

ParslFest 2021 - Matthew Shaxted

Parsl on Parallel Works for Science and Industry



Parallel Works

PRESENTATION AGENDA

1. What is Parallel Works?
2. Platform Updates & Roadmap
3. Parsl in-the-Wild
 - Tier-1 Auto Supplier
 - Software Vendors
 - University Collaborations
 - NOAA

SECTION ONE

WHAT IS PARALLEL WORKS?

MIDDLEWARE HPC PLATFORM

On-Premise Resource Managers
(Clusters, Desktops)



PBS Works™



kubernetes

Parallel Works Platform
(Saas, On-Prem or Hybrid)

App and Workflow Solutions Library

PIDO

SimuLearner

DOE

Sweep



Workflow Execution Environment

Workflow Development Environment

Universal Cluster Interface



Elastic Cloud and HPC Resources
(Cloud and/or HPC)



SINGLE PANE OF GLASS WORKFLOW EXECUTION

The screenshot displays the Parallel Works dashboard interface. At the top, there is a navigation bar with tabs for COMPUTE, RESOURCES, WORKFLOWS, STORAGE, and ACCOUNT. The main content area is divided into several sections:

- Workflows:** A sidebar on the left lists various workflows such as CANTILEVER_DEMO, CONVERGE_RUNNER, DESIGN_EXPLORER, DOE_GEN, ENERGY_STORAGE_PLANT_D..., GMSH2NEK, GMSH_GEN, GRAPHCONN, HELLO_WORLD, JULIA_RUNNER, and JULIA_RUNNER_NOTEBOOK.
- Run Monitor:** A central section containing a Resource Monitor graph. The graph plots Core-Hr Rate (0 to 1500) over time. Below the graph, summary statistics are shown: 2k / 5k Core-Hrs Used, 0 / 50 GB Stored, and 3k Core-Hrs Remain.
- Computing Resources:** A section listing active resources with their status and capacity. For example, AWS_CONVERGE has 16 cores/active worker, GCP_CONVERGE has 4 cores/active worker, GCP_POOL has 4 cores/active worker, LSF_CLUSTER has 2 cores/active worker, and OCI_POOL_PW has 1 core/active worker. A progress bar for GCP_CONVERGE shows 0 active, 0 requested, and 10 stopped cores, with a total capacity of 40 cores.
- File Explorer:** A sidebar on the right shows a file tree with folders like converge, github, jobs, modules, sample_outputs, storage, and workflows.

SINGLE PANE OF GLASS WORKFLOW EXECUTION

The screenshot displays the Parallel Works web interface. The top navigation bar includes 'COMPUTE', 'RESOURCES', 'WORKFLOWS', 'STORAGE', and 'ACCOUNT'. The left sidebar shows a list of workflows, with 'ENERGY_STORAGE_PLANT_D...' selected. The main content area is titled 'ENERGY_STORAGE_PLANT_DESIGN' and features a 'julia' logo. Below the logo, there are input fields for 'Case Name' (energy_storage_plant_design) and 'storage_a_range' (1:5:1), and a 'storage_b_range' (1:5:1) field. An 'Execute' button is visible. A message states 'Workflow Resource (GCP_POOL) Not Started. Please Start Selected Resource on Main Compute Page.' Below this, the section 'Energy Storage Plant Design' provides a detailed description of the workflow, its purpose, and the design parameters. It also lists the input and output of the workflow. On the right side, a file explorer shows the directory structure, including files like 'main.py', 'parsipw.py', 'path.py', 'post.sh', 'template.inputs', 'workflow.xml', and various sub-directories like 'converge_runner' and 'design_explorer'.

Case Name **storage_a_range**

storage_b_range

/storage

Workflow Resource (GCP_POOL) Not Started. [Please Start Selected Resource on Main Compute Page.](#)

Energy Storage Plant Design

The workflow runs design parameter exploration workflow for energy storage plant. The system under study is a energy storage plant that provides products to the customers by using the equipment units. The annual operation costs for such plants are very high and they critically depend on the design of the plant. Thus, it is of interest to discover the optimal design. However, finding the design parameters are computationally challenging because the annual operation cost can only be evaluated via expensive simulations. As such, we utilize the parallelization capabilities of ParallelWorks to enable efficient design parameter exploration.

The design parameters in the system under study is the size of storage A and storage B (see the figure below). In this workflow, for each candidate size combination, we perform 1-year long simulation of the system to evaluate the average annual operation cost. Each simulation (implemented in Julia Language) takes 30 min-1 hour to run, and more than 100 combinations might be of interest. The accumulated operation cost (computed from the simulation) for each design parameter are compared to find the optimal design parameters. Furthermore, the plant operation variables are recorded and visualized for the qualitative assessment of the design.

INPUT:
The workflow takes input (range of sizes of storage A and B) from workflow.xml.

OUTPUT:
The following files are created in the output directory.
Annual operation cost data as text files
Heat map of annual operation over the explored parameter values
Gif animation of plant simulation

SINGLE PANE OF GLASS WORKFLOW DISSEMINATION

The screenshot displays the Parallel Works web application interface. At the top, there is a navigation bar with the logo and menu items: COMPUTE, RESOURCES, WORKFLOWS, STORAGE, ACCOUNT, ADMIN, and SIGN OUT. Below the navigation bar, a search bar is present with the text "Search for workflow names, tags, users...". The main content area is a grid of workflow cards, each representing a different workflow. The cards are organized into two rows. The first row contains six cards: ASHRAE_PREDICT_V2, ASHRAE_TRAIN, CLUSTERING, COLMENA_DEMO, PARSL_STAGE, and TF_FORECAST. The second row contains six cards: TF_MULTIDEVICE_CO, JULIA_RUNNER_NOTEBOOK, CANTILEVER_DEMO, DESIGN_EXPLORER, DOE_GEN, and GMSH. Each card features a unique icon, a title, a subtitle, a description, and a list of tags. The cards are arranged in a grid with a light gray background and a subtle grid pattern. The left sidebar contains a "Category" dropdown menu and a list of "All Items" and "Analysis Category" options. The "All Items" list includes: Built Environment, Development, Finance, Mapping, Manufacturing, Molecular Dynamics, Rendering, and Workflow Templates. The "Analysis Category" list includes: Design Exploration, Energy Modeling, Fluid Dynamics (CFD), Finite Element Analysis (FEA), Machine Learning, Optimization, Parameter Sweep, Rendering, Solar Radiation, and Tolerance Analysis. The "Popular Tools" list includes: 3DCS, Ansys, CalculiX, Converse, Dakota, DOE2, and FEA SWMM.

Parallel Works

COMPUTE RESOURCES WORKFLOWS STORAGE ACCOUNT

ADMIN SIGN OUT

Category

Total 41 Items

Search for workflow names, tags, users...

