

Exploration of Workflow Management Systems

Emerging Features from Users Perspectives

Ryan Mitchell*, Loïc Pottier*, Steve Jacobs[†], Rafael Ferreira da Silva*, Mats Rynge*, Karan Vahi*, Ewa Deelman*

* Information Sciences Institute, University of Southern California, Marina Del Rey, CA, USA

[†] National Ecological Observatory Network (NEON), Boulder, CO, USA

{rmitchel,lpottier,rafsilva,rynge,vahi,deelman}@isi.edu, sjacobs@battelleecology.org

Abstract—There has been a recent emergence of new workflow applications focused on data analytics and machine learning. This emergence has precipitated a change in the workflow management landscape, causing the development of new data-oriented workflow management systems (WMSs) in addition to the earlier standard of task-oriented WMSs. In this paper, we summarize three general workflow use-cases and explore the unique requirements of each use-case in order to understand how WMSs from both workflow management models meet the requirements of each workflow use-case from the user's perspective. We analyze the applicability of the two models by carefully describing each model and by providing an examination of the different variations of WMSs that fall under the task-driven model. To illustrate the strengths and weaknesses of each workflow management model, we summarize the key features of four production-ready WMSs: Pegasus, Makeflow, Apache Airflow, and Pachyderm. To deepen our analysis of the four WMSs examined in this paper, we implement three real-world use-cases to highlight the specifications and features of each WMS. We present our final assessment of each WMS after considering the following factors: usability, performance, ease of deployment, and relevance. The purpose of this work is to offer insights from the user's perspective into the research challenges that WMSs currently face due to the evolving workflow landscape.

Index Terms—Scientific workflow, Workflow Management System, Task-driven, Data-driven.

I. INTRODUCTION

In the last two decades, scientific workflows have become mainstream thanks to their ability to empower scientific discoveries in virtually all fields of science [1]. During this time, key engineering challenges have been solved and a rich set of abstractions and interoperable software implementations have been developed [2]. These advancements have allowed scientists across various fields to begin reaping the benefits of workflow systems [3]. Traditionally, scientific workflows are described as directed-acyclic graphs (DAGs), in which nodes represent computational tasks and edges represent the dependencies of those tasks [4]. The traditional approach for orchestrating DAG-based workflows is to use task-based scheduling algorithms that spawn tasks for execution once their dependencies are satisfied. To keep up with increased computing requirements, workflow systems have developed mechanisms to manage the distribution and execution of

tasks on varied and across computing infrastructures such as local servers, campus computing clusters, high performance computing resources (such as XSEDE [5]), and even popular cloud computing platforms. These developments in workflow management are of primary concern to the field of scientific computing, where scientists often run complex pipelines that scale over hundreds and thousands of tasks [2].

There are four general workflow system use-cases that have been identified [3]:

- Traditional scientific compute workflows, as discussed above;
- Data analytics workflows (including big data and machine learning);
- Sensor and Internet-of-Things (IoT) workflows; and
- Commercial, developer, and business-related workflows.

One of the major trends among scientific applications recently concerns big data analytics and machine learning [4]. These data-oriented workflows pose different challenges when compared to traditional workflow structures, and they often require special features such as data provenance, data reproducibility, and special data ingestion features.

A second type of data-oriented workflow that is gaining prominence in the scientific field is sensor-based workflows. Such workflows process data in a continuous fashion, with data being ingested in near real-time from distributed sources (e.g. sensors that stream data to a central endpoint). Here, the ability to trigger computations based on the arrival of new data is paramount, in addition to the ability to replay processing on previous datasets.

Similar to sensor-based workflows, the rapid expansion of the Internet-of-Things (IoT) field also raises the need for new workflow orchestration models [6]. In contrast to large-scale data analytics (which process large amounts of data in parallel), sensor-based and IoT-based workflows generally process smaller amounts of data at one time with a higher data arrival rate, placing a greater importance on new data as compared to old or late data.

