Towards an AI-driven scientific workflow

Gnosis Research Center @ ILLINOIS TECH

Shazzadul Islam Jaime Cernuda Isa Muradli Anthony Kougkas Xian-He Sun

sislam6@hawk.illinoistech.edu jcernudagarcia@hawk.illinoistech.edu imuradli@hawk.illinoistech.edu akougkas@illinoistech.edu sun@illinoistech.edu

Abstract

Modern scientific computing involves complex, disconnected tools and technologies that require data gathering, resource allocation, application deployment, and output analysis, all of which distract researchers from scientific discovery. Large language models (LLMs) can simplify these workflows by utilizing predetermined code, executed by the agent, known as tools, to interact with external systems. The Model Context Protocol (MCP) provides a standardized framework to define and communicate with these tools; however, scientific computing environments have unique requirements that general-purpose tools cannot address. To address this gap, we present *IOWarp-AI*, a framework designed to enable AI-driven scientific workflows. As part of IOWarp-AI, our key contributions are:

- **IOWarp-MCP**: a suite of MCPs for scientific systems, designed with data reduction techniques for a large volume of HPC data.
- **IOWarp-Observer**: A design for an end-to-end observability layer that captures the entire chain of thought of the LLM and the interactions with the tools and real system.

IOWarp-MCP: Access to Scientific Resources

IOWarp-MCP offers domain-specific MCP servers that seamlessly integrate with HPC workflows. HPC workflows generate massive datasets, use domain-specific file formats (like ADIOS BP5, HDF5), and require complex resource management across distributed systems. Standard AI tools cannot handle these data volumes or understand scientific workflows. Two core design features make this possible:

- **Chunked I/O access** partitions large variables into dynamically sized segments; metadata supports *lazy evaluation* and prevents memory exhaustion.
- **Label-based filtering** applies selective data reduction before ingestion by the AI model, lowering context overhead while preserving semantics.

Use Case 1: EarthScope Dataset Gathering via NDP-MCP

Figure 1: Using IOWarp Context to download and visualize the EarthScope dataset from the National Data Platform correctly with Claude Code + Sonnet 4.

- **Dataset:** Latest EarthScope dataset accessed through the NDP service, including GeoJSON metadata and CSV seismograph records.
- **Execution:** Agents used the NDP-MCP to query the dataset, download both file formats, and parse the seismic data.
- **Prompt:** "Find the latest dataset of the EarthScope organization, retrieve the GeoJSON and CSV files, and plot the seismograph data with each axis in a separate subfigure."
- **Result:** With the IOWarp Context active, Claude Code (Sonnet 4) successfully called the appropriate MCPs, retrieved the files, and generated a multi-axis visualization in $\sim$2 minutes. Without the IOWarp Context, the agent failed to complete the workflow.

Use Case 2: Pipe Flow Data Visualization via ParaView-MCP

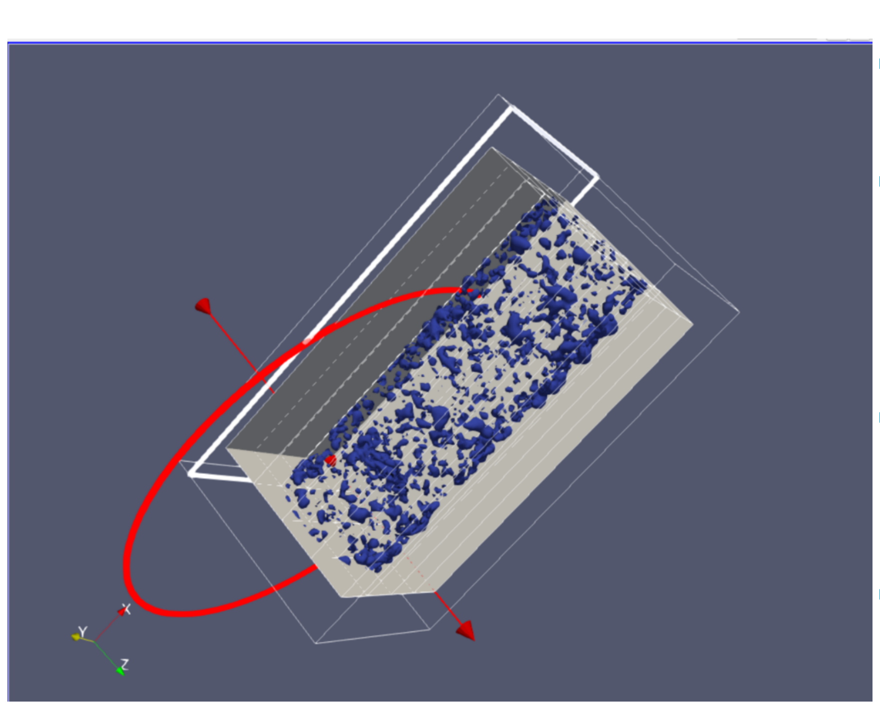

Figure 2: Pipe flow isosurface and orthogonal slices.

- **Dataset:** Incompact3d pipe flow simulation (`data.bp5`) with pressure and velocity fields.
- **Execution:** Agents used **ParaView-MCP**, a service that connects AI systems to the ParaView visualization engine, to load the dataset and generate the 3D visualization.
- **Prompt:** "Explore the pipe flow data. Generates an isosurface of pressure at `pp=0.1` plus three orthogonal slice planes (X, Y, Z)."
- **Result:** Produced a combined 3D flow visualization without manual setup, highlighting internal flow structures.