All Items

Discipline

- Built Environment
- Development
- Finance
- Mapping
- Manufacturing
- Molecular Dynamics
- Rendering
- Workflow Templates

Analysis Category

- Design Exploration
- Energy Modeling
- Fluid Dynamics (CFD)
- Finite Element Analysis (FEA)
- Machine Learning
- Optimization
- Parameter Sweep
- Rendering
- Solar Radiation
- Tolerance Analysis

Popular Tools

- 3DCS
- Ansys
- CalculiX
- Converse
- Dakota
- DOE2
- FEA SWMM

Workflow Cards:

- ASHRAE_PREDICT_V2** (alvaroidalto) - test - tags: test - Add Parallel Workflow
- ASHRAE_TRAIN** (alvaroidalto) - test - tags: test - Add Parallel Workflow
- CLUSTERING** (alvaroidalto) - Cluster data - tags: ml, cluster - Add Parallel Workflow
- COLMENA_DEMO** (alvaroidalto) - Colmena Demo - tags: colmena, ml - Add Parallel Workflow
- PARSL_STAGE** (alvaroidalto) - Parsl Staging Test - tags: test, staging - Workflow Added
- TF_FORECAST** (alvaroidalto) - Tensorflow multidevice coaster cluster - tags: tensorflow, gpu, ml - Add Parallel Workflow
- TF_MULTIDEVICE_CO** (alvaroidalto) - Tensorflow multidevice coaster cluster - tags: tensorflow, gpu, ml - Add Parallel Workflow
- JULIA_RUNNER_NOTEBOOK** (danielapuchall) - Julia Basic Script Runner - tags: test - Workflow Added
- CANTILEVER_DEMO** (demoworkflows) - Dakota Cantilever Demo - tags: sandia, dakota, fea - Workflow Added
- DESIGN_EXPLORER** (demoworkflows) - Design Explorer Visualization - tags: dev, visualization, graphconn - Workflow Added
- DOE_GEN** (demoworkflows) - Design of Experiments - tags: superlearner, ml, graphconn - Workflow Added
- GMSH** (demoworkflows) - Gmsh Mesh Generator - tags: gmsh, mesh, graphconn - Workflow Added

SINGLE PANE OF GLASS WORKFLOW DEVELOPMENT

The screenshot displays the Parallel Works IDE interface. The top navigation bar includes 'COMPUTE', 'RESOURCES', 'WORKFLOWS', 'STORAGE', and 'ACCOUNT'. The main workspace is divided into a file explorer on the left, a code editor in the center, and a terminal at the bottom.

File Explorer (Left): Shows a project structure with folders like 'converge', 'github', 'jobs', 'modules', 'sample_outputs', 'storage', 'workflows', and 'cantilever_demo'. The 'cantilever_demo' folder is expanded, showing files like 'main.py', 'parslpw.py', 'path.py', 'post.sh', and 'template.inputs'.

Code Editor (Center): Displays the content of 'main.py' with the following code:

```
1 import parsl
2 import os
3 import time,sys
4
5 from parsl.app.app import python_app, bash_app
6 from parsl.data_provider.files import File
7 #from path import Path
8 from parsl.data_provider.pwfiles import Path
9 from parslpw import pwconfig,pwargs
10
11 @bash_app
12 def run_cantilever (inputs=[], outputs=[],stdout='std.out',stderr='std.err'):
13     return '''
14         if=%s
15         cf=%s
16         of=%s
17         python $cf $if > $of
18     ''' % (inputs[0].filepath,inputs[1].filepath,outputs[0].filepath)
19
20 @bash_app
21 def summarize_results (inputs=[], outputs=[],stdout='std.out',stderr='std.err'):
```

Terminal (Bottom): Shows the current shell prompt as 'demoworkflows@go-centos-user-1: /pw\$'.

Status Bar (Bottom Right): Displays 'Ln 1, Col 1 LF UTF-8 Spaces: 4 Python'.

LOCATION INDEPENDENT MULTI-SITE WORKFLOW W/ PARSL

Simple automation scripts enable complex workflows to run across multiple Cloud service providers using multiple applications, with automated data movement.



