

Workflow Manager

in Open OnDemand



Harshit Soora · Alan Sussman

University of Maryland



Outline

- 01** | **Motivation**
- 02** | **Project Manager in Open OnDemand**
- 03** | **Workflow Manager in OOD**
- 04** | **Current v4.1 Demo**
- 05** | **Phase II**
- 06** | **Future Directions**

Motivation



Multi-Step Pipelines are Common

- Researchers routinely **Chain tasks**
- Preprocessing → Simulation → Analysis → Visualization
- Manually **Track** job-ids, **Wait** to finish and **Resubmit** by hand



Existing Tools Require Scripting

- Existing workflow like Peagsus, Airflow, Nextflow, Drona **demand learning** a DSL or programming language
- This creates a barrier for domain scientists who are experts in their field, not in software engineering

Motivation



Multi-Step Pipelines are Common

- Researchers routinely **Chain tasks**
- Preprocessing → Simulation → Analysis → Visualization
- Manually **Track** job-ids, **Wait** to finish and **Resubmit** by hand



Existing Tools Require Scripting

- Existing workflow like Peagsus, Airflow, Nextflow, Drona **demand learning** a DSL or programming language
- This creates a barrier for domain scientists who are experts in their field, not in software engineering



No Integrated GUI for HPC Workflow

- Visual tools like Galaxy focus on Bioinformatics
- KNIME targets business BI
- A general-purpose, scheduler-agnostic, browser-based tool, **natively integrated** into an HPC portal



OOD Already Has the Foundation

- Open OnDemand's Project Manager already provides Launchers, Scripts, and Shared directories
- Adding a **visual workflow layer** lets users connect them

Project Manager in Open OnDemand



Launchers

- Contains Script that submit jobs to the scheduler
- Customize scheduler params, resources, and script behavior through environment variables

All sbatch scripts, input data, and output data stored in the project directory that can be freely organizable per user needs

Project Manager in Open OnDemand



Launchers

- Contains Script that submit jobs to the scheduler
- Customize scheduler params, resources, and script behavior through environment variables



Collaboration

- Share projects via shared directories
- UNIX permission-based access control
- Template support for admins

All sbatch scripts, input data, and output data stored in the project directory that can be freely organizable per user needs

Project Manager in Open OnDemand




Launchers

- Contains Script that submit jobs to the scheduler
- Customize scheduler params, resources, and script behavior through environment variables



Collaboration

- Share projects via shared directories
- UNIX permission-based access control
- Template support for admins



Workflows

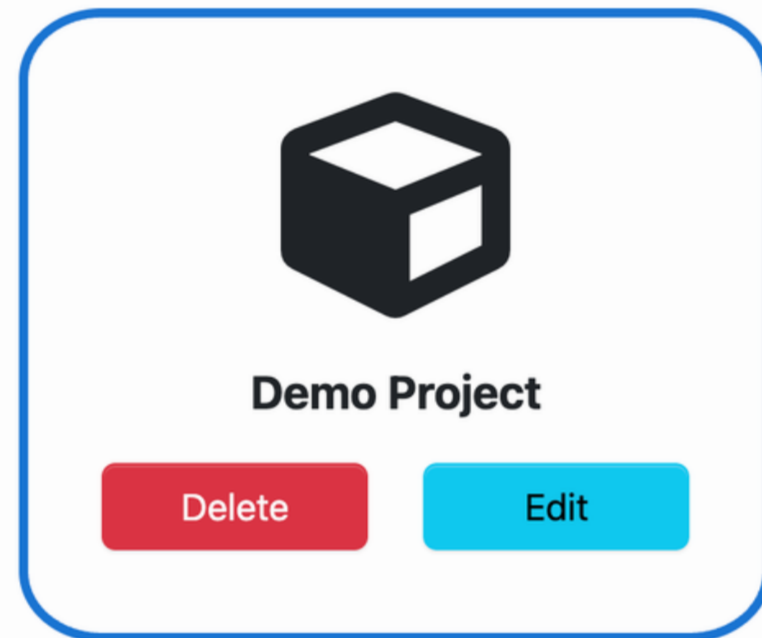
- Connect launchers with **DAG-based** dependency relationships
- Build self-contained systems that can be run with a single click through a visual interface
- **Real-time** status indicator and Error Transparency

All sbatch scripts, input data, and output data stored in the project directory that can be freely organizable per user needs

Project Manager — Interface

OSC OnDemand Files ▾ Jobs ▾ Clusters ▾ Interactive Apps ▾ My Interactive Sessions All Apps </> Develop ▾ Help ▾ Logged in as hsoora Log Out

Project Manager



A card representing a demo project. It features a 3D cube icon at the top, the text "Demo Project" in the middle, and two buttons at the bottom: a red "Delete" button and a cyan "Edit" button.

Create a new project



Create a new project from a
template



Import a shared project



powered by



OnDemand version: 4.1.3

[Software License Notice](#)

Project Manager — Interface

Demo Project

[Back to Projects](#)

Launchers

A

Launch Show

Edit Delete

B

Launch Show

Edit Delete

[New Launcher](#)

Workflows

Demo Workflow

Show

Edit Delete

Active Jobs

44535553 Running 44535554 Queued

Completed Jobs

44535547

Project Directory: *agwzqz3k*

/users/PZS0694/hsoora/ondemand/data/sys/dashboard/projects/agwzqz3k

Name	Size	Date
..		
.ondemand		2026-03-05 09:52:09 AM
slurm-44535553.out	36 B	2026-03-05 09:56:03 AM
A.sh	186 B	2026-02-20 05:22:26 PM
B.sh	186 B	2026-02-20 05:22:26 PM
slurm-44535547.out	83 B	2026-03-05 09:55:17 AM

[Open in files app](#)

What is a DAG?