Furthermore, in the developer and commercial communities, workflows are becoming increasingly important as developers have started to automate more complex tasks in the software development lifecycle. On the commercial side, businesses are turning to in-house workflow management systems to analyze

their data., and have started to create a new generation of workflow management systems that are DAG oriented and are able to perform batch processing to their own specification (e.g., LinkedIn’s Azkaban [7] or AirBnb’s Apache Airflow [8]). In addition to being used by companies for management or development purposes, these custom-developed solutions are also being used by scientists as building blocks [9], enabling them to create their own light-weight workflow management solutions.

While all of these general use-cases play an important part in defining the next generation of workflows and their requirements, this paper intends to focus on use-cases that apply to the scientific community, specifically traditional compute-based scientific workflows, big data workflows, and sensor-based workflows.

Based on the distinct characteristics and requirements of the different use-cases described, we can generally classify workflow management systems (WMSs) into two distinct categories: traditional, *task-driven* WMSs and modern, *data-driven* WMSs. In the traditional *task-driven* approach, workflow tasks are triggered for execution once all of their parent tasks have completed. A more recent development is the *data-driven* approach, in which tasks in the workflow are triggered by data input and output, rather than task completion dependencies. In this paper, we conduct a study on the key requirements and features that have driven the development of this new paradigm. We explore how workflow systems have addressed recent challenges presented by these new workflow use-cases, and identify open questions that have not yet been addressed by today’s workflow management solutions. We first describe the paradigms in a WMS-agnostic and platform-agnostic manner, and we then present real world use-cases that have benefited from these state-of-the-art workflow system implementations. Note that we do not aim at performing a feature-by-feature comparison of workflow systems. Instead, our goal is to provide our own hands-on experience in dealing with such challenges from the WMS’s and user’s perspectives. It is also important to note that, during the last two decades, many of the authors of this paper have been involved in the scientific workflow community and have contributed to the development of workflow management systems (most notable of which is the Pegasus Workflow Management System [10]). Though Pegasus falls into the traditional, task-driven workflow management model, the authors of this paper are excited and intrigued by the new approaches and use-cases that have recently emerged.

For each general use-case, we have tried to focus on representative workflow systems. We are aware that a plethora of workflow systems have been developed in the recent years, but it is not possible to account for every one of them. The purpose of this paper is to broadly classify use-cases from a scientific computing perspective and help the reader identify the workflow system that is suited for their research needs.

More specifically, this work makes the following contributions:

- 1) We depict the differences between the traditional task-

driven approach and the next-generation data-driven approach for workflow systems. We present several existing WMS, along with their respective features, that exemplify each model. We also provide a detailed discussion about new trends and innovations with the task-driven model.

- 2) We describe three real-world use-cases, match them with a representative WMS, and evaluate selected features of each WMS that benefit the given use-case. We also present the potential challenges that users may face when deploying these use-cases on different cyberinfrastructures.
- 3) We summarize our discussions and present our findings in a manner that aims to assist an end-user in classifying their own workflows and choosing a WMS.

This paper is organized as follows. In Section II, we provide an overview of the background and related work. Section III provides an overview of the requirements for both workflow management models by describing systems which implement each paradigm. In Section IV, we describe real-world use-cases, and describe how a particular WMS (Pegasus, Makeflow, Apache Airflow, and Pachyderm) is beneficial to implementation. Section V is dedicated to illustrating the experience from the perspective of both the user and the WMS. This is done by comparing the usability, performance, and relevancy of each use case. Finally, Section VI summarizes our findings about future scientific workflow management developments.

II. BACKGROUND AND RELATED WORK

A. Traditional Scientific Workflows

One of the first models to represent a sequence of different computations is the directed acyclic graph (DAG) model [11]. An example of the *task-driven* approach, computational tasks, as represented by nodes in the DAG, are the primary units. Many popular WMS in the scientific community, such as Kepler [12], Makeflow [13], Taverna [14], and Pegasus [10] rely on a task-driven approach using a DAG representation. A DAG is a very simple and natural representation that allows WMS to apply efficient scheduling [3] and data management optimizations.

In the past, scientific workflows were traditionally compute-intensive but thanks to GPUs and new memory technologies, many data-intensive scientific workflows have been developed [4]. In addition, from biology to astronomy, the number of scientific domains embracing WMS is constantly growing [2]. The requirements and uses from one community to another, however, are not consistent [3].