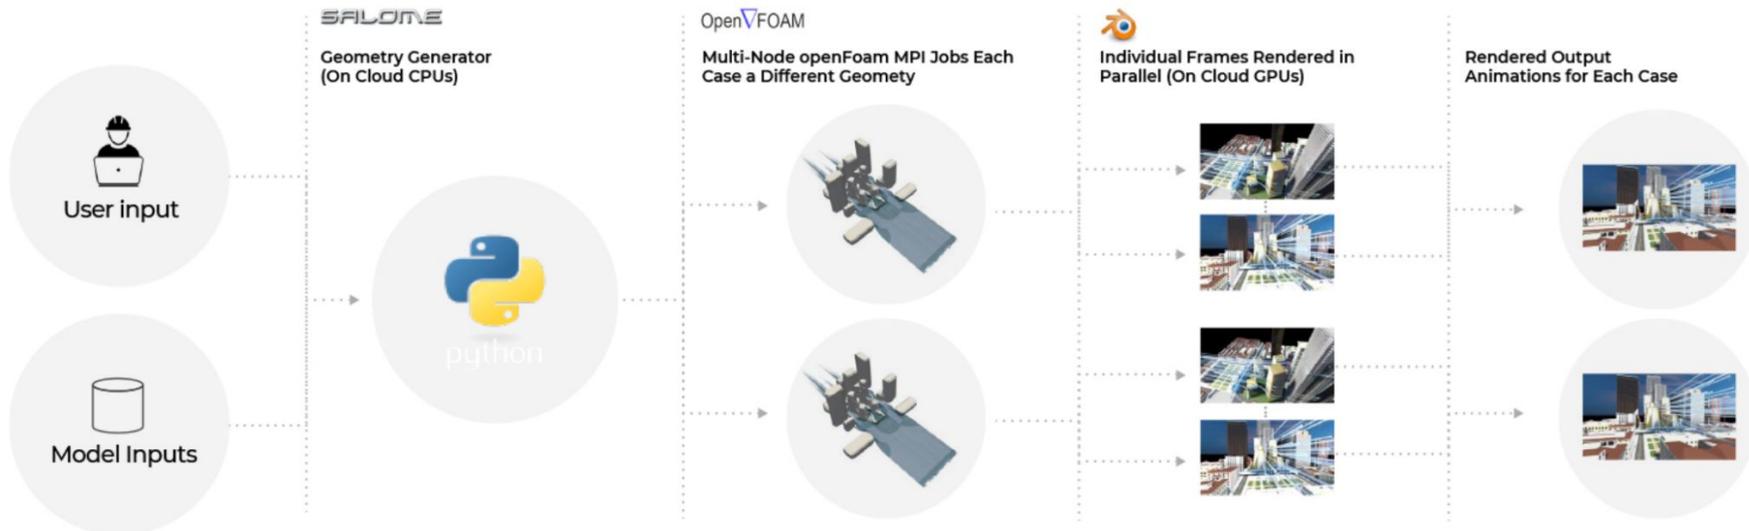
PUBLIC CLOUD



BARE-METAL HPC



PUBLIC CLOUD

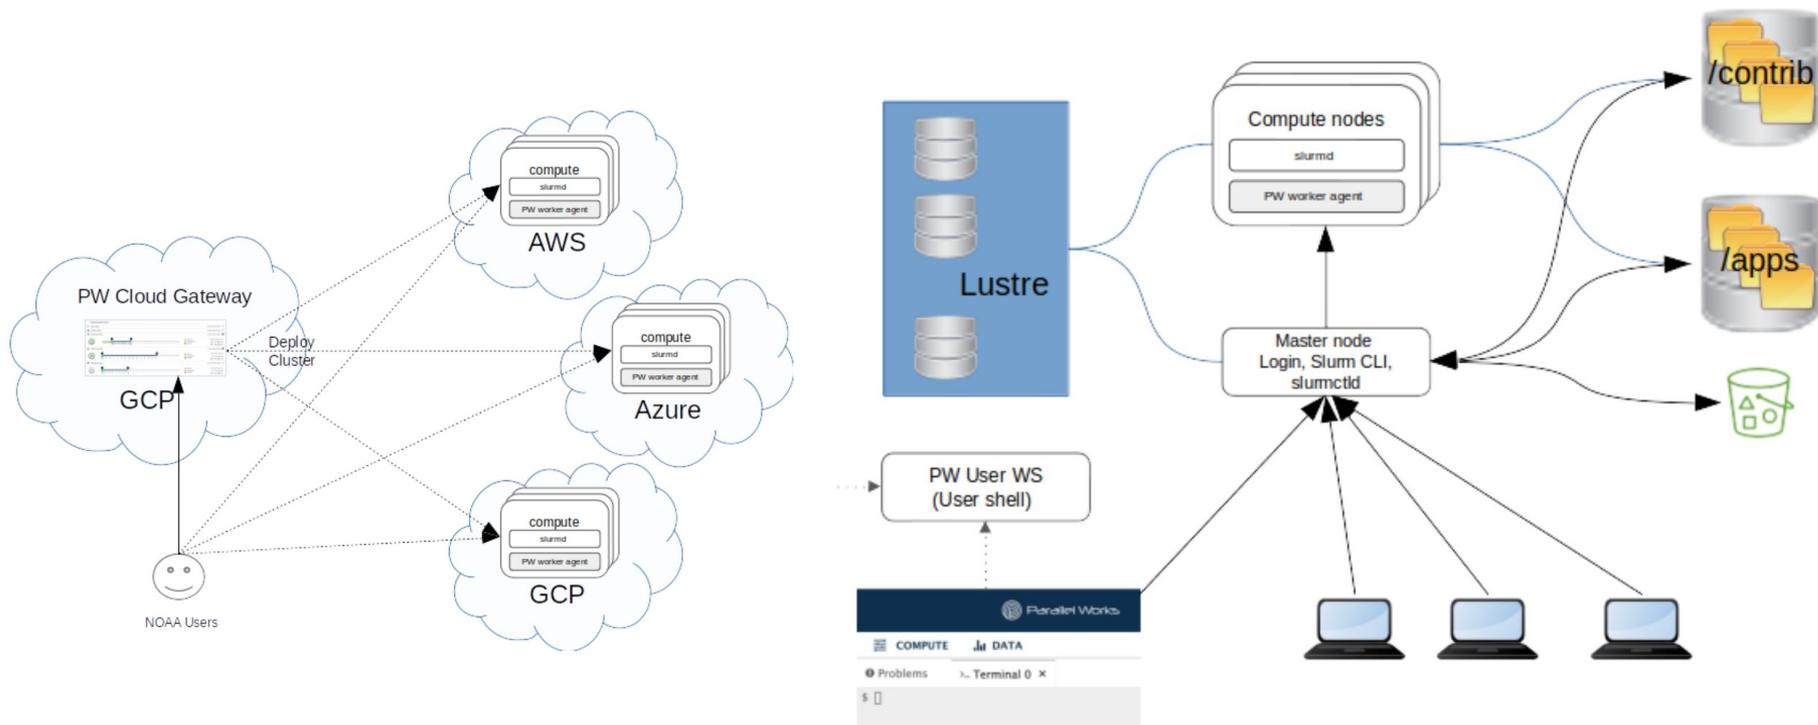


SECTION TWO

PLATFORM UPDATES & ROADMAP

CLUSTERS-IN-THE-CLOUD

Transient fully provisioned slurm/lustre clusters on AWS, GCP, Azure and OCI.



NEW UNIFORM “ELASTIC” PROVIDERS

Elastic providers running over Swift “Coaster” shell-based service - custom provider integrated into Parsl.

Elastic Resource Providers:



Google Compute



Amazon Web Services



Microsoft Azure



Oracle Cloud



VMware vSphere



R-Systems



Penguin POD



Slurm Cluster



PBS Cluster



IBM LSF Cluster



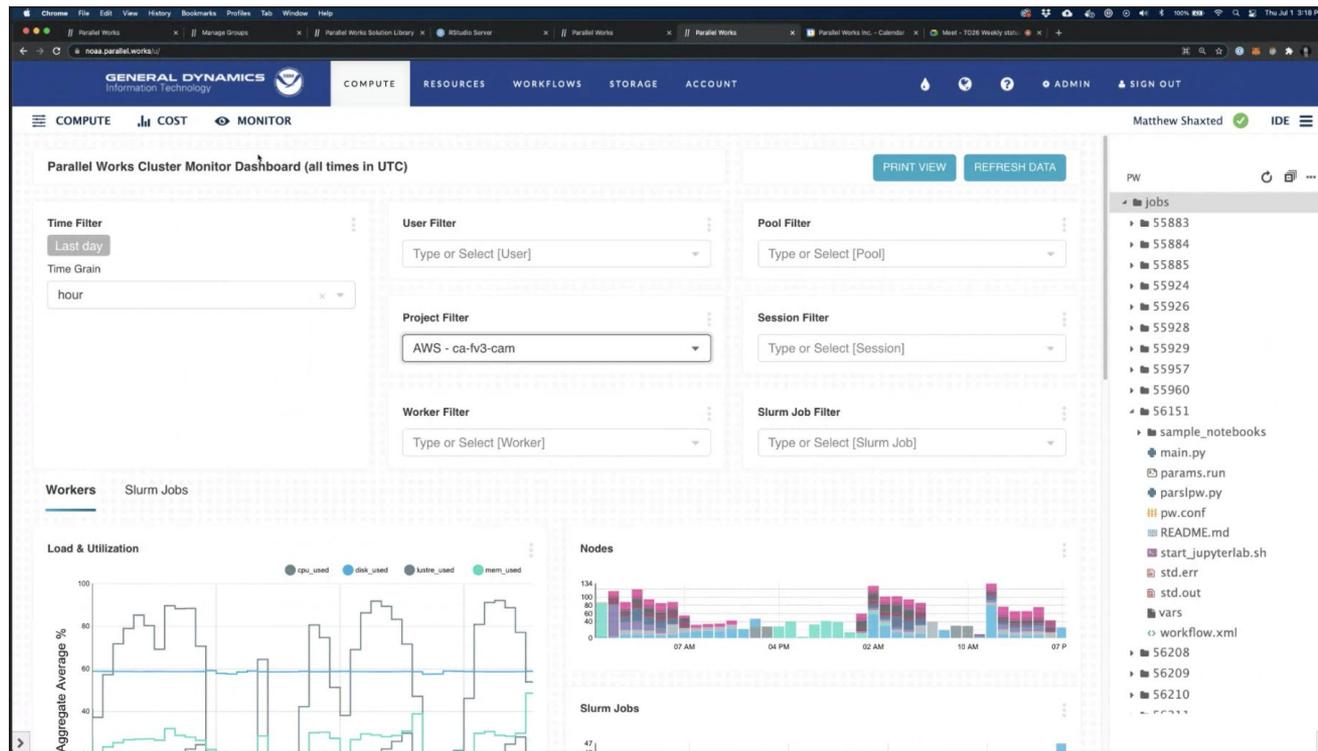
Smart Pool



Passive

CLOUD COST & ALLOCATION ENFORCEMENT

Uniform monitoring capability for tracking all costs and node health information across all resource pools and enforcement of project allocations to Parallel Works users.



TECHNOLOGY ROADMAP

- Policy-Driven Task Placement (e.g. Smart Pools)
- Moving Platform to HA with Kubernetes
- Uniform cluster-in-the-cloud provider w/ terraform
- Exploring use of native parsl for compute providers
- FuncX native integration to the platform
- Supporting Additional Workflow Fabrics (e.g. Dask, Jetstream, Prefect.io, Rocoto)

SECTION THREE

PARSL IN-THE-WILD

TIER-1 AUTOMOTIVE SUPPLIER

Multi-stage parsl workflow with conditional inputs running a combinatorial sweep of cases on primarily cloud resources connected to on-premise license servers.

Case Name
NewCase

Simulation Type

Input Parameters for Output CSV
turbDia,turbinePhi,turbine_ER,turbineBsr,turbineD2t

Turbine Type
VNT

Steady State
Yes No

Steady State Model [.gtm] Steady State Export Template [.exp]
/gtauto_data/engines/steady_state/... /gtauto_data/engines/steady_state/...

Output Parameters for Output CSV
BMEP[bar],Torque[N-m],Power[kW]

Transient
Yes No

Transient Model [.gtm] Transient Export Template [.exp]
/gtauto_data/engines/transient/1p2... /gtauto_data/engines/transient/trans...

