject --dependencies 12345 a.out
  - qsub -q cache-quad -n 512 -t 60 -A yourproject --dependencies 12345:12346 a.out

- Is there an advantage to setting job dependencies?
  - Dependent jobs accumulate score more quickly
- How many jobs can be submitted at once?

```
qstat -Q
Name            Users       Groups      MinTime     MaxTime     MaxRunning  MaxQueued   MaxUserNodes  MaxNodeHours
cache-quad      None        None        None        06:00:00    10          20          None          None
flat-quad       None        None        None        06:00:00    10          20          None          None
```

  - On Theta, max queued is set to 20

j1

j2

j3

multiple
independent
short
jobs

Argonne ▲
NATIONAL LABORATORY

ENSEMBLE JOBS ON THETA

```
for x in {0..100}
do
    # background one aprun call per job
    aprun -n 1 -N 1 myapp.py --args $x &
    sleep 1
done
wait
```

Note:

▪ There is a system limitation of 1000 simultaneous apruns per Cobalt script job

▪ You should include a sleep call between aprun invocations

For more details of running ensemble jobs on Theta, see Paul Rich's slides
https://www.alcf.anl.gov/files/rich-ensemble-jobs-2017_1.pdf

WORKFLOW ACROSS ALCF SYSTEMS

▪ How can I set job dependencies across multiple resources (e.g. Mira and Cooley)?
  – ALCF does not provide a means of doing this currently
  – However, Cetus and Cooley mount the Mira filesystems. Analysis jobs run on Cooley without needing to transfer the data.

Cetus	Mira	Cooley

/gpfs/mira-fs0, /gpfs/mira-fs1

Argonne
NATIONAL LABORATORY

SOFTWARE FOR MANAGING WORKFLOWS

This manual effort can be ameliorated by scripting your job submission so that new jobs are automatically submitted as you drop below the max_queued limit. Multiple toolkits are available for this purpose.

Balsam (https://www.alcf.anl.gov/balsam)
Balsam is an ALCF project for managing jobs and workflows on our systems. Balsam allows users to define campaigns of many jobs with interdependencies and control how they execute on the systems, either by submitting them to the job scheduler, or running them within a single job for higher throughput.

Swift (http://swift-lang.org)
Swift provides a C-like language for expressing workflows that consist of command-line invocations, and an engine for managing their execution. For a recent Python implementation of Swift concepts, also see Parsl (http://parsl-project.org).