Graph

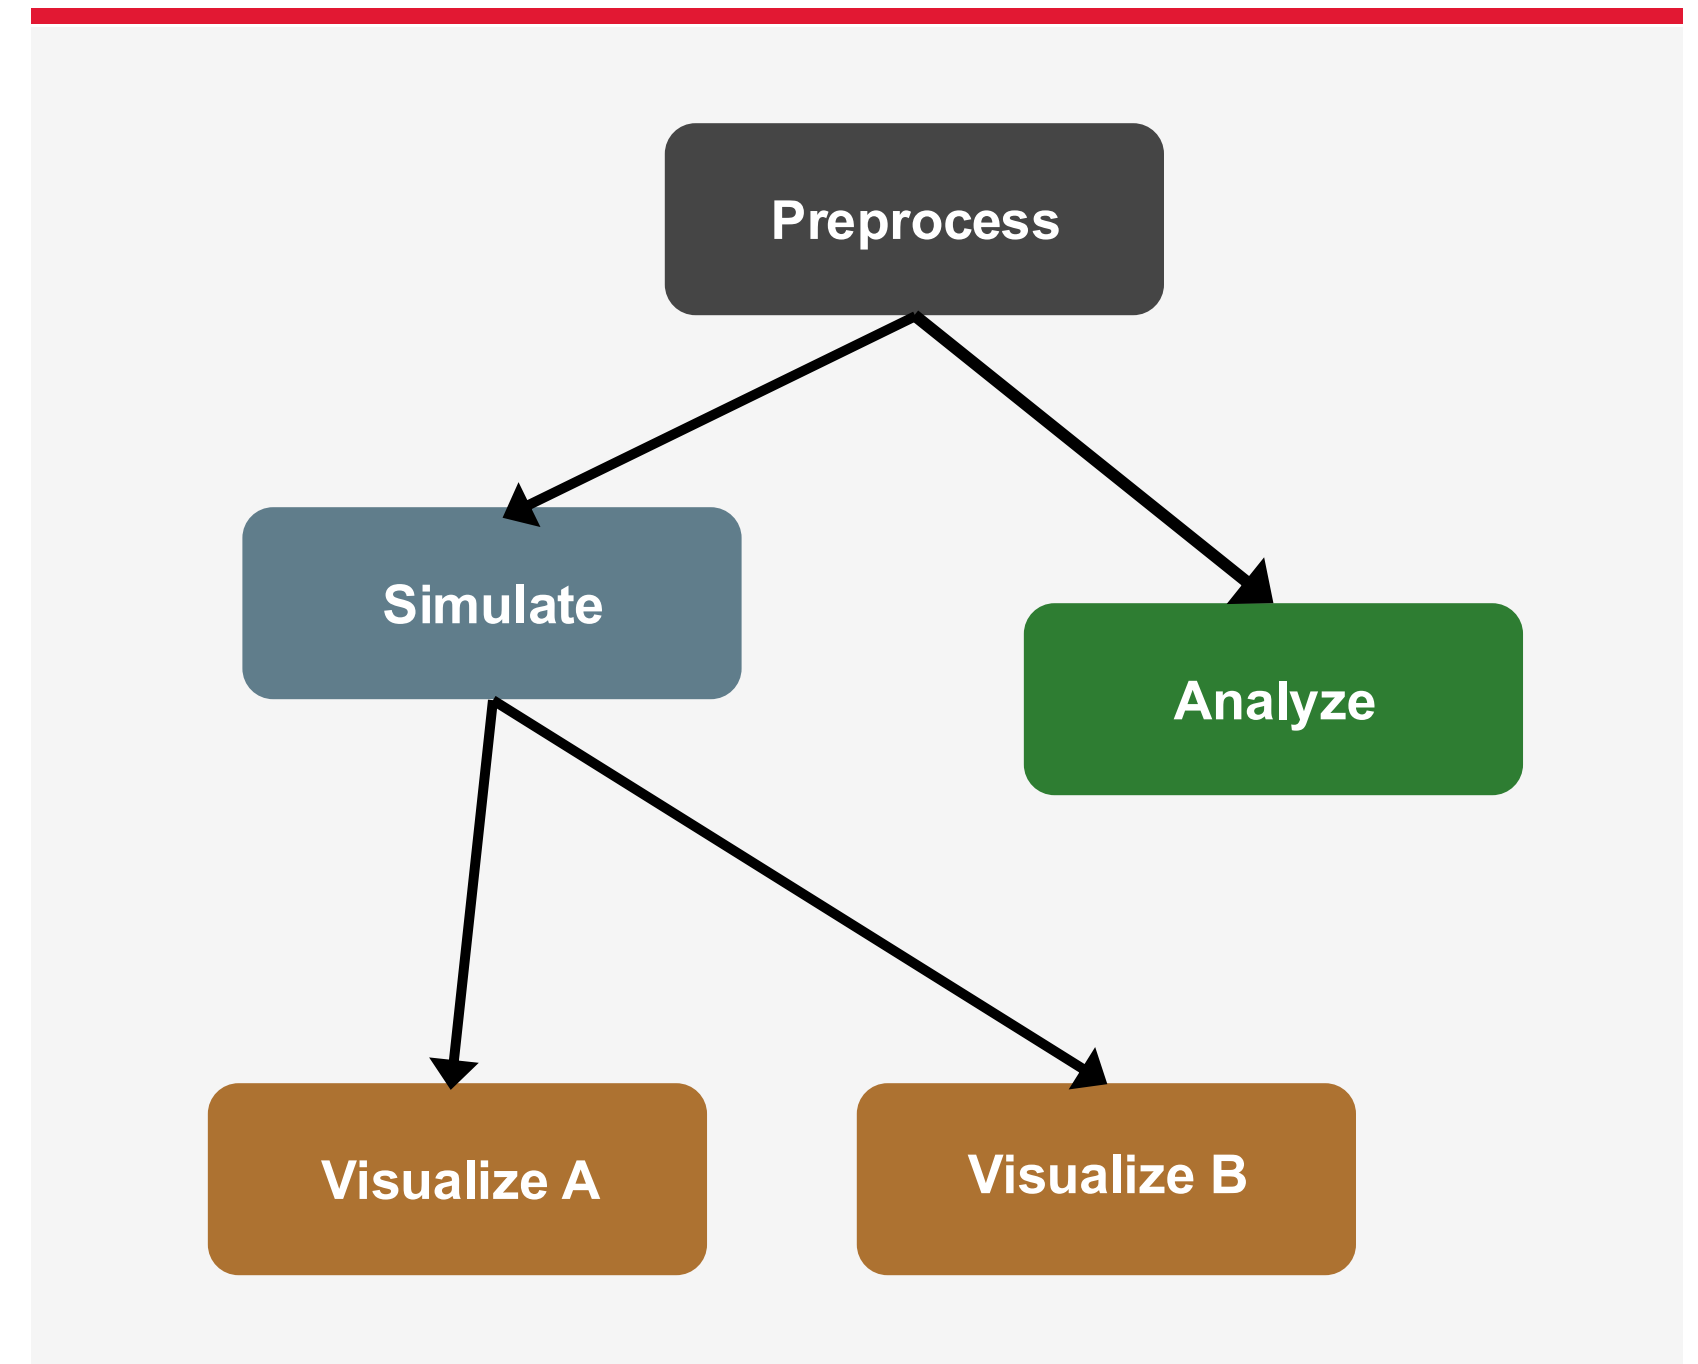
Nodes can be considered as tasks / Launchers
Edges express dependencies used in Scheduler