Output Parameters for Output CSV
Tqe Grad @ 1s[Nm/s],P1E @ transient end[bar]

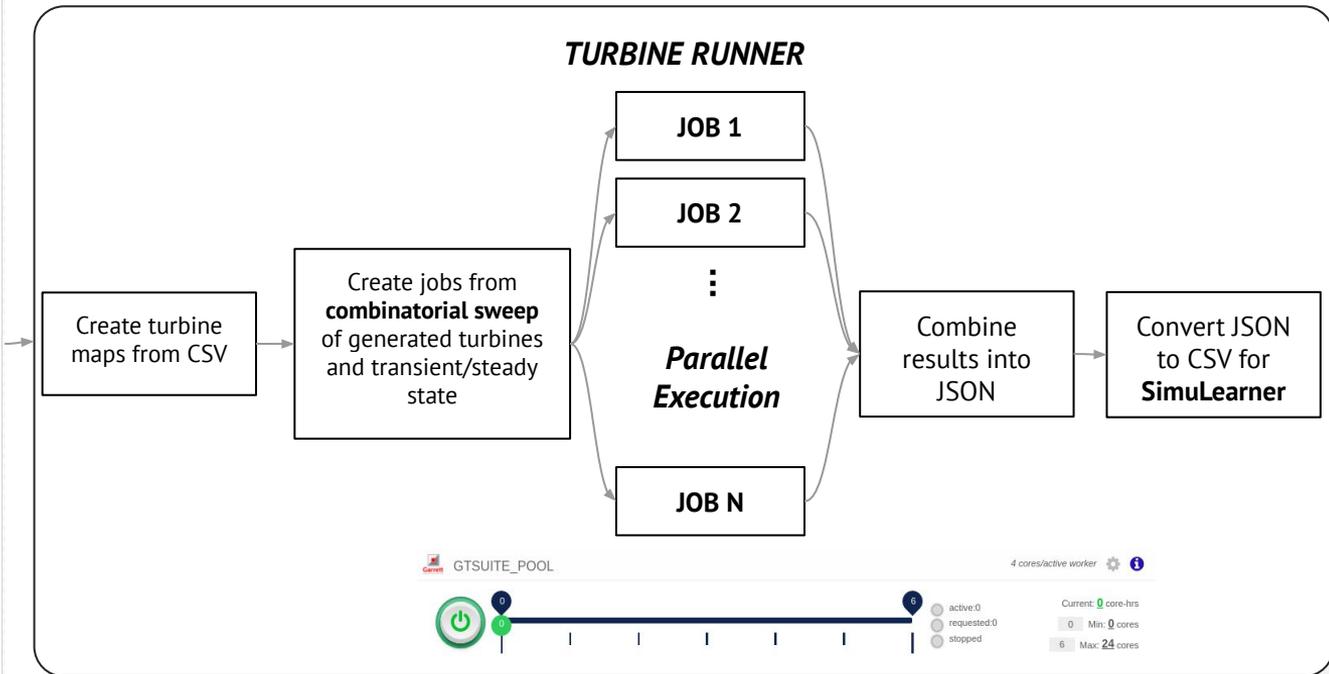
Compressor Map

Compressor Map(s) [.comp]
/gtauto_data/compressors

Turbine Definition

Turbine CSV
/sample_data/turbine_vnt_short.csv (759)

✓ Execute /sample_data



Parallel Works Account Owner

Parallel Works COMPUTE RESOURCES WORKFLOWS STORAGE ACCOUNT MARKET HELP SIGN OUT

HOME

Workflows

ANSYS ANSYS_RUNNER
Ansys Runner Workflow

DOE_GEN
Design Of Experiment Generator

TIVE_OPTIMIZER
GTAutomation Workflow

V_RUNNER_TWIN
GTAutomation Workflow

P_RUNNER_TWIN
GTAutomation Workflow

NUMECA_GEN
Numeca Case Generator

NUMECA_GENRUN
Numeca Case Generator And Runner

NUMECA_RUN
Numeca Case Runner

SL_FIT
Fit Superlearner ML Model

SL_GA_OPTIMIZE
Surrogate Optimization With SL Mo...

SL_PREDICT
Predict With Superlearner ML Model

SL_SCORE
Score Superlearner ML Model

IAP_GEN
ITAutomation Workflow

Run Monitor

Resource Monitor

Core-Hr Rate

7k / 30.5k Core-Hrs Used

206 / 50 GB Stored

23.5k Core-Hrs Remain

GOOGLE_POOL GTSUITE_POOL NUMECA_POOL POD_POOL

Computing Resources

GOOGLE_POOL 4 cores/active worker

GTSUITE_POOL 4 cores/active worker

NUMECA_POOL 40 cores/active worker

POD_POOL 200 cores/active worker

PW

- Ashraf
- gtauto_data
- jan
- jobs
- modules
- Numeca2
- sample_data
- storage
- Anslys2_success
- Anslys3_success
- Anslys1
- Anslys2
- Ashraf
- GSTTwin
- Numeca1
- PROJ_20190920_192226
- test2
- TWR_test
 - 1.txt
 - 2.txt
 - GST19.zip
 - GVT GTP.ipynb
 - GVT.ipynb
 - objective_function.py
 - PROJ_20190920_192226.csv
 - PROJ_20190920_192226.jsor
 - Selected data.csv
 - turb-JanTest1-20200117-10
 - turb-NewCase-20200214-22
 - tw01opt.wbpz
 - tw01opt.wbpz.159E92Ac
- GCT

Account containing multiple private workflows, S3 data stores, and various computing resource pools.

Parallel Works Account Owner | File Manager | Globus

Parallel Works COMPUTE RESOURCES WORKFLOWS STORAGE ACCOUNT MARKET HELP SIGN OUT

HOME

Workflows

search

- ANSYS ANSYS_RUNNER Ansys Runner Workflow
- DOE_GEN Design Of Experiment Generator
- TIVE_OPTIMIZER GTAutomation Workflow
- V_RUNNER_TWIN GTAutomation Workflow
- P_RUNNER_TWIN GTAutomation Workflow
- NUMECA_GEN Numeca Case Generator
- NUMECA_GENRUN Numeca Case Generator And Runner
- NUMECA_RUN Numeca Case Runner
- SL_FIT Fit Superlearner ML Model
- SL_GA_OPTIMIZE Surrogate Optimization With SL Mo...
- SL_PREDICT Predict With Superlearner ML Model
- SL_SCORE Score Superlearner ML Model
- IAP_GEN ITAutomation Workflow

GT_CSV_RUNNER_TWIN

Case Name
NewCase

Simulation Type

Turbine Type
TwinScroll

Steady State
Yes No

Steady State Model [.gtm]
/gtauto_data/engines/steady_state/1p2L_Steady_state.gtm (432K)

Steady State Export Template [.exp]
/gtauto_data/engines/steady_state/steady_state_export.exp (4.3K)

Transient
Yes No

Transient Model [.gtm]
/gtauto_data/engines/transient/1p2L_Transient2_1.gtm (486K)

Transient Export Template [.exp]
/storage/GST/Templates/transient_export_V1.exp (2.8K)

Compressor Map

Compressor Map(s) [.comp]
/gtauto_data/compressors/C301(4)T56AR035-EngineSurge-Confidential-190624.comp (3.1K)

Turbine Definition

Turbine CSV
/jobs/51872/turbine.csv (1.6K)

Execute

/Ashraf

GT CSV Runner:

Runs GT cases with different turbines defined in a CSV file. Each case is defined as a row in the *Turbine* CSV file:

```

turbDia turbinePhi turbine_ER turbineBs turbineD2t
30.0
0.2
1.8

```

Demo Account IDE

PW

- Ashraf
- gtauto_data
- jan
- jobs
- modules
- Numeca2
- sample_data
- storage
- Ansys2_success
- Ansys3_success
- Ansys1
- Ansys2
- Ashraf
- GSTTwin
- Numeca1
- PROJ_20190920_192226
- test2
- TWR_test
- 1.txt
- 2.txt
- GST19.zip
- GVT GTP.ipynb
- GVT.ipynb
- objective_function.py
- PROJ_20190920_192226.csv
- PROJ_20190920_192226.json
- Selected data.csv
- turb-JanTest1-20200117-10
- turb-NewCase-20200214-22
- tw01opt.wbpz
- tw01opt.wbpz.159E92Ac
- GCT

GT Automation tool combines turbine generation with GT Suite compressor / turbine map executions.