SOFTWARE FOR MANAGING WORKFLOWS

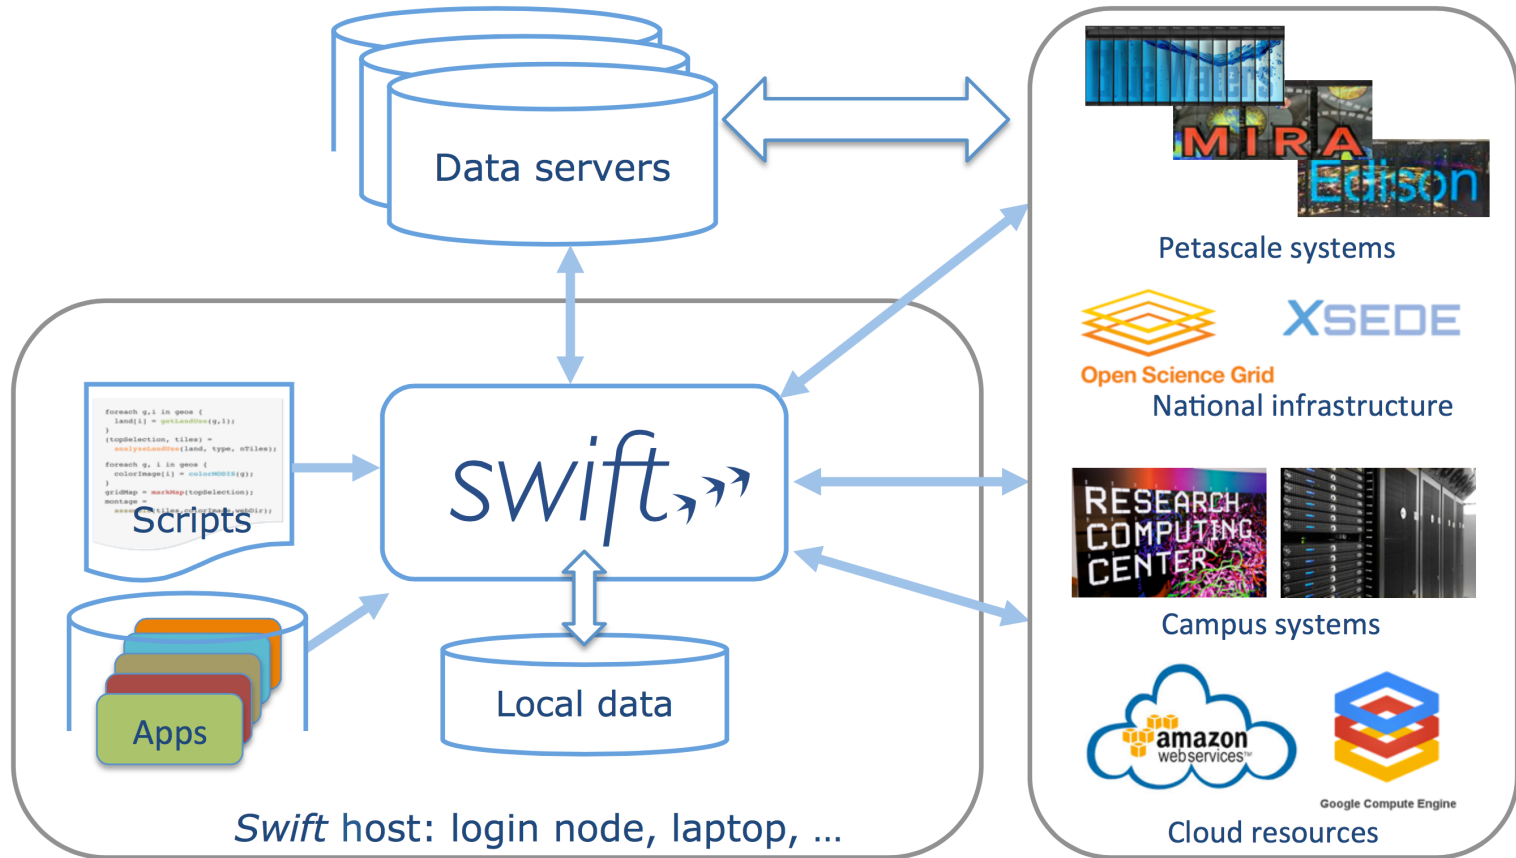**Fireworks (https://pythonhosted.org/FireWorks)**
FireWorks is a free, open-source code for defining, managing, and executing workflows. Complex workflows can be defined using Python, JSON, or YAML, are stored using MongoDB, and can be monitored through a built-in web interface. Workflow execution can be automated over arbitrary computing resources, including those that have a queueing system. FireWorks has been used to run millions of workflows encompassing tens of millions of CPU-hours across diverse application areas and in long-term production projects over the span of multiple years.

Pegasus/Condor (https://pegasus.isi.edu)
Pegasus bridges the scientific domain and the execution environment by automatically mapping high-level workflow descriptions onto distributed resources. It automatically locates the necessary input data and computational resources necessary for workflow execution. Pegasus enables scientists to construct workflows in abstract terms without worrying about the details of the underlying execution environment or the particulars of the low-level specifications required by the middleware

Many others…

Argonne
NATIONAL LABORATORY

Swift enables execution of simulation campaigns across multiple HPC and cloud resources

Data servers

Scripts

swift

Apps

Local data

Swift host: login node, laptop, …

Petascale systems

Open Science Grid   XSEDE

National infrastructure

RESEARCH COMPUTING CENTER

Campus systems

amazon webservices™   Google Compute Engine

Cloud resources

See https://www.alcf.anl.gov/swift

Argonne NATIONAL LABORATORY

When do you need HPC workflow?

Example application: protein-ligand docking for drug screening

O(10) proteins implicated in a disease

X

O(100K) drug candidates

...

= 1M docking tasks...