Following these trends, new requirements have emerged in the scientific workflow user community, some of which include easy deployment on several cloud and HPC platforms and efficient data management with a strong data reproducibility aspect. In this vein, workflow systems have evolved and adopted new technologies to ensure better reproducibility and easier deployment, including support for containers and cloud-based execution. Several new workflow systems have also adopted an API approach in which users programmatically

define the workflow instead of giving it an abstract definition. An example of such a system is Parsl [15], a Python scripting workflow system that enables users to quickly define their workflows by directly annotating their Python codes.

B. Data-Oriented Workflows

With the scientific workflow community growing in size, more WMS are being developed in response to the community's need for specific features and evolving workflow paradigms. Driven by the popularity of data analytics and machine learning systems, these new WMS are more data-oriented than their traditional counterparts, and have many interesting and modern features not shared by the more traditional systems.

Many studies on data-oriented workflow management have focused on the MapReduce [16] approach and its most known implementation: Apache Hadoop YARN [17]. YARN aims at decoupling resource management from the programming model. From this, several tools have been developed in recent years: Apache Apex [18] is a data-oriented solution built on top of Hadoop and YARN that allows users to express both streaming and batch data pipelines with a DAG-based representation. Besides MapReduce-based solutions, several generalist data-oriented workflow management solutions have been developed, including Nextflow [19] and Pachyderm [20].

Nextflow [19] uses a domain-specific language that allows users to quickly prototype workflows running in containers. An interesting feature of Nextflow is the complete integration with several versioning platforms such as GitHub, which enables the workflow to check for updates and pull data from a given repository. We explain the features offered by Pachyderm in greater details, later in this work.

In this paper, we attempt to distill these recent developments in the field of scientific workflows by presenting a basic comparison between the traditional paradigm of WMS (including Pegasus and Makeflow), a more modern approach of this classic paradigm with Apache Airflow and the advent of newer systems intended for more specific purposes (such as Pachyderm). We also aim to aid users in the increasingly complex process of narrowing down which paradigm and representative WMS best suit their use-case by examining a distinct set of representative workflow systems in a holistic manner.

III. WORKFLOW MANAGEMENT MODELS

In this section, we further define the two models mentioned in the introduction, and provide an overview of their respective requirements and key features. We also introduce several representative workflow management systems using those models.

A. Task-driven Model

Model: The task-driven approach has traditionally been based on the principles of a directed acyclic graph, or DAG. The idea behind the task-driven model is to break a large and complex workflow into a sequence of individual computational tasks. A WMS using the task-driven model usually defines




Figure 1. Two examples of DAG-based task-driven workflows. On 1(b), if in v_2 a given condition is true then v_3 is executed, otherwise v_4 is spawned. In addition, the execution v_6 fails, but as v_5 successfully finished, v_8 is spawned.

the workflow as a graph, with the computational tasks as the primary entity of work, and edges representing data and control dependencies (see Figure 1). A task can start its execution if, and only if, all of its predecessors have successfully completed. The task-driven model is simple to understand and well adapted to heavy computational tasks, and is widely-used by HPC frameworks [21], as well as by numerous theoretical scheduling studies [22]. It is also very efficient, thanks to years of research focused on workflow optimization, scheduling strategies and data management, both from a practical and theoretical points of view. In addition, due to the decades of research and software development many task-driven WMS are mature and production-ready, such as Pegasus [10] and Makeflow [13]. This model allows users to target many different platforms, from large-scale HPC systems to distributed cloud platforms, grid infrastructures, or local clusters.

Recently, several new workflow management systems have been developed to bring more flexibility to the classic task-driven model and to address new requirements raised by modern workflows, with unique features including complex workflow-level and task-level scheduling abilities, conditional task execution, and improved error management and mitigation. To evaluate the task-driven model, we consider three production WMS, Pegasus [10], Makeflow [13], and Airflow [8].

Workflow management systems: Pegasus [10] and Makeflow [13] are two well-established workflow management systems designed to manage and optimize the execution of large-scale scientific workflows on distributed resources, they provide container support on various commercial clouds such as AWS or Microsoft Azure, as well as HPC systems via support of various cluster job managers, including Slurm and PBS. An important difference between them is that Pegasus allows control and data flow to express workflows while Makeflow allows only data flow. Control flow is the ability to define individual tasks and their dependencies in a workflow, often by specifying the executable or type of data transformation, whereas data flow defines the data being passed by the task instead of the task or transformation itself.