Parallel Works Account Owner x +

goofs.parallel.works/u/

Parallel Works COMPUTE RESOURCES WORKFLOWS STORAGE ACCOUNT MARKET HELP SIGN OUT

HOME

Workflows

search

ANSYS ANSYS_RUNNER
Ansys Runner Workflow

DOE_GEN
Design Of Experiment Generator

TIVE_OPTIMIZER
GTAutomation Workflow

V_RUNNER_TWIN
GTAutomation Workflow

P_RUNNER_TWIN
GTAutomation Workflow

NUMECA_GEN
Numeca Case Generator

NUMECA_GENRUN
Numeca Case Generator And Runner

NUMECA_RUN
Numeca Case Runner

SL_FIT
Fit Superlearner ML Model

SL_GA_OPTIMIZE
Surrogate Optimization With SL Mo...

SL_PREDICT
Predict With Superlearner ML Model

SL_SCORE
Score Superlearner ML Model

IAP_GEN
ITAutomation Workflow

__init__.py

Workflow Summary

- Started: 2020-02-24 23:39:43
- Completed: 2020-02-25 00:00:57
- Workflow duration: 0:21:14
- Owner: root
- host: garrettdemo
- rundir: /pw/jobs/52287/runinfo/000
- tasks_failed_count: 0
- tasks_completed_count: 4

[View workflow DAG -- colors grouped by apps](#)

[View workflow DAG -- colors grouped by task states](#)

[View workflow resource usage](#)

App Summary

Name	Count
run_bash_app	4

1w 1m 6m YTD 1y all

Task

Running

Pending

PW

- Ashraf
- gtauto_data
- jan
- jobs
- modules
- Numeca2
- sample_data
- storage
- Ansys2_success
- Ansys3_success
- Ansys1
- Ansys2
- Ashraf
- GSTTwin
- Numeca1
- PROJ_20190920_192226
- test2
- TWR_test
 - 1.txt
 - 2.txt
 - GST19.zip
 - GVT GTP.ipynb
 - GVT.ipynb
 - objective_function.py
 - PROJ_20190920_192226.csv
 - PROJ_20190920_192226.jsor
 - Selected data.csv
 - turb-JanTest1-20200117-10
 - turb-NewCase-20200214-22
 - tw01opt.wbpz
 - tw01opt.wbpz.159E92Ac
- GCT

When a workflow is launched, resources activate elastically and monitor appears to view progress.

Parallel Works Account Owner x +

goofs.parallel.works/u/

Parallel Works COMPUTE RESOURCES WORKFLOWS STORAGE ACCOUNT MARKET HELP SIGN OUT

HOME

Workflows

search

ANYSYS ANSYS_RUNNER Ansys Runner Workflow

DOE_GEN Design Of Experiment Generator

TIVE_OPTIMIZER GTAutomation Workflow

V_RUNNER_TWIN GTAutomation Workflow

P_RUNNER_TWIN GTAutomation Workflow

NUMECA_GEN Numeca Case Generator

NUMECA_GENRUN Numeca Case Generator And Runner

NUMECA_RUN Numeca Case Runner

SL_FIT Fit Superlearner ML Model

SL_GA_OPTIMIZE Surrogate Optimization With SL Mo...

SL_GA_OPTIMIZE_BACKUP Surrogate Optimization With SL Mo...

SL_PREDICT Predict With Superlearner ML Model

SL_SCORE Score Superlearner ML Model

IAP_GEN ITAutomation Workflow

JSON DATA

object > configs > 0 > simulations > 0 > metadata > 0 >

```
object {2}
  metadata [2]
    0 {2}
    1 {2}
  configs [2]
    0 {2}
      metadata [13]
        simulations [1]
          0 {2}
            metadata [2]
              0 {2}
                key : type
                value : Steady State
              1 {2}
            parameters [27]
              0 {2}
                key : Case[No Unit]
                value [2]
                  0 : 1
                  1 : 2
              1 {2}
              2 {2}
              3 {2}
              4 {2}
              5 {2}
              6 {2}
              7 {2}
              8 {2}
              9 {2}
```

Demo Account IDE

PW

- sample_data
 - turbines
 - 46_6DM3_TwScr_SteadyState_V
 - 46_6DM3_TwScr_Transient_V1
 - Ansys3-Success-v19.5-20200
 - Ansys3_2-20200124-113353.
 - Ansys3_test2-20200124-113
 - Ansys3-20200123-195252.tx
 - Ansys3-8process-20200127-(
 - Case1-20200123-204629.txt
 - Case1-20200123-204759.txt
 - data.csv
 - doe-alvaro-test-20200120-17
 - doe-alvaro-test-20200120-17
 - doe-mcs-20200116-022519.t
 - doe-mcs-20200116-022519.j
 - GT-alvaro-20200217-171349
 - GT-alvaro-20200217-173137
 - GT-alvaro-20200217-213943
 - GT-alvaro-20200217-231005
 - GT-alvaro-20200217-235115
 - GT-alvaro-20200221-164522
 - GT-alvaro-monday-20200224
 - GT-alvaro-opt-20200120-194
 - G_0200130-
 - GT-mcstwintest-20200214-15
 - GT-mctwintest-20200214-21.
 - GT-NewCase-20200113-1442
 - GT-NewCase-20200124-1847
 - GT-NewCase-20200124-2026
 - GT-NewCase-20200124-2132
 - inputs_ony.csv
 - numeca-test-20200204-1818

Results post-processed directly into JSON and pushed into private S3 data lake bucket.

Parallel Works Account Owner x +

goofs.parallel.works/u/

Parallel Works COMPUTE RESOURCES WORKFLOWS STORAGE ACCOUNT MARKET HELP SIGN OUT

HOME

Workflows

search

ANSYS ANSYS_RUNNER
Ansys Runner Workflow

DOE_GEN
Design Of Experiment Generator

TIVE_OPTIMIZER
GTAutomation Workflow

V_RUNNER_TWIN
GTAutomation Workflow

P_RUNNER_TWIN
GTAutomation Workflow

NUMECA_GEN
Numeca Case Generator

NUMECA_GENRUN
Numeca Case Generator And Runner

NUMECA_RUN
Numeca Case Runner

SL_FIT
Fit Superlearner ML Model

SL_GA_OPTIMIZE
Surrogate Optimization With SL Mo...

SL_GA_OPTIMIZE_BACKUP
Surrogate Optimization With SL Mo...