Use Case 3: LAMMPS (BP5) Data Analysis via ADIOS-MCP

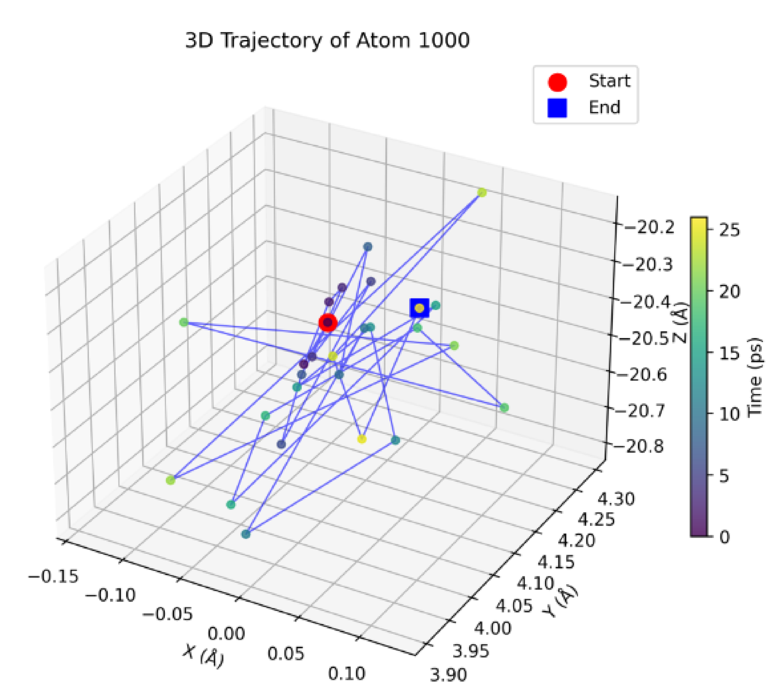

Figure 3: Pipe flow isosurface and orthogonal slices.

- **Dataset:** LAMMPS simulation of an FCC gold crystal, heated from 1 K to 2500 K over 100 ps, stored in BP5 format.
- **Execution:** The agent accessed the dataset through **ADIOS-MCP**, a service built on the ADIOS2 framework that allows AI systems to read and analyze BP5 files efficiently.
- **Prompt:** "Explore the dataset. Plot the trajectory of a single atom (atom ID as parameter). Output PNG."
- **Result:** Execution completed in ~2.5 minutes with Claude Sonnet 3.5, compared to ~12 minutes manually. Generated parameterized script and trajectory plot.

Use Case 4: Automated IOR Deployment via Jarvis-MCP

- **Execution:** Agents used **Jarvis-MCP**, a pipeline orchestration that automates setup, configuration, and execution, to evaluate automated **IOR** deployment (parallel I/O benchmark) across seven agent-model combinations and five prompt variations.
- **Prompt:** "Initialize the Jarvis pipeline, configure node allocation from Jarvis.hostfile, build the environment, append the IOR package, and deploy."
- **Result:** Agents deployed IOR in $\sim$0.5–2 min vs. $\sim$5–10 min manually, highlighting tradeoffs in **speed**, **reliability**, and **completeness**.

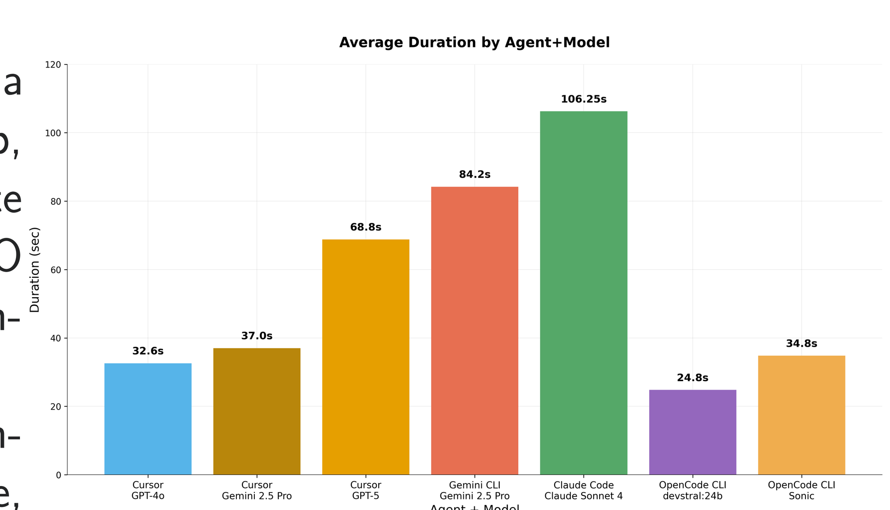

Figure 4: IOR deployment evaluation.

IOWarp-Observer: Reproducible AI Science

AI agents are powerful but inherently non-deterministic; they may not always produce the same result twice. Scientific reproducibility requires detailed documentation of computational processes.**IOWarp-Observer** addresses this challenge.

- Connects to any AI agent and captures all interactions: **user** $\leftrightarrow$ **agent** and **agent** $\leftrightarrow$ **tool**.
- Logs stored in a **time-series DB** with **dashboards** for monitoring and debugging.
- Ensures **traceability, reproducibility, and transparency**.

Figure 5: Visualization of the IOWarp-Observer log collection.

Conclusion

IOWarp-AI streamlines HPC workflows with domain-specific MCPs, delivering faster and reproducible deployments. Our results demonstrate significant reductions in execution time and configuration complexity, enabling researchers to focus more on scientific discovery.

Acknowledgments

This material is based upon work supported in part by the National Science Foundation (NSF), Office of Advanced Cyberinfrastructure (OAC) under Grants 2411318 and 2313154.

IOWarp-MCP

Funded by

National Science Foundation

IOWarp Community