...then hundreds of detailed MD models to find 10-20 fruitful candidates for wetlab & APS crystallography

Expressing this many task workflow in Swift

For protein docking workflow:

```
foreach p, i in proteins {
    foreach c, j in ligands {
        (structure[i,j], log[i,j]) =
            dock(p, c, minRad, maxRad);
    }
}
scatter_plot = analyze(structure)
```

*To run:*

```
swift -site cooley,blues dock.swift
```

Large-scale applications using Swift

A Simulation of super-cooled glass materials

B Protein and biomolecule structure and interaction

C Climate model analysis and decision making for global food production & supply

D Materials science at the Advanced Photon Source

E Multiscale subsurface flow modeling

F Modeling of power grid for OE applications

All have published science results obtained using Swift

BALSAM

We are developing Balsam to manage workflows on systems at ALCF (and elsewhere). What is Balsam?

- Workflow system for managing large campaigns of interdependent jobs (unlimited queue depth)
- Service for managing flow of jobs into scheduler (automated job submission)
- Flexible launcher for high throughput execution by side-stepping the queue (user control)
- Built on Python and Django
  - Powerful Django object-relational mapper (ORM)
  - Support for many databases (sqlite, MySQL, Postgres)
- Handles data transfers for stage-in and stage-out
- Modular interface to job sources, schedulers, and transfer utilities

Argonne
NATIONAL LABORATORY

BALSAM INSTALLATION ON THETA

```
export PATH=$PATH:$HOME/bin:/opt/intel/python/2017.0.035/intelpython35/bin
conda config --add channels intel
conda create --name balsam_env intelpython3_full python=3
source activate balsam_env
cp /opt/cray/pe/mpt/7.6.0/gni/mpich-intel-abi/16.0/lib/libmpi* ~/.conda/envs/
balsam_env/lib/
export LD_LIBRARY_PATH=~/.conda/envs/balsam_env/lib:$LD_LIBRARY_PATH


git clone git@xgitlab.cels.anl.gov:datascience/balsam.git
cd balsam
pip install -e .
```

Argonne
NATIONAL LABORATORY

BALSAM USE CASES

Add an application to the database

```
balsam app --name myapp --desc "my application" --exec myapp.py
```

Add a job to the database (options mirror qsub/aprun options)

```
balsam job --name myjob --workflow balsam_tutorial --application myapp \
          --wall-min 60 --num-nodes 1 --ranks-per-node 1
```

Launch the Balsam service to automatically flow jobs from the database to the scheduler (subject to user-defined limits)

```
balsam service
```

Submit Balsam launcher jobs manually to run jobs from the database continuously in the context of a single Cobalt job

```
qsub balsam launcher --consume-all --max-ranks 4
```

Argonne ▲
NATIONAL LABORATORY

BALSAM COMMAND LINE INTERFACE (CLI)

```
$ balsam ls
                             job_id |  name |      workflow | application |   state
------------------------------------------------------------------------------------
43e4d1be-fc30-4618-91e0-65b100bc18a8 | myjob | balsam_tutorial | myapp     | CREATED
77d5e32e-8595-4162-ac21-1b86835a6cd0 | myjob | balsam_tutorial | myapp     | CREATED
ebcdc382-d548-4d1d-8c6a-53f3ac59d64e | myjob | balsam_tutorial | myapp     | CREATED
e705d917-b196-4cf4-9b6f-c2ea453e14ae | myjob | balsam_tutorial | myapp     | CREATED
35a570d9-2120-42a8-9984-387285827606 | myjob | balsam_tutorial | myapp     | CREATED
50d8bea1-977a-4028-84ac-022686c7949b | myjob | balsam_tutorial | myapp     | CREATED
eb613637-efc2-49c0-9105-bac8cd48fdd8 | myjob | balsam_tutorial | myapp     | CREATED

$ balsam rm jobs --id 43e4d1be-fc30-4618-91e0-65b100bc18a8

$ balsam --help
  app              add a new application definition
  job               add a new Balsam job
  dep               add a dependency between two existing jobs
  ls                list jobs, applications, or jobs-by-workflow
  modify           alter job or application
  rm               remove jobs or applications from the database
  qsub             add a one-line bash command or script job
  killjob          Kill a job without removing it from the DB
  mkchild          Create a child job of a specified job
  launcher         Start an instance of the balsam launcher
  dbserver         Start/stop database server process
```

21

Argonne
NATIONAL LABORATORY

MORE INFO

▪ Code available at xgitlab.cels.anl.gov/datascience/balsam

▪ Documentation

    ▪ Generated docs exist in repo under balsam/docs

    ▪ Online documentation at balsam.alcf.anl.gov (soon)