SL_PREDICT
Predict With Superlearner ML Model

SL_SCORE
Score Superlearner ML Model

MAP_GEN
GTAutomation Workflow

Run Monitor

Show 10 Rows complete

ID	Workflow	Status	Workspace	Creation Time	Time (min)	View	Stop	Redo
52287	GT_CSV_RUNNER_TWIN	Complete	My Parallel Workspace	11:39pm 2/24/2020	21.4			
52279	GT_MAP_RUNNER_TWIN	Complete	My Parallel Workspace	8:28pm 2/24/2020	21.9			
52274	GT_MAP_RUNNER_TWIN	Complete	My Parallel Workspace	4:07pm 2/24/2020	15.9			
52269	TURB_MAP_GEN	Complete	My Parallel Workspace	3:09pm 2/24/2020	5.3			
52268	GT_MAP_RUNNER_TWIN	Complete	My Parallel Workspace	9:16pm 2/21/2020	14.8			
52266	GT_MAP_RUNNER_TWIN	Complete	My Parallel Workspace	8:47pm 2/21/2020	9.9			
52265	GT_MAP_RUNNER_TWIN	Complete	My Parallel Workspace	8:13pm 2/21/2020	15			
52264	GT_MAP_RUNNER_TWIN	Complete	My Parallel Workspace	4:45pm 2/21/2020	5.5			
52263	GT_MAP_RUNNER_TWIN	Complete	My Parallel Workspace	3:30pm 2/21/2020	10.4			
52262	GT_MAP_RUNNER_TWIN	Complete	My Parallel Workspace	3:14pm 2/21/2020	11.1			

Showing 1 to 10 of 25 entries (filtered from 100 total entries)

< Previous 1 2 3 Next >

Resource Monitor

Computing Resources

GOOGLE_POOL 4 cores/active worker

0 30

active: 0 requested: 0 stopped

Current: 0 core-hrs
0 Min: 0 cores
30 Max: 120 cores

PW

- Ashraf
 - CFD
 - Mesh
 - Volvo26
 - doe-CaseName-20200220-12
 - doe-CaseName-20200220-12
 - doe-CaseName-20200220-12
 - numeca.png
 - T343D78T67VH11p5_VN3D78
 - gtauto_data
 - Jan
 - jobs
 - modules
 - Numeca2
 - sample_data
 - storage
 - workflows

Access or rerun any previous workflow executions.

Parallel.Works Account Owner x File Manager | Globus

app.globus.org/file-manager?destination_id=3e860cf8-3df3-11ea-ab4c-0a7959ea6081&destination_path=%2F~%2F&origin_id=b3913bfe-3df2-11ea-9710-021304b0cca7&origin_path=%2Fstorage%2F

File Manager

Collection: scdemo's PW Globus Access Path: /storage/ Matthew 16" Laptop /~/

select all up one folder refresh list view

NAME	LAST MODIFIED	SIZE	
m_s3_bucket	11/18/2019 02:02pm	-	>
new_data.csv	01/21/2020 07:58am	389 B	>
notebooks	02/24/2020 05:34pm	-	>
parsl_tutorial	01/15/2020 11:40am	-	>
pod_store	11/18/2019 01:49pm	-	>
rsystems_store	01/22/2020 12:43pm	-	>
training_data.csv	02/19/2020 04:19pm	6 KB	>

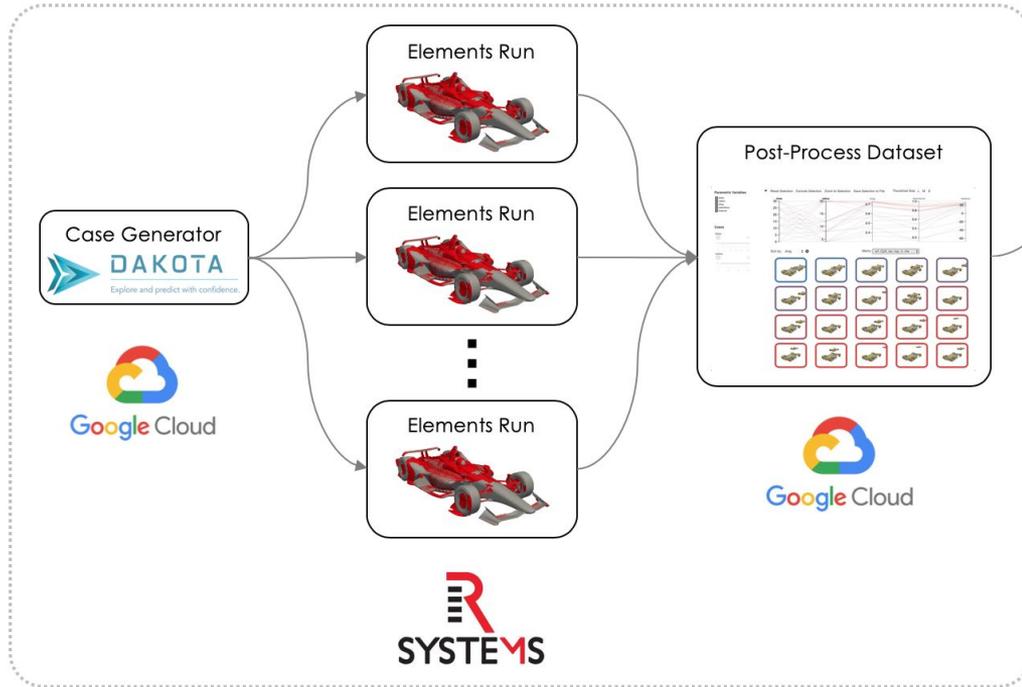
- Permissions
- Transfer or Sync to...
- New Folder
- Rename
- Delete Selected
- Download
- Open
- Upload
- Get Link
- Show Hidden Items
- Manage Activation

Start Transfer & Sync Options Start

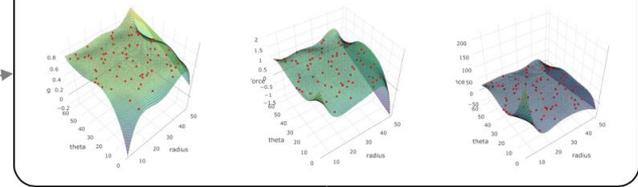
Large datasets and directories synced in Parallel Works user environment with Globus.

AEROKIT DESIGN EXPLORATION

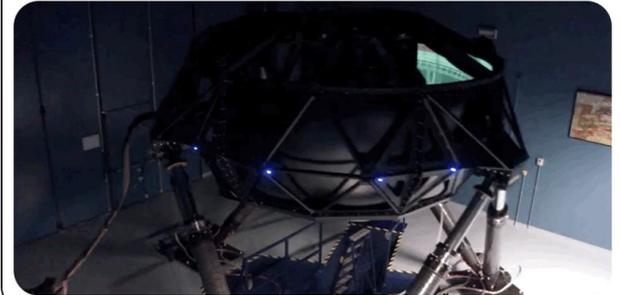
1) Generate Training Data



2) Generate Response Surfaces & Predictive Models



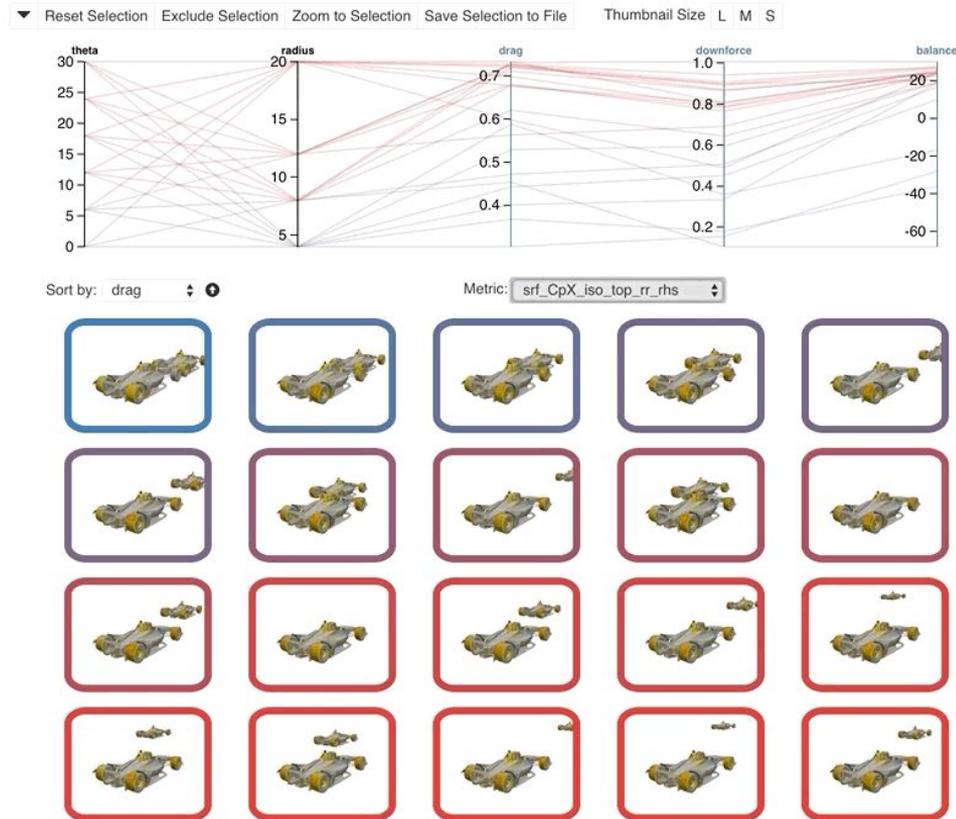
3) Point Predictions to Actuator Movements