Directed

Edges have a direction that represent flow of dependency

Acyclic

No cyclic dependencies are allowed



Workflow Manager in OOD



No Scripting Language Required

Unlike previous workflows users build workflows entirely through a graphical drag-and-drop interface.



Scheduler Agnostic by Design

Internally delegates to the HPC scheduler using native job dependency (job-IDs). No meta-scheduler layer, no internal scripting language for translation. Works with SLURM, PBSPro, LSF out of the box.

Workflow Manager in OOD



No Scripting Language Required

Unlike previous workflows users build workflows entirely through a graphical drag-and-drop interface. Zero DSL to learn.



Scheduler Agnostic by Design

Internally delegates to the HPC scheduler using native job dependency (job-IDs). No meta-scheduler layer, no internal scripting language for translation. Works with SLURM, PBS, LSF out of the box.



Visual DAG Builder in the Browser

Interactive SVG canvas with grid snapping, zoom, real-time cycle detection, and live job status polling (30s).



Integrated with Project Manager

Workflows live inside projects alongside launchers, scripts, and data. Collaboration, templates, shared directories, and job history are all built-in. Not a standalone tool.

Workflow v4.1— Demo

[Drag-and-drop DAG builder · Live job status](#)

Edit Launcher

Editing Launcher

Cluster

ascend ▼

Edit

Script