▪ Misha Salim will present use of Balsam for hyperparameter optimization on Thursday

HIGH ENERGY PHYSICS EXAMPLE

ATLAS detector at LHC
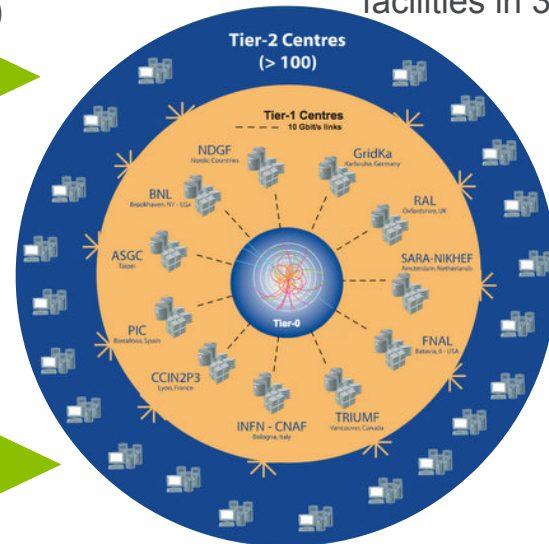

Experimental data (>25PB/year)

LHC Computing Grid (analysis)

>170 computing facilities in 36 countries

ALCF

Simulation data

HIGH ENERGY PHYSICS EXAMPLE

▪ The ATLAS experiment uses more than 1 billion compute hours per year

▪ This consists of integration, event generation, showering, simulation, reconstruction, and analysis

▪ This (very loosely coupled) workflow propagates step-by-step through the ATLAS production system based on requests from users.

▪ Simulation with Geant4 accounts for 60% of ATLAS's computing. We wanted to leverage Mira to offload some of this computing, and chose to target event generation (not simulation) as a first step. This task required coordination of multiple workflow-like steps:
 - fetching job descriptions from ATLAS (manual process)
 - serial phase-space grid integration on local cluster
 - transfer of data from local cluster to ALCF
 - large-scale event generation on Mira, based on incoming integration grids
 - transfer of data from ALCF to local cluster
 - post-processing of output data to exportable format
 - transfer of data from local cluster to WLCG

Argonne
NATIONAL LABORATORY
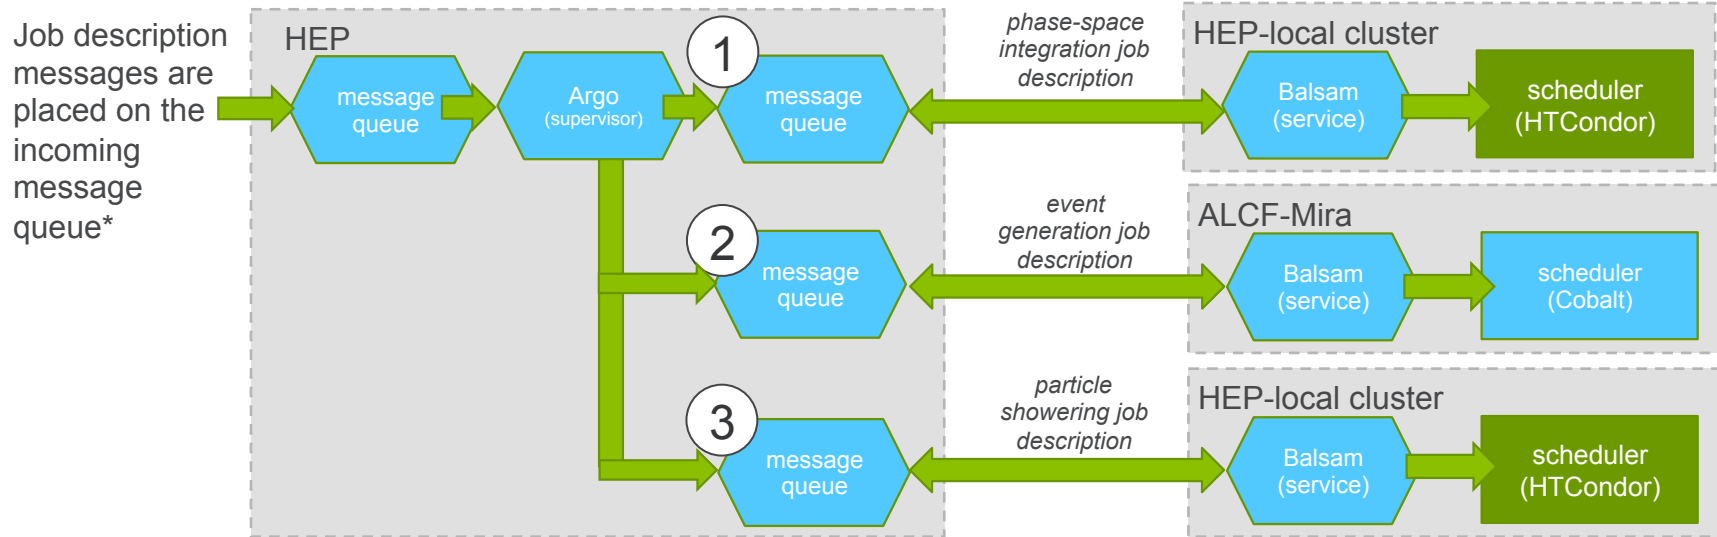
EXAMPLE: HEP EVENT GENERATION