Apache Airflow [8] is an Apache Software Foundation project, originally developed by Airbnb and released in 2016,

that aims to provide a lightweight workflow management solution to easily model, maintain, and monitor workflows. In contrast to the traditional task-driven model in which a DAG describes data and/or control exchanges, in Airflow a task is not supposed to exchange data with other tasks—they can only exchange metadata (i.e., only control flow) [8]. Airflow is very modular and provides many pre-built interfaces (*Hooks*) to common clouds and database systems such as Amazon S3, Google Cloud, or HDFS among others, and has a modular execution engine for computational tasks (*Operators*). Users are able to utilize multiple clouds on a single deployment. Airflow also supports containerized execution and orchestration by way of a Kubernetes [23] operator.

Airflow is an excellent example of a next-generation WMS that has unique features as discussed above. For example, in Airflow [8], each DAG is associated with parameters such as a workflow's start date, an end date, a number of retries in case of failure, the delay between each retry, among others. `cron` expressions allow the user to describe a schedule interval. The scheduler runs in the background as a daemon and will pick up or kick off any DAG according to their start dates, end dates, and schedule intervals. Another interesting task triggering concept extending the possibilities of the task-driven model is the ability to trigger tasks without satisfying dependencies (e.g., users can execute a task if one or all of its predecessors have failed, see the red task in Figure 1(b)). This ability can be seen as an exception handling mechanism for workflow execution. Finally, another major refinement when compared to the traditional task-driven approach is the support of *conditional execution*, where a branch of the workflow is executed only if a given condition is satisfied (see Figure 1(b)).

B. Data-driven Model

Compared to the previous approach, data-driven WMS in general provide better data provenance and versioning, better support for cloud-based storage, easier data ingestion, and a good scaling capability via the adoption of highly scalable orchestration solutions such as Kubernetes [23].

Model: In this model, data are the primary units, as opposed to tasks. A data-driven workflow can be represented as a DAG but instead of tasks being individual nodes, a node represents a *data repository* and the edges are the computational tasks. A data repository can be seen as a directory where data are structured as objects or files. A task describes the processing steps to be performed on the data in the incoming repository. A pipeline corresponds to the edge in the classic DAG representation (see Figure 2). Let v_i and v_j be two distinct data repositories, let $e_{i,j}$ be the edge from v_i to v_j . Then, v_i stores the input data used by the pipeline $e_{i,j}$ and v_j stores the output data produced by this pipeline. A successful pipeline will create a new repository for output files, and can be used as input directory for another pipeline. Connecting pipelines to each-other by way of data repositories defines a completed workflow.

Often, in data-oriented workflows, users start processing newer data as soon as they become available, without to have

to trigger the workflow themselves. In certain cases, WMS using the data-driven model can be considered "active," as they proactively monitor each repository for new data, and trigger individual computational pipelines as needed. Notice that this model differentiates from stream-based workflows since computational tasks are not constantly running on computing nodes waiting for the next chunk of data to process.




Figure 2. An example of a data-driven workflow. A piece of data is passing through different tasks, here called pipelines, that compute on the data.

Workflow management system: Pachyderm [20], [24] aims to enable reproducible, collaborative, and scalable data science through a more innovative approach to workflow management. A Pachyderm workflow, called a pipeline, is organized around data repositories (nodes in the DAG) containing data. Whereas Hadoop-based solutions are usually optimized for MapReduce processing, Pachyderm is data- and language-agnostic, meaning that it is not limited to a given data format or programming language to process the data.

Using the active approach previously described, Pachyderm runs a pipeline on the its data and waits asynchronously for new commits (i.e., new data to be processed by the pipeline). Pachyderm is built on top of numerous software layers and runs on top of widely-used commercial cloud providers (Amazon S3, Microsoft Azure, Google Cloud, etc.)