- ✓ A.sh
- B.sh
- C.sh
- D.sh
- E.sh

Account

PZS0694 ▼

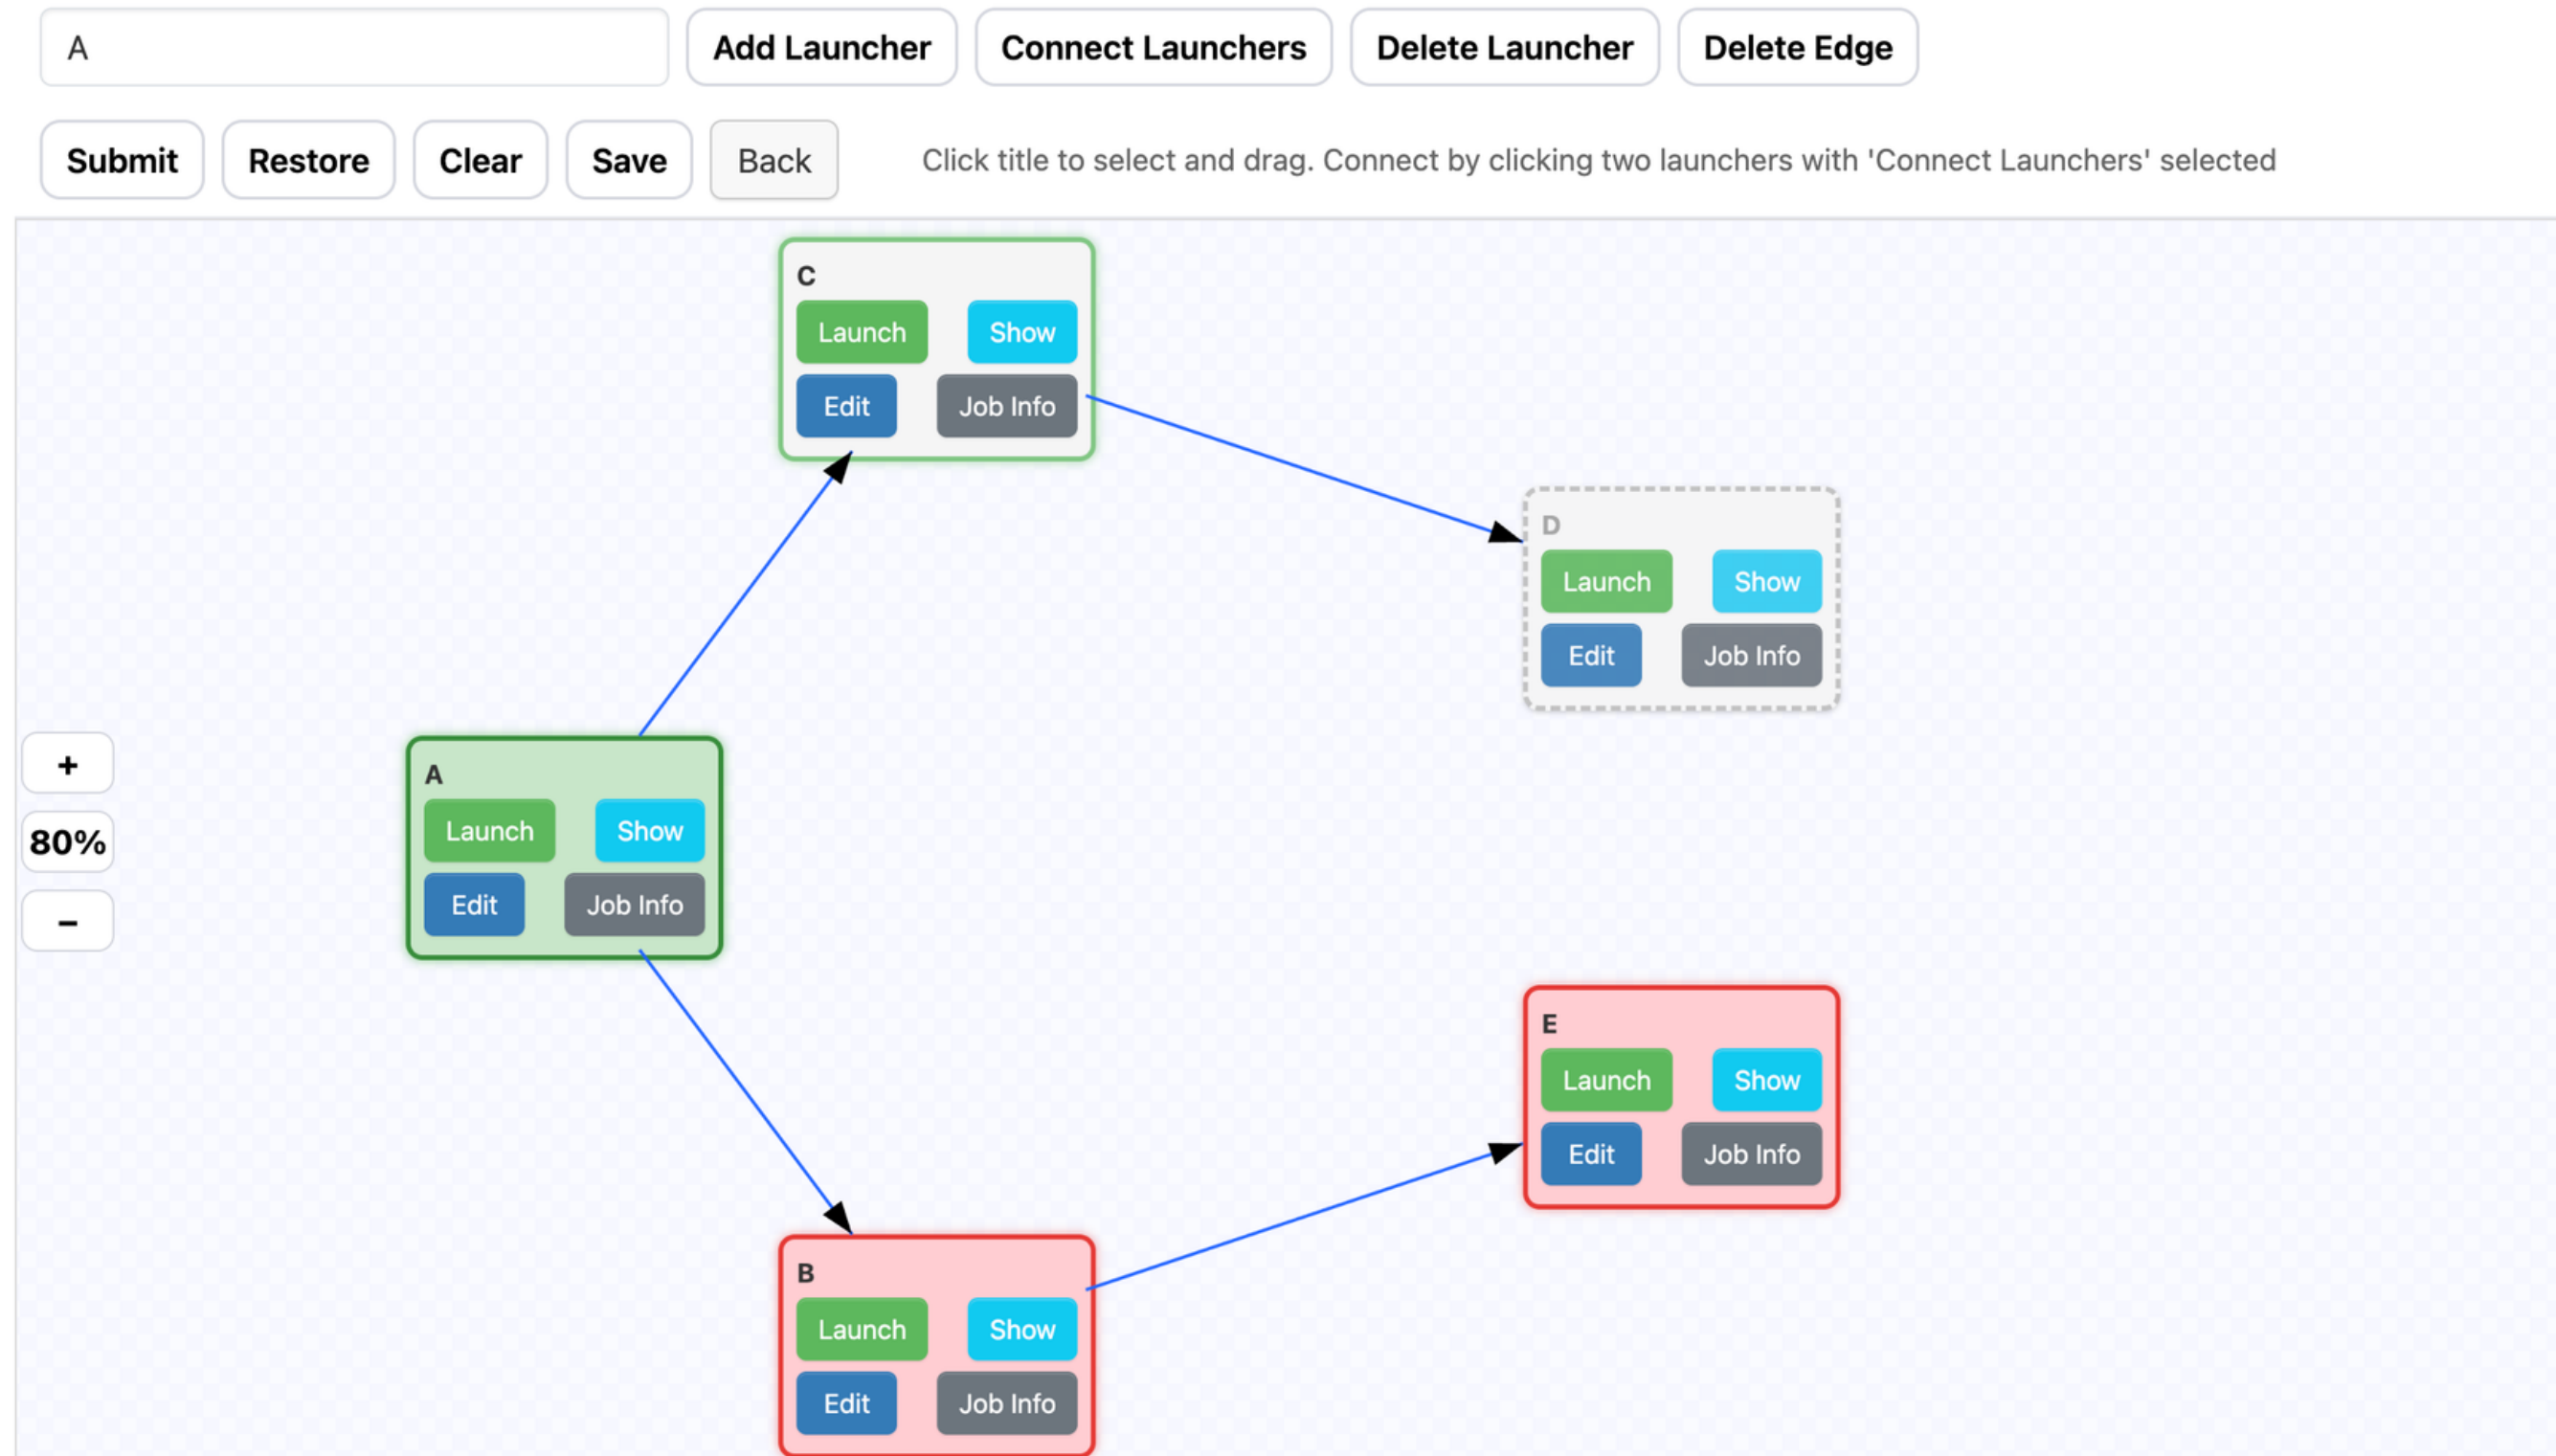
Edit Remove

Add new option

Back

Save

Submit Workflow



Job Info

A

Add Launcher Connect Launchers Delete Launcher Delete Edge

Submit Restore Clear Save Back

Click title to select and drag. Connect by clicking two launchers with 'Connect Launchers' selected

44628367	
Message	Current job state is 'CANCELLED' because of 'DependencyNeverSatisfied'. Job depends on job-id 44628366(failed).
Id	44628367
Status	completed
Job name	sbatch
Job owner	hsoora
Accounting	pzs0694
Procs	1
Queue name	gpubackfill-exp
Wallclock time	less than 1 minute
Wallclock limit	1 hour
Submission	2026-03-09 17:51:17 -0400