Monte Carlo-based generation of particle collision events such as occur in the ATLAS detector at the Large Hadron Collider, using Alpgen.

Consists of three stages:

1. Generation of phase-space integration grid
2. Generation of weighted events
3. Unweighting of events

MULTI-RESOURCE EVENT GENERATION WORKFLOW

Job description messages are placed on the incoming message queue*

HEP

message queue → Argo (supervisor)

1 message queue ← *phase-space integration job description*

HEP-local cluster
Balsam (service) → scheduler (HTCondor)

2 message queue ← *event generation job description*

ALCF-Mira
Balsam (service) → scheduler (Cobalt)

3 message queue ← *particle showering job description*

HEP-local cluster
Balsam (service) → scheduler (HTCondor)

Final events are transferred manually to ATLAS; this could be automated in future

* Note: To date, all jobs have been injected manually; integration with external services (e.g. PanDA) is incomplete
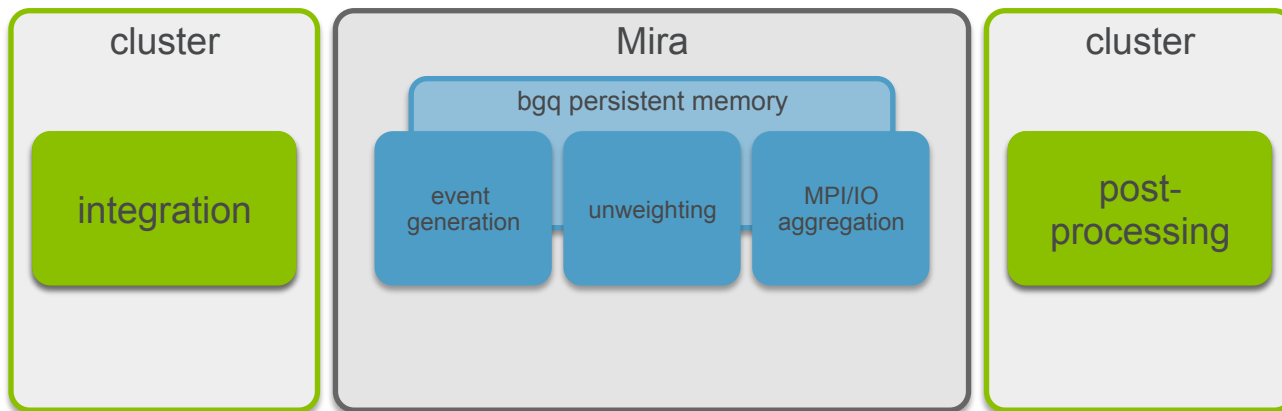
Argonne
NATIONAL LABORATORY