C. Reproducibility

As scientific workflows use increasing amounts of data, reproducibility of results and data provenance become crucial. Task-driven workflows often take the approach that the input data and the description of the workflow are sufficient to reproduce the results [1]. In the data-driven model, each data repository is versioned ensuring a complete data reproducibility and allowing users to execute workflows on each data version available. Note that, this feature has a non-negligible cost in terms of storage. For example, Pachyderm uses a Git-inspired [25] data-versioning system, so users add data to repositories via a *commit* and these data are then processed by the tasks (see Figure 2). Pachyderm versions the pipeline specifications and all the data processed, allowing users to rollback and execute any pipelines on any data that have been versioned. Coupled with a native containerized execution, this enable fully reproducible data pipelines.

D. Cloud-based orchestration

Many task-driven workflows pre-date the cloud, thus these virtual resources were often treated as an additional execution environment. On the other hand data-driven workflows were born cloud-ready and utilize cloud-based container orchestration solutions such as Kubernetes [23], which allow efficient

handling of large amounts of data in order to have a reproducible and an easy deployment procedure on different cloud providers. In Pachyderm's case, each pipeline is contained in a Docker image and all pipelines managed through Kubernetes worker pods.

IV. REAL-WORLD USE CASES

To examine each workflow management paradigm and its associated workflow management systems, we chose three real-world workflow use-cases based on the four general workflow categories described in the introduction, choosing to omit the commercial and developer use-case as it is less pertinent to scientific workflows and interests. These real-world use-cases are as follows: (i) a traditional scientific compute workflow used to evaluate Pegasus and Makeflow; (ii) a data streaming workflow to evaluate Airflow; and (iii) a sensor-based workflow to evaluate Pachyderm.

A. 1000 Genomes: Traditional Compute Workflow

The 1000 Genomes workflow is a traditional DAG-based bioinformatics workflow, fetching and parsing data from the 1000 Genomes Project [26]. The workflow aims to analyze mutational overlaps in humans, ultimately allowing statistical evaluation of potential disease-related mutations. The Project's Phase 3 and superpopulations data is downloaded and parsed (*Individuals* and *Populations* tasks), sorting amino acid substitutions and determining their potential phenotypic effects (*Sifting* tasks). Analysis is then performed in the *Frequency overlap* and *Pair overlap* tasks (see Figure 3).




Figure 3. 1000 Genomes Workflow.

The workflow itself can be considered an example of a “classic” scientific DAG-based workflow because of its consistent dependencies and static nature. The workflow satisfies all DAG properties—each computational task in the workflow depends on the completion of previous parent tasks, and no changes to the workflow structure occurs. Additionally, the workflow is not dynamically triggered, i.e. it starts based on a user's command, and all workflow input data are known *a priori*. This workflow use-case has no particular requirements, only requiring data-flow dependencies and an available network connection to retrieve the dataset.

B. CASA: Streaming Data Workflow

Streaming data workflows have become increasingly popular in recent years. The University of Massachusetts' Collaborative Adaptive Sensing of the Atmosphere (CASA) project utilizes a sensor-based streaming data workflow for their Dallas/Fort Worth (DFW) weather radar testbed, which aggregates data from eight short-range weather radar sensors, providing higher data precision, accuracy and timeliness versus other, longer-range radar systems [27].




Figure 4. CASA Wind Velocity Workflow.

This paper focuses on a subsection of the CASA workflow pipeline [27], the calculation of maximum wind velocity and the sending of a customizable email alert to warn certain entities such as hospitals or airports of impending high wind velocity. The CASA workflow takes input data from each weather radar sensor (radar_N.netcdf.gz files), unzips the data (unzip tasks), computes the maximum wind velocity around the DFW area (max_velocity task) and creates a graphical image of this data (merged_netcdf2png task). It then outputs velocity data in geojson files (mvt task). Finally, these files are used to create high wind velocity alerts in the pointalert task (see Figure 4).

The CASA use-case has two basic requirements. Most importantly, the streaming property of the workflow requires workflow triggering every 75 seconds using new just-in-time data ingested since the last workflow run. Furthermore, computational executables are stored in a Docker container. While the use of the container is not required, its usage allows for easy portability and reproducibility.