80%

Phase II — Motivation

Existing v4.1 Job Order Dependency

Input/Output Ports

- Connect output of one launcher as input to another via edges
- Leverage environment variables to expose output directories
- Dependent launcher should handle the incoming port / environment variable in its sbatch script

Phase II Input/Output Dependency without any internal scripting

Orthogonal Edge Routing

- Convert straight edges to right-angle orthogonal paths
- User-adjustable routing to manage overlaps
- Split edges into angles, snap angles back into straight line

Phase II — Motivation

Existing v4.1 Job Order Dependency

Input/Output Ports

- Connect output of one launcher as input to another via edges
- Leverage environment variables to expose output directories
- Dependent launcher should handle the incoming port / environment variable in its sbatch script

Phase II Input/Output Dependency without any internal scripting

Orthogonal Edge Routing

- Convert straight edges to right-angle orthogonal paths
- User-adjustable routing to manage overlaps
- Split edges into angles, snap angles back into straight line

Visual Dependency Info

- Display output directory name while hovering on an edge
- Port name visible on hover to guide connection decisions

Automatic Script Translation

- Dependency is translated via environment variables into sbatch scripts
- No scripting language needed thus scheduler agnostic design

Phase II — Demo

[Orthogonal edges · Input Output Port connections](#)

Environment Variable as Ports

Script

A.sh

Edit

Cluster

ascend

Edit

Name

output

Environment Variable: output

hsoora/outA

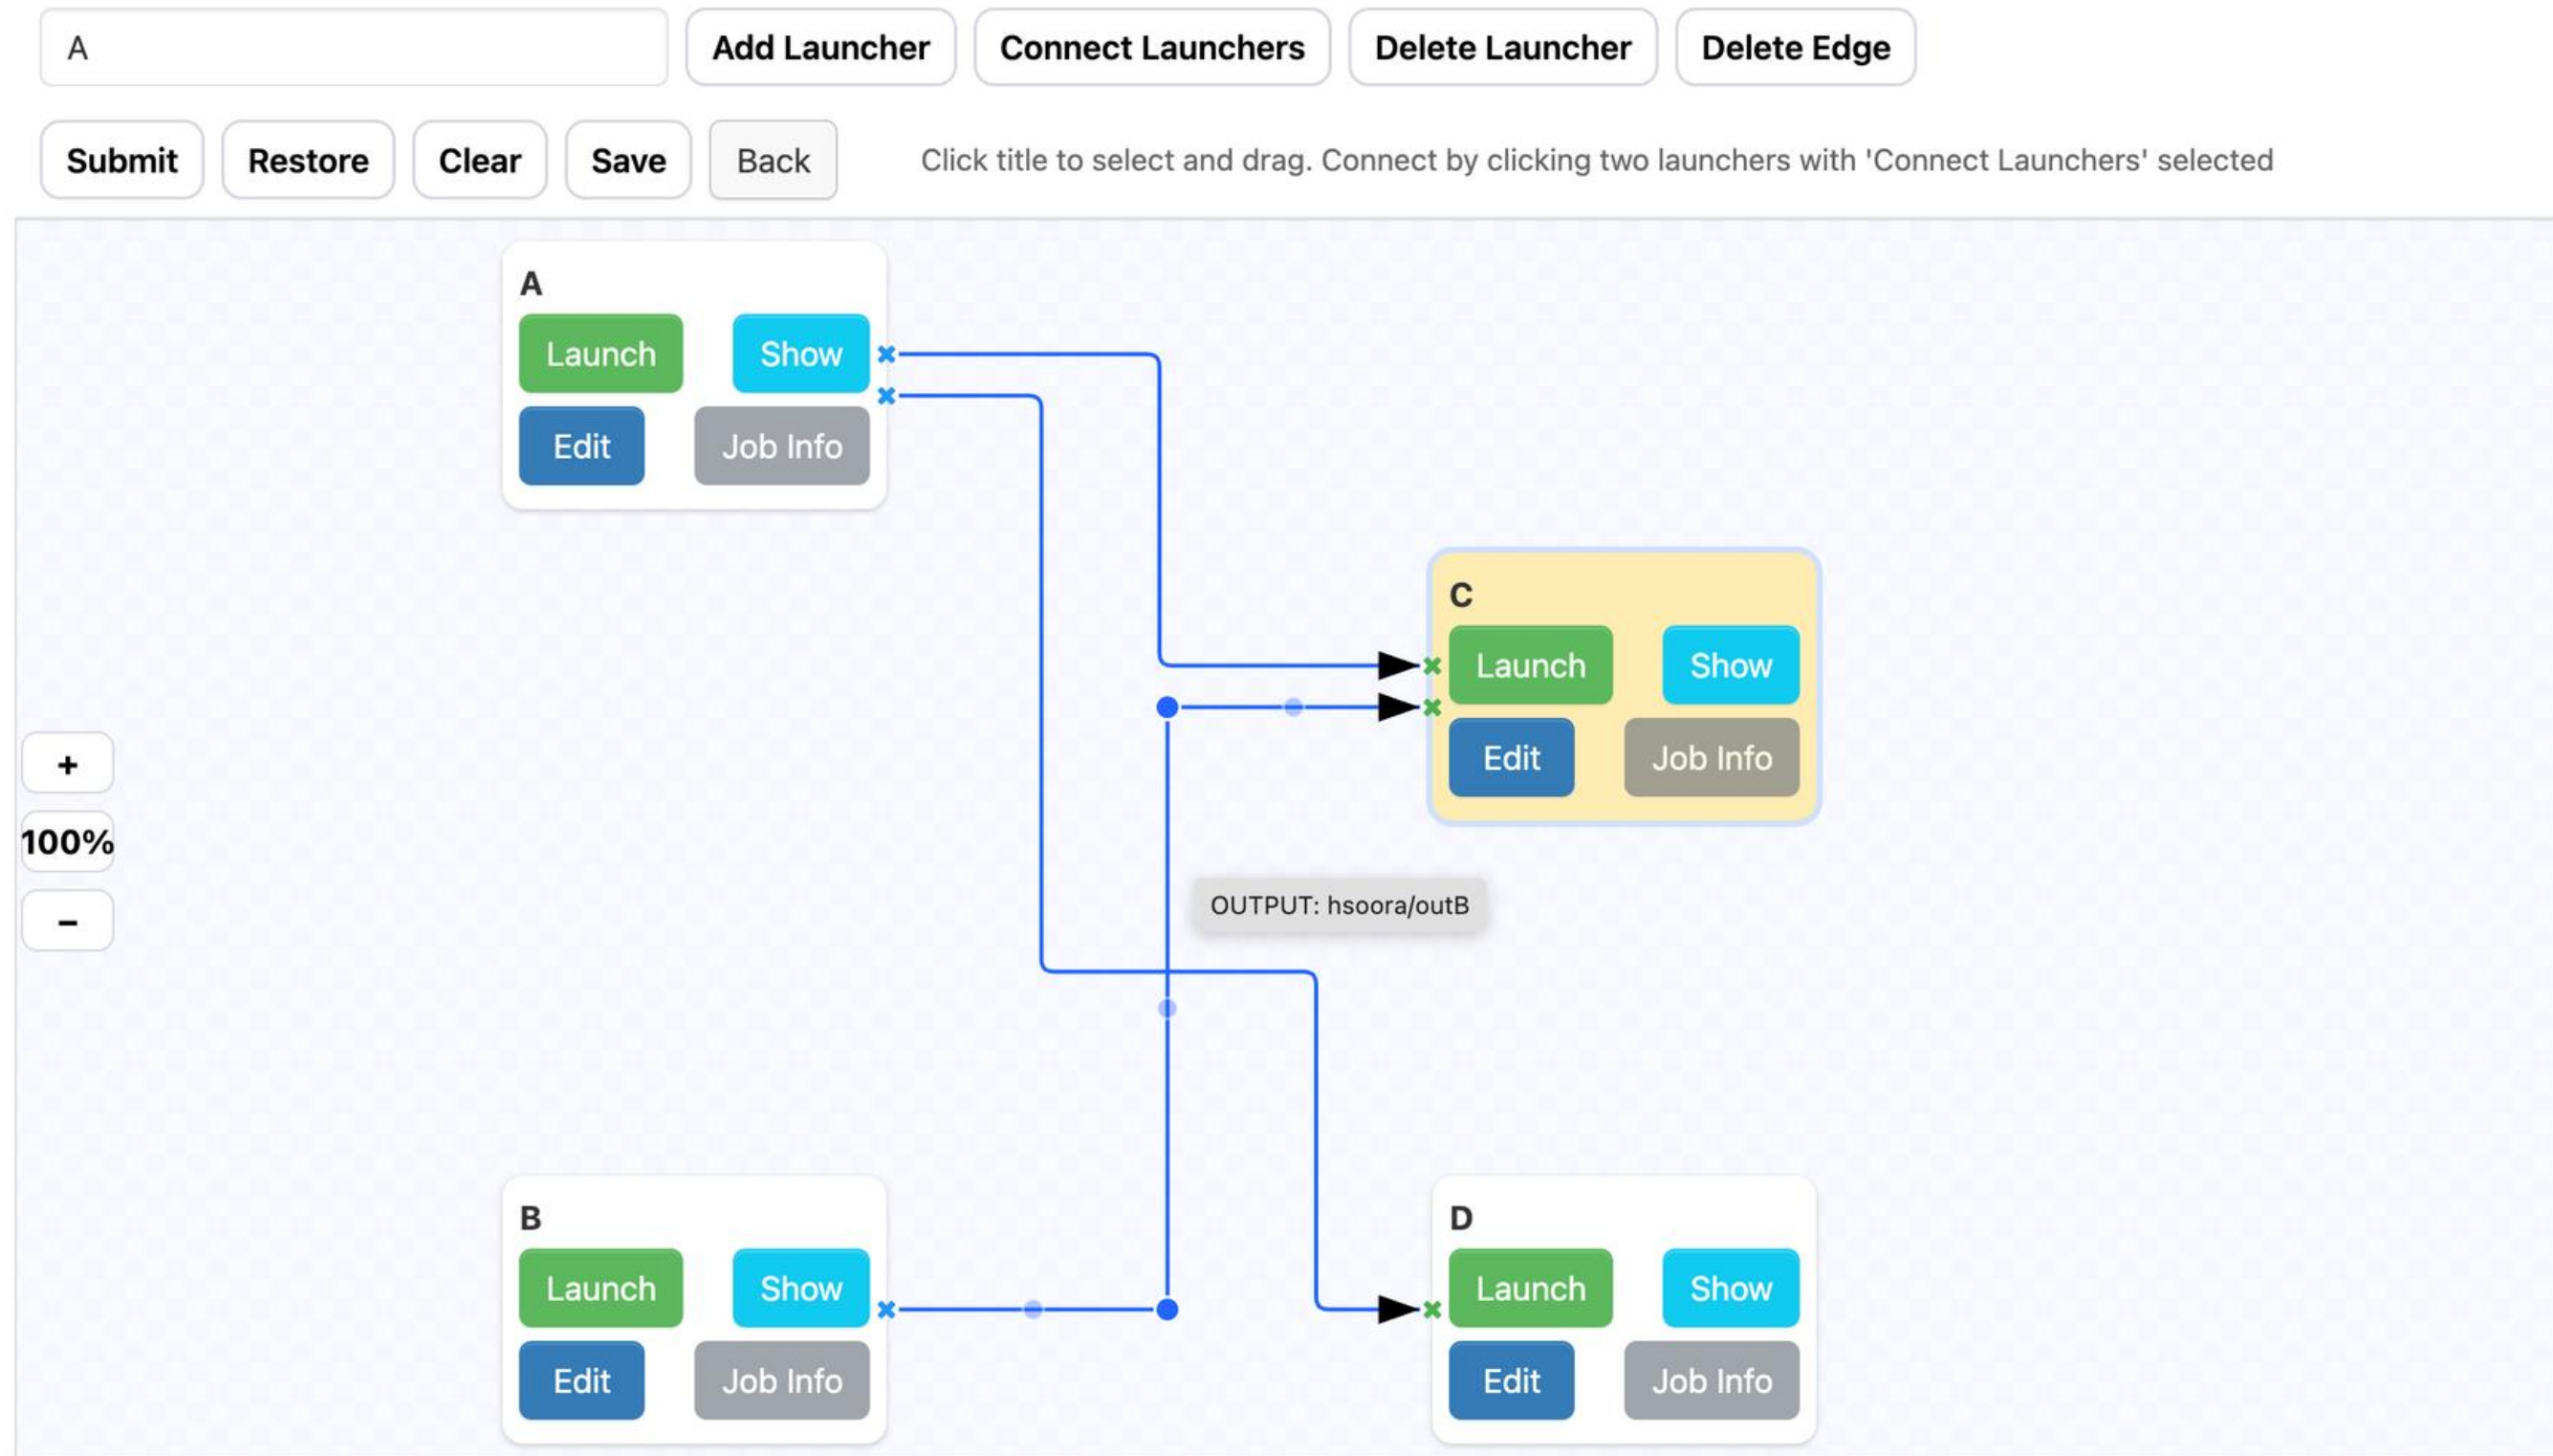
Port Type

None Input Port Output Port

Input Ports Remaining: 5 Output Ports Remaining: 3


Edit Remove

Orthogonal Edges as Output



Design Diagram

Front-End (JavaScript, Html)




Detect cycles on new edges
between launchers

Draw dependencies via
drag-and-drop. Show
success / failure.

Back-End (Ruby on Rails)




Attach current launcher's
job-id to its dependents



Use edge dependencies
to determine launcher
execution order

External



Future Directions

Accessibility

In Progress

WCAG-compliant UI, full keyboard navigation

User Documentation

In Progress

Expressive user documentation, demo examples, design documentation and feature tracking

Future Directions

Accessibility

In Progress

WCAG-compliant UI, full keyboard navigation

User Documentation

In Progress

Expressive user documentation, demo examples, design documentation and feature tracking

Conditional Edges

Planned

Conditional dependencies to submit launchers based upon if current launcher has failed or passed

Parallel Submission

Planned

Submit independent tasks simultaneously using native parallel execution
Ex. sbatch a & b & c

Future Directions

Accessibility

In Progress

WCAG-compliant UI, full keyboard navigation

User Documentation

In Progress

