Running Hydraulics Simulations at Scale Using Inductiva Python API

Luís Sarmento, Hugo Penedones, Sérgio Santos, and Paulo Barbosa
Inductiva.AI - Porto, Portugal
e-mail: {sarmento | hpenedones | ssantos, | pbarbosa}@inductiva.ai

***Abstract*—This demo will introduce participants to the Inductiva cloud-based High Performance Computing (HPC) platform, designed for large-scale simulations. Inductiva provides pre-installed ready-to-use simulation models for marine, coastal, and hydraulic projects. Researchers and engineers can run these simulations immediately using Python scripts. In addition to scripting, users can also interact through an intuitive web interface that simplifies data management, resource allocation, and cost control, giving them full oversight of their simulations.**

Keywords-Cloud-based Computing, Numerical Simulation, Python API.

I. Introduction

Despite advancements in simulation technologies, engineers and scientists still face challenges in running numerical models, from configuring simulators and setting up suitable hardware to managing large data volumes.

To help engineers and scientists run large-scale simulations, companies like Rescale [1] and Sabalcore [2] introduced cloud-based solutions that enable simulation jobs on powerful cloud machines. However, these platforms often lack model-specific support, particularly in fields like hydraulics, coastal engineering, and marine sciences, and are primarily optimized for single simulation runs, making them less suitable for workflows requiring hundreds or thousands of simulation variations. Such scalability is essential for exploring design spaces, conducting sensitivity analyses, or generating training data for Physics-Artificial Intelligence (AI) models.

Inductiva offers a cloud-based HPC infrastructure controlled via simple Python scripts and a user-friendly web interface, streamlining complex simulation tasks so experts can focus on their models without concern for underlying infrastructure.

This demo introduces Inductiva's Application Programming Interface (API)-driven infrastructure through several use cases, and demonstrates why Inductiva is an ideal solution for automating complex simulation workflows. In Sections II and III, we will cover how to submit simulations on Inductiva's platform using Python, with practical examples of hydraulics models like Reef3D [3] and SWASH [4], both of which come pre-installed. We will also explore how the platform allows users to generalize simulation parameters and run simulations in parallel to explore different scenarios. Sections IV and V cover Inductiva's task and cost-management features, while Section VI presents supported hydraulics simulators. Section VII concludes the work.

II. The Infinite Laptop Experience

The workflow with Inductiva is simple: users start running their simulation on their local machine and when more computational power is needed—such as for running a longer version of the simulation—users can redirect their workloads to cloud machines equipped with hundreds of central processing units (CPUs), using Inductiva's Python API. Consider the example in **Figure 1**, where a user switches from running a Reef3D simulation on their local machine to a cloud machine with 112 virtual CPUs (vCPUs) provided by Inductiva, using a few lines of Python code:

```
import inductiva

my_machine = inductiva.resources.MachineGroup(
    machine_type="c2d-highcpu-112",
    data_disk_gb=20)
my_machine.start()

reef3d = inductiva.simulators.REEF3D()

simulation_task = reef3d.run(
    input_dir="./3D_Dam_Break_with_Obstacle",
    on=my_machine,
    n_vcpus=112,
    storage_dir="3D_dam_break_with_obstacle")

simulation_task.wait()
simulation_task.download_outputs()

my_machine.terminate()
```

Figure 1. Starting a Reef3D simulation on the cloud from a local machine.

In addition to running simulations on a single large machine, as shown above, users can also set up groups of machines to operate as a Message Passing Interface (MPI) cluster. This allows them to scale their computations even further without worrying about the resource allocation or formalities required by traditional HPC facilities.

III. A Tool for Optimal Design and Physics-Machine Learning (ML)

Inductiva's API-driven computational infrastructure provides a fresh approach for thinking about marine and coastal engineering projects. By orchestrating multiple simulation jobs in parallel with simple Python scripts, users can systematically explore different parameter configurations or scenarios for a single base simulation. With Inductiva's *templating mechanism*, users can transform specific parameters in the simulation configuration files into variables that they can easily adjust programmatically through Python.

As an example, the code below (see **Figure 2**) illustrates how to run five SWASH simulations on five different machines, each testing a different value for the base water level. The API's TemplateManager class is used to automatically adjust the water level in the SWASH simulation configuration

Courtesy of IARIA Board and IARIA Press. Original source: ThinkMind Digital Library https://www.thinkmind.org

files. All simulations are executed in parallel, running in the background to allow for simultaneous processing.

