

1 Abstract

This paper...

Explain briefly the paper and what it does.

2 Introduction

Scientific Workflow Management Systems (SWMSs) are an essential tool for automating, managing, and executing complex scientific processes involving large volumes of data and computational tasks¹. Traditional SWMSs employ a linear sequential approach, in which tasks are performed in a pre-defined order, as defined by the workflow. While this linear method is suitable for certain applications, it might not always be the best choice: processing sequentially can prove inefficient in cases where some processes can be done in parallel, or where the next step should adapt to the previous one. For these use-cases a dynamic scheduler is required, of which *Managing Event Oriented Workflows*[1] (MEOW) is one.

MEOW employs an event-based scheduler, in which jobs are performed non-linearly, triggered based on events². By dynamically adapting the execution order based on the outcomes of previous tasks or external factors, MEOW provides a more efficient and flexible solution for processing large volumes of experimental data³.

What else should be included in the introduction?

2.1 Problem

In its current implementation, MEOW is able to trigger jobs based on changes to monitored files. This covers a wide range of scenarios, specifically where the data processing workflow involves the creation, modification, or removal of files. By monitoring file events, MEOW's event-based scheduler can dynamically execute tasks as soon as the required conditions are met, ensuring efficient and timely processing of the data.

While file events work well as a trigger on their own, there are several scenarios where a different trigger would be preferred or even required, especially when dealing with distributed systems or remote operations. To address these shortcomings and further enhance MEOW's capabilities, the integration of network event triggers would provide significant benefits in several key use-cases.

Firstly, network event triggers would allow for manual triggering of jobs without the need for direct access to the monitored files. This is particularly useful in scenarios where human intervention or decision-making is required before proceeding with the subsequent steps in a workflow. While it is possible to manually trigger job using file events by making changes to the monitored directories, this might lead to an already running job accessing the files at the same time, which could cause problems with data integrity.

Secondly, incorporating network event triggers would facilitate seamless communication between parallel jobs, ensuring that tasks can efficiently exchange information and synchronize their progress.

Finally, extending MEOW's event-based scheduler to support network event triggers would

¹citation?

²citation?

³citation?

enable the simple and efficient exchange of data between workflows running on different machines. This feature is particularly valuable in distributed computing environments, where data processing tasks are often split across multiple systems to maximize resource utilization and minimize latency. By leveraging network event triggers, MEOW would be better equipped to manage complex workflows in these environments, ensuring seamless integration and streamlined data processing

2.2 Background

2.2.1 The structure of MEOW

The MEOW event-based scheduler has three main parts: *monitors*, *handlers*, and *the conductor*.

Monitors monitor for triggering events. They are initialized with a number of *patterns*, which describe the triggering event. When a pattern's triggering event occurs, the monitor signals to the conductor that the pattern has been triggered.

Handlers perform actions and jobs on behalf of the scheduler. They are initialized with a number of *recipes*, which describe the action to be taken. The handler starts a job when signal to do so by the conductor.

The conductor handles the jobs queue. It is initialized with a number of rules, which a pattern paired with a recipe. When a monitor sends it a triggered pattern, the rules are checked for that pattern. If one or more rules contain that pattern, the corresponding recipes are triggered in their handler.

2.2.2 The `meow_base` codebase

Specific (but not too granular) implementation details of `meow_base`.

3 Method

Explain the code I wrote and why I made those choices.

- The `socket` library
- Expanding on existing code, reusing boiler-plate code
- Test-driven development
- Experiments with triggering on packet
 - Removes the ability to send arbitrary data

4 Results

Does it work? How well?

4.1 Discussion

With the hindsight of the results, what could I have done better?

4.2 Future Work

What should someone do if they want to fix my mistakes, or expand on them further.

5 Conclusion

Did I succeed in what I wanted to do?

References

- [1] David Marchant. “MEOW - Enabling Dynamic Scheduling of Scientific Analysis”. PhD thesis. University of Copenhagen, May 2021.