EVENT GENERATION ON MIRA
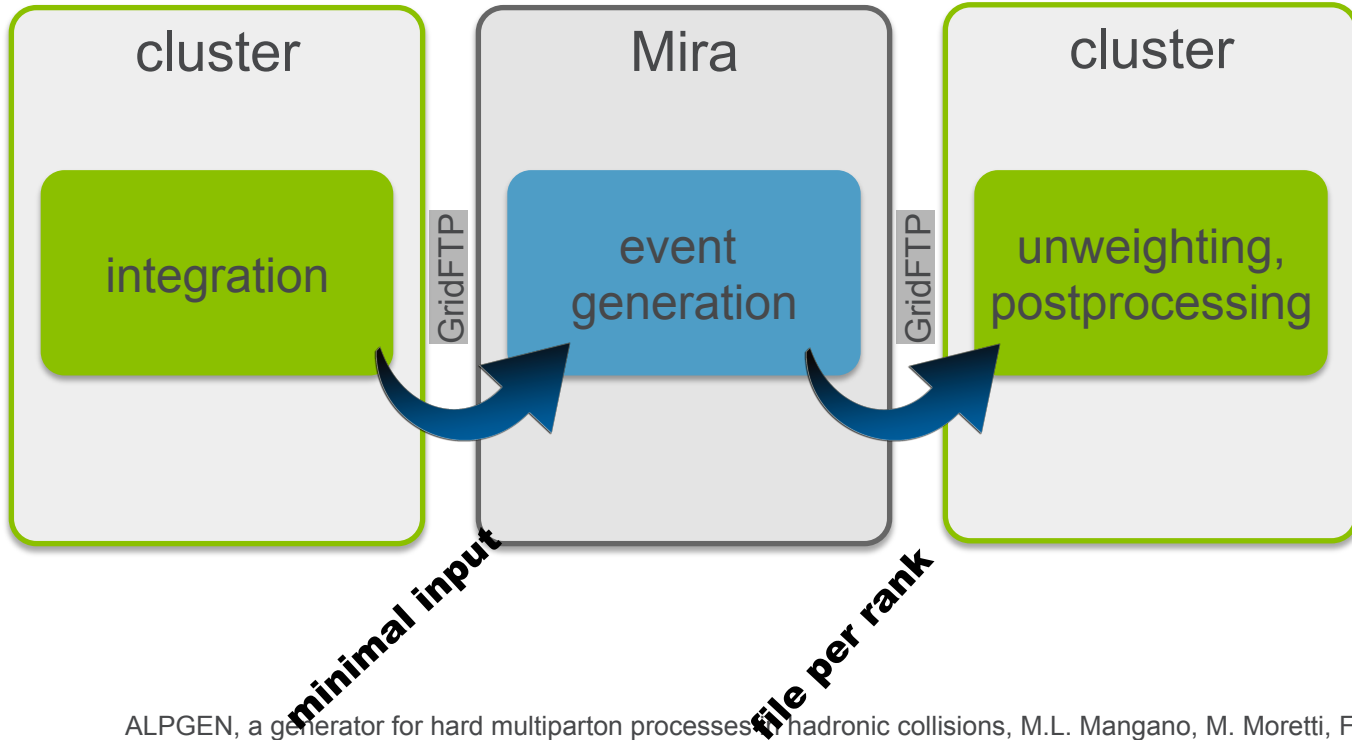
ALPGEN, a generator for hard multiparton processes in hadronic collisions, M.L. Mangano, M. Moretti, F. Piccinini, R. Pittau, A. Polosa, JHEP 0307:001,2003

EVENT GENERATION ON MIRA

- Combine multiple application invocations in single script job
- Use persistent memory for exchanging data between invocations
- Aggregate data from persistent memory to filesystem

For details of container-based workflows for HEP work on Theta, see Taylor Childers' talk on Thursday

RESULTS

- **100M+ hours used for ATLAS event generation through Balsam**
- Unprecedented rates of event generation (largest event generation jobs on Mira requested >1T events)
- More complex/rare events than could be produced on the Grid
- Mira has become the primary site for ATLAS event generation
- 100 million compute hours have been offloaded from the WLCG to Mira, freeing this time to be used for other purposes
- Events that would have taken years to produce on the Grid were generated within a couple months, accelerating the simulation and analysis pipeline and therefore publications

EnergyFEC
Machine: MIRA
Total core hours used per category
2015-07-01 to 2016-06-26

0% <= x < 16.7% of Production Resources
16.7% <= x < 33.3% of Production Resources
33.3% <= x <= 100% of Production Resources

EnergyFEC_3 Machine: THETA

Allocation Core Hours:	40,000,000	Data Dates: 2017-07-01 through 2018-02-25
Usage Core Hours:	36,051,682 (90.1%)	Allocation Dates: 2017-07-01 through 2018-06-30