C. NEON: Sensor-based Data-driven Workflow

The National Ecological Observatory Network (NEON) is an NSF open-science facility collecting ecological data from sensors across the US with the objective to study ecological processes and changes. NEON's instrument data pipeline takes raw sensor data from terrestrial and aquatic sensors and processes it for publication. Raw sensor data ranges from resistance values and voltage at a low frequency of collection, to high frequency sonic anemometer data. The workflow converts this data into the appropriate unit of measurement using calibration coefficients, and performs QA/QC steps on

the data to ensure the quality (see Figure 5). The workflow is a linear workflow that, as long as data is coming through the pipeline, gathers data from multiple sources (metadata, calibration data and raw sensor data) and process them to create significant output results.

An essential requirement of the NEON instrument data pipeline is the ability to reprocess data. This is necessary for several reasons, including improved algorithms, more recent calibration data, or late data. Being able to process data that arrives at a later date is of utmost importance to this workflow, and NEON notates missing data with a "null flag" in their data repositories as a placeholder.




Figure 5. NEON processing workflow. The rectangles represent the pipelines and the cylinders the data repositories.

V. HOLISTIC EVALUATION

In this section, we present the results of a holistic evaluation of WMS for the models discussed above. We do not attempt to make a singular recommendation, nor is this an analysis of the different features and limitations of each WMS. As expressed above, different use-cases have different feature requirements, and it's impossible to dictate a singular WMS that meets all requirements possible. Rather, this work aims to help users decide which workflow management model, task-driven or data-driven fits their needs best, and provide an example WMS for each model to further demonstrate how a particular model can compliment a given workflow. In order to perform a thorough evaluation of each workflow management system, we used the following criteria.

Setup and deployment: We tested the installation process on a local cluster and using a cloud platform (AWS), using publicly-available documentation and user-facing support channels.

Workflow implementation: We examined the level of knowledge and effort required to model the relevant use-case and workflow.

Workflow execution: We researched features relating to workflow triggering, scalability, workflow resiliency, and how each WMS handles failures.

Data management: We studied how data is managed throughout the workflow execution, including whether data is transferred between computational tasks or workflows.

To test each WMS, two types of deployments were used in order to account for the majority of use-cases. First, initial WMS installation, testing, and workflow modeling was done on a local system. This deployment was used for initial

testing as many WMS intended for scientific workflows are still deployed and used on local computers or clusters, and this use-case has its own unique challenges. Second, Airflow and Pachyderm, both considered cloud-ready WMS, are discussed with special attention paid to cloud deployment, which presents separate challenges, and is increasingly important as some scientific users have begun transitioning their workflows from local to cloud deployments.

A. Traditional Task-driven WMS: Pegasus and Makeflow

In the case of the 1000 Genome use-case, the workflow uses a Python workflow generation script to enumerate files in the 1000 Genome dataset, and assign them as inputs of various tasks inside the workflow pipeline without knowing the exact file names, facilitating an easy workflow modeling process.

Pegasus and Makeflow are excellent WMS for the 1000 Genome use-case, due to its traditional workflow structure. Pegasus Python API allows the user the flexibility and ease-of-use of the Python language in creating a workflow pipeline. Whereas Pegasus targets flexible composition of DAG workflows via APIs with emphasis on functionality, Makeflow has a more rigid, yet simple, workflow modeling structure based on GNU make. In this case, dependencies between workflow tasks are automatically inferred from the data flow specified in the workflow description file, which alleviates the user's burden on defining task dependencies.

Setup and deployment: Pegasus installation is relatively simple due to its availability on different official repositories. Since Pegasus relies on HTCondor [28] as a task scheduler and interface to other cluster managers, additional effort is required to properly configure and describe the resources. Pegasus also supports cloud deployments on commercial cloud providers, and NSF-cloud infrastructures [27]. Similarly, due to its simple workflow structure model, Makeflow installation is relatively easy—though it is assumed a workload manager is already available (e.g., WorkQueue). Both Pegasus and Makeflow were not natively designed for cloud support, rather cloud resources generally are manually deployed in the cloud (e.g, AWS or Azure).

Workflow implementation: Pegasus provides a rich set of APIs (Python, Java, R, and Perl) for modeling workflows. These APIs provide a versatile mechanism for modeling large-scale workflows ($O(10^6)$ tasks), which is not often practical via graphical interfaces—though the entry barrier for non-expert users is higher. In Makeflow workflows are defined like 'Makefiles'. This structure is fairly simple for defining workflows where the data flow drives the tasks dependencies. The drawback of this approach is the limited flexibility for defining complex workflow patterns or control flows compared to more complex APIs.