Achieving the Vision

1. Generate a parametric model
2. Monte carlo training dataset
3. Generate response surfaces
4. Create predictive model

- Outputs of parameter sweep put into parametric design explorer (right).
- Allows design engineers to parametrically explore run results
- For example, as the trailing car moves closer into the slipstream, drag decreases.



Achieving the Vision

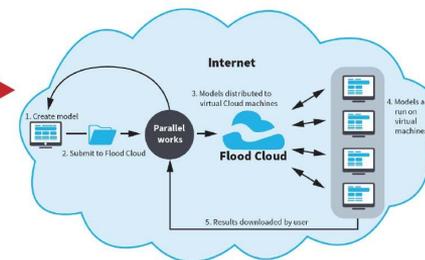
1. Generate a parametric model
 2. Monte carlo training dataset
 3. Generate response surfaces
 4. Create predictive model
- Modular approach - parameter sweep workflows plug into Sandia's Dakota design exploration methods.
 - Complex Dakota routines easy to run via Parallel Works form execution.



The screenshot shows the Dakota configuration interface within the Parallel Works environment. The interface is titled 'AMOSEDALE_UAK18_2CAR_DAKOTA' and includes a 'Scenario Name' field set to 'Scenario1'. It is divided into two main sections: 'Dakota Parameters' and 'Workflow Inputs'. The 'Dakota Parameters' section includes a 'Design Study' dropdown set to 'Design of Experiment', an 'Evaluation Concurrency' slider set to 10, and a 'Number of DOE Experiments' field set to 10. The 'Workflow Inputs' section includes 'FDS case Input Files (.zip, .tar.gz or.tgz)' set to '152: case.tgz', 'Cases file' set to '204: Case1-sweep-np64.csv', 'Number of Processors per Case' set to 64, and 'Number of refinement levels' set to 7. An 'Execute' button is located at the bottom left, and there are icons for refresh, save, and delete at the bottom right.

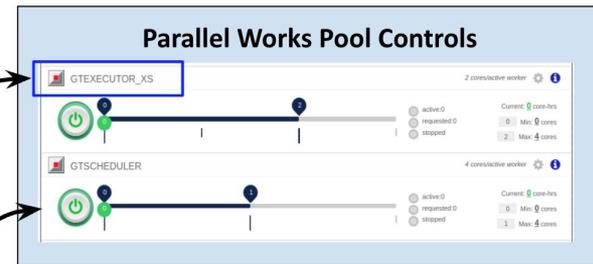
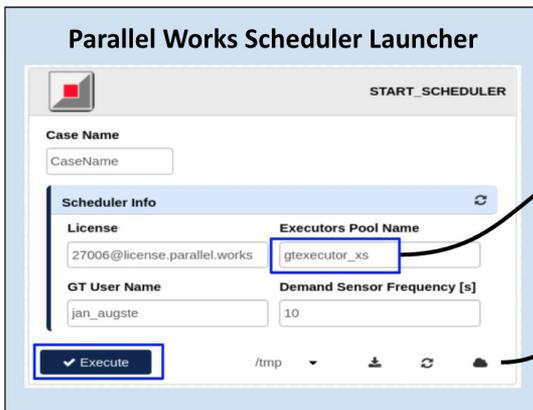
SOFTWARE VENDORS “POWERED BY PARALLEL WORKS”

Software vendors develop Parsl workflows on Parallel Works, and then integrate our REST API into their softwares to get a quick on-ramp to a cloud solution.



JACOBS Flood Cloud
Innovation by Jacobs • Powered by Parallel Works

Flood Cloud, Powered By The Parallel Works API



Use the cloud icon to change / select the scheduler's pool

FLOOD CLOUD WORKFLOW W/ STREAMING & BILLING MANAGER

```
simnum = 1
sims = []
with open("example_upload/simulationlist.csv","r") as f:
    with open("FloodResults.csv","w") as output:
        outfile = csv.writer(output, delimiter=',', quotechar='"', quoting=csv.QUOTE_MINIMAL)
        outfile.writerow(["jobid","task","log","result","fmout","stdout","stderr","solver_versions","misc_metadata"])

    keys = f.readline().rstrip() # read headers
    for line in f.readlines():
        outputs = []
        inputs,output_log,cmd_line,run_dir = line_function(line)
        stream_in,stream_out,stream_csv = get_streams(line, simnum)
        output_dir = get_output_dir(line)
        inputs += stream_in
        outputs += stream_out
        try:
            print(stream_in)
        except:
            print(stream.__repr__())
        #fmlog = "./logs/fm/{}.log".format(simnum) Moved fm log to original .lfi log file
        stdout = "./logs/out/{}.out".format(simnum)
        stderr = "./logs/err/{}.err".format(simnum)
        results = "./results/sim{}.zip".format(simnum)
        output_zip = File(cwd + "/results/./sim{}.zip".format(simnum))
        output_out = File(cwd + "/./" + stdout)
        output_err = File(cwd + "/./" + stderr)
        outputs.append(output_zip)
        outputs.append(output_out)
        outputs.append(output_err)
        #outputs.append(File(cwd + "/example_upload/./" + output_log))

        outfile.writerow([jobid,simnum,stream_csv,results,output_log,stdout,stderr,solver_versions,misc_metadata]) #add jobid and streamed file path

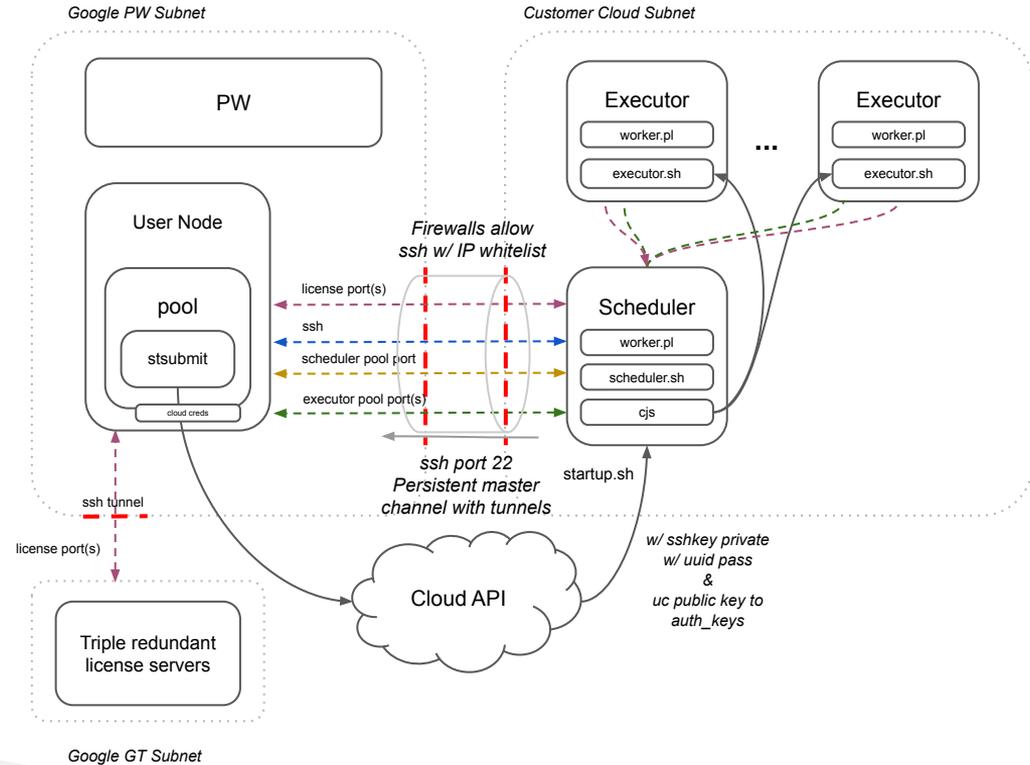
        script = gen_run_script(command_line=cmd_line,run_dir=run_dir,out_zip="sim{}.zip".format(simnum),solver_versions=solver_versions,output_dir=output_dir, log_file=output_log, simnum=simnum)
        x = app(script, inputs=inputs, outputs=outputs, stdout=stdout,stderr=stderr)
        sims.append(x)
        simnum +=1

results = [sim.result() for sim in sims]
#subprocess.run(["zip",""])

print(pwargs)
print("Flood Modeller Task Execution Complete")

PWlogger = get_PW_logger(filename="billing.log")
set_file_logger(filename="billing.log")
set_file_logger(filename="billing.log", name="interchange")
..
```