Expressive user documentation, demo examples, design documentation and feature tracking

Conditional Edges

Planned

Conditional dependencies to submit launchers based upon if current launcher has failed or passed

Parallel Submission

Planned

Submit independent tasks simultaneously using native parallel execution
Ex. sbatch a & b & c

Distributed Workflows

Ideas

Running launchers in separate Clusters and bookkeeping input-output dependencies

Workflow Merging

Ideas

Merge two or more workflows depends on merging multiple launchers of the same type

Feature Comparison

Feature	Pegasus	Airflow	Nextflow	Snakemake	OOD Workflow
Language	Python API	Python	Groovy DSL	Python DSL	GUI (No code) ✓
Execution Model	DAG Mapping	Schedule DAG	Dataflow	Rule-based	Scheduler native ✓
Scripting Needed	Yes	Yes	Yes	Yes	No ✓
GUI Builder	No	No	No	No	Yes ✓
Scheduler Agnostic	Partial	N/A	Yes	Yes	Yes ✓
HPC Native	Yes	Poor	Yes	Yes	Yes ✓
Live Monitoring	Dashboard	Web UI	Tower (paid)	CLI	Built-in ✓
Learning Curve	High	Medium	Medium	Low-Med	Low ✓
Data Movement	Auto staging	Manual/Hooks	Symlink/Stage	Shared FS	Scheduler native
Distributed Compute	HTCondor / Grid	Celery / K8s	Native / Cloud	Native Slurm	Scheduler native
Workflow Size	Massive (10 ⁶ tasks)	Medium	Large (10 ⁵ tasks)	Large (10 ⁵ tasks)	Small (10 ³)
Dynamic Branching	High (Programmatic)	High (Pythonic)	High (Dataflow)	High (Python DSL)	Basic / Limited

Questions?

Thank you for your time!



Code v4.1



Phase II



Documentation



LinkedIn