Using a simple “FOR loop", users can launch thousands of variations of a base simulation case. This capability allows users to: **i)** explore different parameters in optimal design projects, **ii)** perform sensitivity analysis on models, or **iii)** generate large-scale synthetic data for training Physics-Machine Learning (ML) models (see one of our tutorials [5] where we explain how to use the API to generate the data required to replicate the work by [6]).

```
import inductiva

# Start 5 preemptible machines
machines = inductiva.resources.MachineGroup(
    machine_type="c2-standard-30",
    num_machines=5,
    spot=True)
machines.start()

swash = inductiva.simulators.SWASH()

water_levels = [3.5, 3.75, 4.0, 4.5, 5.0]

# Launch one simulations for each water level.
for i, water_level in enumerate(water_levels):
    target_dir = f"./my_outputs/swash-sim-{i}"
    inductiva.TemplateManager.render_dir(
        source_dir=template_dir,
        target_dir=target_dir,
        water_level=water_level)

    simulation_task = swash.run(
        input_dir=target_dir,
        sim_config_filename="input.sws",
        on=machines)
```

Figure 2. Running five variations of a SWASH simulation in parallel.

Unlike many providers of cloud resources for simulation that focus on providing resources to run *a single* simulation, Inductiva distinguishes itself by offering an efficient and convenient solution for users who need to run multiple variations of a base simulation case. This capability empowers users to explore different scenarios—something that is typically overlooked by other platforms.

IV. Task and Data Management

One of the key challenges when running simulations at scale is managing all the simulation tasks, especially when it comes to tracking inputs, outputs, and statuses across different projects. To address this, Inductiva provides the **Inductiva Console** [7], a web-based **User Interface (UI)** that enables users to monitor the status of all tasks, organize them by project, download outputs, and track resource usage, including costs. Users can also get an overview of all their activity over time and drill down into specific dates to review the simulation tasks executed on those days.

V. Performance Metrics and Cost Control

Numerical simulation is costly. The performance-to-cost ratio of the computational infrastructure, along with the project’s budget, determines how many and how large simulations can be run. To address this, Inductiva maintains full transparency on all costs, including computation, storage, and auxiliary tasks (e.g., compressing and moving output data).

The platform also helps users save by optimizing resource selection and configuration. A key strategy is using “spot instances”—cloud machines that can be preempted by the provider but are up to 91% cheaper than standard instances [8].

VI. Available Hydraulics Simulators

Table 1 lists the simulation models currently supported by Inductiva for marine, coastal, and hydraulic sciences. New simulators are regularly added. To suggest a specific simulator to be integrated, please e-mail *support@inductiva.ai*.

TABLE I. Ready to use simulators available via Inductiva API

Simulator / Model	Versions
DualSPHysics	5.2.1
Reef3D	24.02
SCHISM	5.11
SPlisHSPlasH	2.13.0
SWAN	41.45
SWASH	9.01A, 10.01
XBeach	1.23, 1.24

VII. Conclusion

Inductiva offers an efficient and scalable solution for engineers and researchers to run simulations. It integrates cloud resources into a Python API and an intuitive web interface. This enables users to focus on analysis and innovation in their respective fields engineering. Future work will focus on enhancing platform stability to support increasingly complex simulations, lowering the barrier of entry by making the platform more user-friendly, and integrating additional simulators.

Free experimental access to Inductiva can be obtained by registering on Inductiva’s website or directly via [7].

VIII. Conclusion

References

[1] Rescale, *Why rescale: Cloud-based high-performance computing*, Accessed: 2024-11-11, 2024.

[2] Sabalcore, *Sabalcore cloud computing solutions*, Accessed: 2024-11-11, 2024.

[3] Reef3D, *Reef3d: Open-source computational fluid dynamics*, Accessed: 2024-11-11, 2024.

[4] SWASH, *Swash: Simulating waves nearshore*, Accessed: 2024-11-11, 2024.

[5] *Generating synthetic data for training physics-ml models*, https://tutorials.inductiva.ai/generating-synthetic-data/synthetic-data-generation-1.html, Accessed: 2024-19-23.

[6] A. Sanchez-Gonzalez *et al.*, “Learning to simulate complex physics with graph networks,” in *International Conference on Machine Learning*, 2020.

[7] *Inductiva web console*, https://console.inductiva.ai/, Accessed: 2024-19-23.

[8] *Spot vms: Affordable compute instances suitable for batch jobs and fault-tolerant workloads*, https://cloud.google.com/spot-vms, Accessed: 2024-19-23.

Courtesy of IARIA Board and IARIA Press. Original source: ThinkMind Digital Library https://www.thinkmind.org