GT SUITE CLOUD EXECUTION ENGINE



GT Customers submit jobs to the GT scheduler.

The scheduler senses the load and starts executor workers with `cog-job-submit (cjs)`.

The executor script is called on the spawned executor nodes, connecting them to the Scheduler and allowing them to checkout the proper GT licenses (by accessing the license server ports on the Scheduler node).

Executors shutdown automatically when not in use. When scheduler is shutdown, all tunnels are removed.

UNIVERSITY COLLABORATIONS

LOAD_FLEXIBILITY_IN_POWER_GRID

Case Name **Flexibility Range**

Model File (.jl)

Power System Data File (.jl)

/storage

Workflow Resource (GCP_POOL) Not Started. [Please Start Selected Resource on Main Compute Page.](#)

Load Shifting Flexibility in Power Systems

This workflow simulates the effect of spatial load shifting flexibility on the pricing behavior of electricity markets.

Input:

Case Name: Name of the case study

Flexibility Range: A tuple indicating the maximum allowed flexibility level for each instance (in the format of start:step:end)

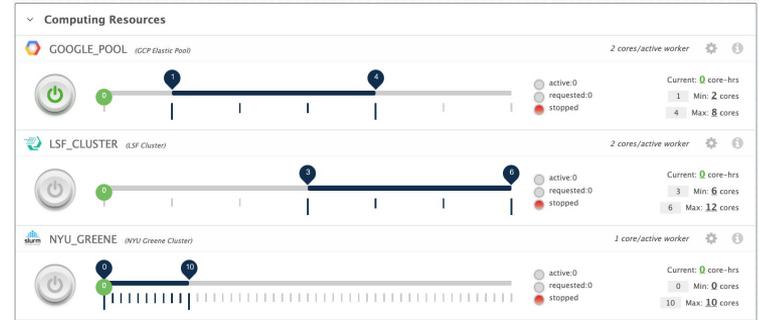
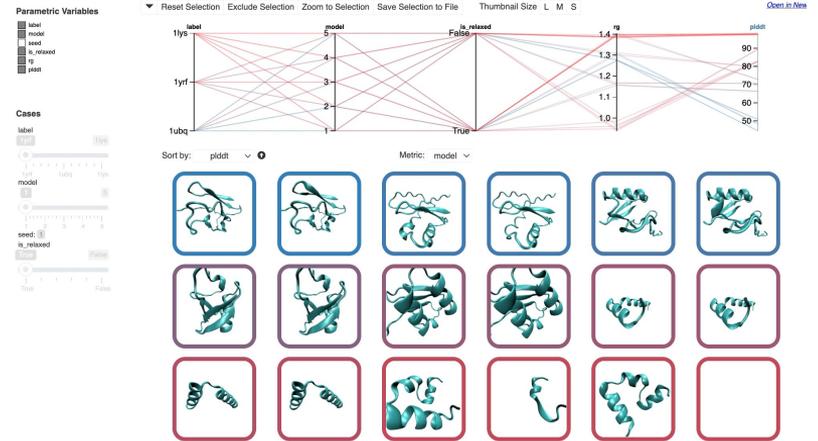
Model File (.jl): A Julia file containing the market clearing optimization formulation for the power system.

Power System Data File (.jl): A Julia file containing data of the power system.

Output:

The price map of each instance with different flexibility levels.

*Images courtesy of Hocky Group at NYU
and Zavala Group at U Wisc.*



NOAA RDHPCS PLATFORM DEPLOYMENT

The screenshot displays the NOAA RDHPCS platform deployment interface. The top navigation bar includes 'COMPUTE', 'RESOURCES', 'WORKFLOWS', 'STORAGE', and 'ACCOUNT'. The user is logged in as 'Matthew Shaxted'.

Workflows:

- OCEAN_PARCELS_DEMO: OceanParcels Lagrangian Trajecto...
- START_JUPYTERLAB: Start JupyterLab in Slurm Cluster
- START_RSERVER: Start R Server

Resource Monitor:

A line graph showing 'Nodes Active' over time. The y-axis ranges from 0 to 80. The x-axis shows various resource pools: AWS_CA_BUDGET_TEST, AWS_TEST_POOL, AZCLUSTER_NOAA_SHARED, AZCLUSTER_NOAA_SINGLE, GCLUSTER_NOAA, and GCLUSTER_NOAA_BUDGET_TEST. The graph shows several spikes in activity, with the highest peak reaching approximately 80 nodes active.

Computing Resources:

Resource Name	Provider	Configuration	Active Cores
AWS_TEST_POOL	AWS	(AWS Test Pool - project-ca-cloudmgmt)	2 cores/active worker
AZCLUSTER_NOAA_SHARED	Azure	(Azure CycleCloud Cluster - project-cz-cloudmgmt)	
AZCLUSTER_NOAA_SINGLE	Azure	(Azure CycleCloud Cluster - project-cz-cloudmgmt)	
GCLUSTER_NOAA	GCP	(GCP Slurm Provider Running NOAA/GDIT GCP Account - project-cg-cloudmgmt)	
GCLUSTER_NOAA_CLOUDMGMT	GCP	(GCP Slurm Provider Running NOAA/GDIT GCP Account - project-cg-cloudmgmt)	
GCLUSTER_NOAA_GPU	GCP	(GCP Slurm Provider Running NOAA/GDIT GCP Account - project-cg-cloudmgmt)	
GCLUSTER_NOAA_SHARED	GCP	(GCP Slurm Provider Running NOAA/GDIT GCP Account - project-cg-cloudmgmt)	
PCLUSTER_NOAA	AWS	(Pcluster Provider Running NOAA/GDIT AWS Account - project-ca-cloudmgmt)	
PCLUSTER_NOAA_SHARED	AWS	(Pcluster Provider Running NOAA/GDIT AWS Account - project-ca-cloudmgmt)	

A power button icon and a progress bar are visible below the resource list. The progress bar shows 145 nodes active, with 0 requested and 0 stopped.

File Explorer:

- PW
 - jobs
 - storage
 - jobs
 - misc
 - my_keys
 - project_keys
 - pw_api_python
 - GCP_IMAGE.txt
 - workflows



Parallel Works

INTERESTED IN TRYING PARALLEL WORKS?

Can get demo accounts created on current system, and/or
put you on a preview list for upcoming release.

SHAXTED@PARALLELWORKS.COM