29

AUTOMATIC BETWEEN-SHOT ANALYSIS OF FUSION EXPERIMENT DATA

DIII-D Tokamak User Facility

Argonne Leadership Computing Facility
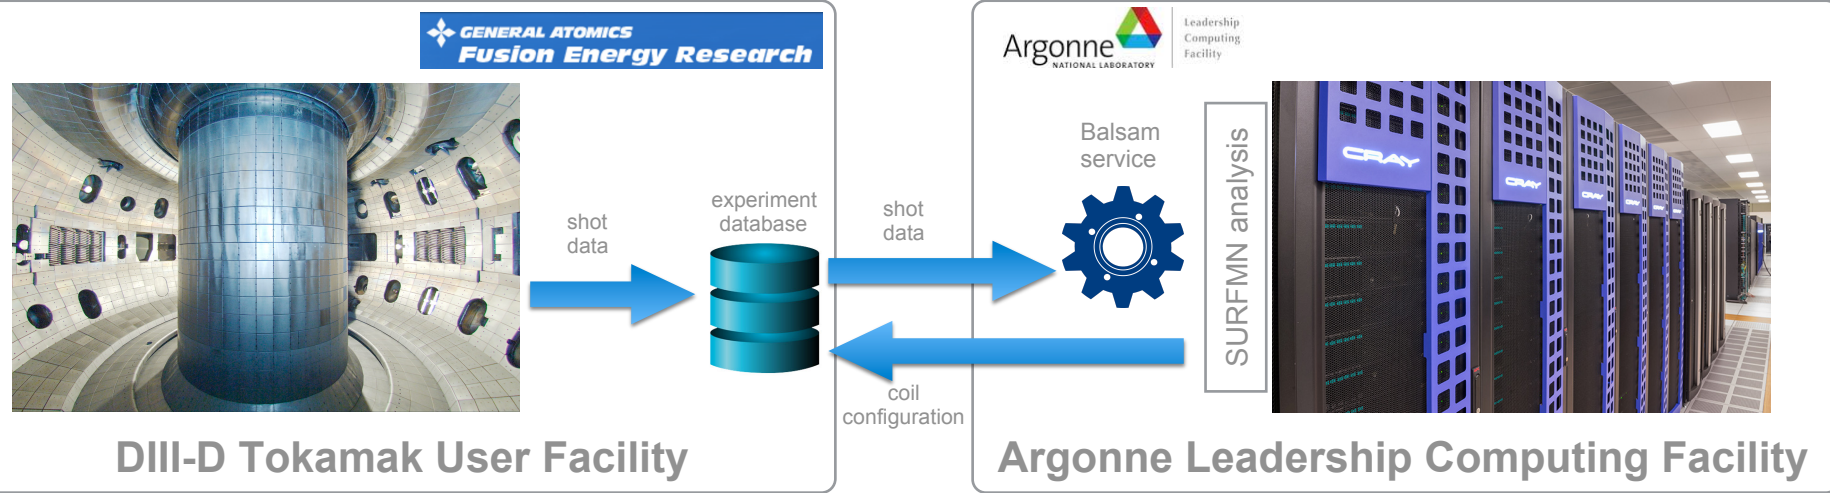
- ‣ Scientists configure experimental "shots" every 15 minutes
  - • A shot is an attempt to magnetically confine high temperature plasma
  - • The timing/current of magnetic coils are configured to control the plasma during a disruption to avoid damage to the containing vessel (applicable to DIII-D and future reactors)
  - • Analyses indicate how to optimize coil configuration for confinement
- ‣ Each shot triggers an automatic, real-time analysis job at ALCF
- ‣ GA scientists integrate analysis results into configuration for next shot
- ‣ Analysis at ALCF enables more complex analyses (128^2 FFT vs 32^2) to be completed faster, improving the accuracy of results and allowing analyses to inform every shot instead of every other

Faster analysis time allows analysis results to be integrated into magnet configuration for subsequent shots. Higher resolution analyses improve configuration accuracy.

SUMMARY

- Workflows are inevitable in computational science

- Small-scale workflow management should be easy and is normally managed using scripting

- Large-scale workflow management can benefit from a workflow system

  - In the transition from small to large workflows, you might build your own workflow system

- The Data Science group at ALCF can help scale your workflows on our systems

We developed a new intersection control paradigm called AIM.

Argonne
NATIONAL LABORATORY