Workflow execution: Both WMS supports different workflow execution environments such as containers, and batch scheduler support, which are key for enabling large-scale executions. Pegasus is built for reliability and integrity, featuring several

different types of workflow recovery methods, provenance data, and checkpointing abilities. Makeflow is more simple, it automatically retries failed tasks, but does not provide support for checkpointing or sophisticated error recovery methods. Pegasus workflows also feature several “catalogs” listing the locations of key data resources, executable ‘transformation’ resources, and compute resources, allowing for easy workflow portability with which only the resource configuration needs to be changed.

Data management: Pegasus provides advanced mechanisms to efficiently manage data movement during workflow execution. During the workflow planning phase, Pegasus identifies data locations and augments the workflows with data transfer jobs for staging input and output data from/to storage resources. A wide range of protocols are supported, including access to cloud object storage, via Globus services, etc. Data throttling allows for increased throughput performance. In Makeflow, data is assumed to be directly accessible from the computing node (e.g., shared filesystem) or fetched from a remote source—support is limited to common Internet protocols. However, depending on the execution resources used (e.g. Docker container), Makeflow does support data staging in/out of individual computational tasks.

B. Recent Task-driven WMS

One example of a more recent (i.e., more flexible, supporting containerized execution and cloud-oriented) task-driven WMS is Airflow. Airflow’s *operators* allow users to develop their own notions of what a workflow “task” means. While many traditional WMS limit computational tasks to executables, in Airflow the notion of task ranges from simple tasks such as sending an email, to more complex tasks, such as running executables inside a container or managing a cloud deployment. In addition to many built-in operators ready to use, the triggering system used by Airflow is also extremely beneficial to complex workflows. More robust than a cron-based implementation on top of a traditional WMS, Airflow’s scheduler allows many tasks to be run at custom times or intervals, and monitors them as such. This concept is extremely beneficial for the CASA workflow, which is started every 75 seconds in order to process new streaming data.

Setup and deployment: Airflow is written on Python, which is the only required software prerequisite. Installing Airflow on a local cluster is relatively easy via the pip package-management. However, additional software such as the high-performance execution queue Celery might be needed to adapt Airflow to the user’s requirements. From a cloud perspective, Airflow primarily supports Google Cloud (GC), with extended support for Azure and AWS. Deployment on cloud platforms is simple, thanks to provided scripts (with a slight edge toward GC) but complex features such as cloud autoscaling still require configuration or external tools.

Workflow implementation: Through the extensive library of built-in operators, workflow implementation in Airflow is

extremely flexible, and the Python API structure used to write workflows is easy to learn and understand.

Workflow execution: Airflow’s scheduling features allow users a robust way to trigger their workflows. Airflow also has a built-in task retries parameter, like many other WMS and external software such as Apache Mesos enables checkpointing features.

Data management: Airflow features a basic way to pass data from task-to-task inside a workflow, but this isn’t as robust as other WMS’s data management features. As such, tracking provenance is also more difficult.

C. Data-driven WMS: Pachyderm

As described in the previous sections, Pachyderm is a prime example of a recently-developed WMS using the data-driven management paradigm. Written to allow portable data pipelines, with reproducibility and data provenance, Pachyderm is not explicitly designed for scientific workflows, but can be co-opted for scientific tasks. As detailed in Section III, Pachyderm’s workflow modeling process is interesting in that each task is defined individually as a data pipeline, and several pipelines can be combined together to form a complete workflow.

One important feature of Pachyderm with regards to the NEON workflow pipeline is Pachyderm’s robust data provenance tracking. The NEON project aims at publishing versioned data sets on an online portal to support open-science research. Although NEON can produce versions of these data sets without a solution like Pachyderm, Pachyderm’s provenance tracking features allow NEON to inform the public what has changed between versions of the data sets, and why. NEON also needs to re-run their pipelines with field calibration data. For sensors that are calibrated regularly in the field this is inconvenient, as NEON has to reprocess when those calibrations are stored. For instruments that self-calibrate, and have no defined calibration period, this is borderline impossible without an on-demand reprocessing capability. Pachyderm allows NEON to do this, by triggering only the pieces of the pipeline necessary to run when a calibration change comes in via the commit system.

Setup and deployment: Similarly to other WMS, Pachyderm requires several dependencies, including Kubernetes, which itself requires expert knowledge (just as with HTCondor). Pachyderm is intended for cloud deployments, and several utilities exist to automate installation on cloud platforms like AWS. However, Pachyderm has limited local cluster deployment, requiring a Kubernetes cluster and S3-compatible storage. Compared to Airflow, which allows more fine-grain resource management, Pachyderm is more restrictive—once its Kubernetes cluster is deployed on a given cloud, all pipelines execute on this cloud.

Workflow implementation:

Because Pachyderm triggers pipelines whenever data is committed to a given input repository, file and folder structures

inside of these repositories might need to be organized for optimal data flow and pipeline triggering. Another resulting requirement of this commit-based triggering system is that each pipeline must output files or objects to an output repository. These triggering features can be very beneficial to scientific workflows, as highlighted with the NEON use-case.

Workflow execution: Pachyderm has the ability to scale over a Kubernetes cluster, which is extremely simple. Pachyderm automatically retries each data pipeline based upon task exit code. However, for certain workflows, error management might not be as easy as other WMS, since actual workflows are buried in containers inside Kubernetes pods.

Data management: Data passing is a strong point of Pachyderm, with data flow and provenance being one of the headline features of the WMS. Provenance also can be tracked back through a workflow, with PFS repositories tracking files and their respective commits.

VI. CONCLUSION

This work's main objective is to describe interesting trends and concepts in next-generation scientific and commercial workflow management systems, from a user's perspective. To this end, we analyzed how two different workflow management paradigms, namely the task-driven and data-driven paradigms, can be applied to real-world use-cases.

From the four generic use-cases detailed in the introduction, we carefully described three real-world use-cases. With a traditional scientific workflow with the 1000 Genome Project workflow, and two different sensor-based workflows with the CASA and NEON soil workflows, we presented a cross-section of traditional and next-generation workflows to evaluate trends and developments in the workflow space.

For the task-driven and data-driven paradigms, we selected representative workflow management systems. Pegasus, Makeflow and Apache Airflow represent the task-driven model, while Pachyderm represents the data-driven model. Each model is thoroughly and holistically evaluated, along with their associated workflow management systems. While performing this evaluation, we examined the rise of new technologies and innovations, including containers and the cloud, and new workflow management use-cases, such as big data analytics, large-scale science and machine learning. Using several real-world use-cases, we highlighted how each WMS's unique features can be an asset to certain next-generation workflows, and emphasized how these features set each WMS apart from one another.

Future work consists of exploring more real-world use-cases such as IoT workflows or large-scale data analytics, as well as more WMS solutions and approaches, e.g. Apache Kafka for real-time data streaming. We would also like to evaluate how techniques developed by these next-generation WMS could benefit to traditional scientific workflows.

Acknowledgments. This work is funded by NSF contract #1842042: "Pilot Study for a Cyberinfrastructure Center of Excellence". The National Ecological Observatory Network is a program sponsored

by the National Science Foundation and operated under cooperative agreement by Battelle Memorial Institute.

REFERENCES

- [1] E. Deelman, T. Peterka *et al.*, "The future of scientific workflows," *The International Journal of High Performance Computing Applications*, 2018.
- [2] C. S. Liew, M. P. Atkinson, M. Galea, T. F. Ang, P. Martin, and J. I. V. Hemert, "Scientific workflows: Moving across paradigms," *ACM Computing Surveys (CSUR)*, vol. 49, no. 4, p. 66, 2017.
- [3] A. Barker and J. Van Hemert, "Scientific workflow: a survey and research directions," in *International Conference on Parallel Processing and Applied Mathematics*. Springer, 2007, pp. 746–753.
- [4] J. Liu, E. Pacitti, P. Valduriez, and M. Mattoso, "A survey of data-intensive scientific workflow management," *Journal of Grid Computing*, vol. 13, no. 4, pp. 457–493, 2015.
- [5] A. Rouhani, E. Bernhardsson, and E. Freider. Extreme science and engineering discovery environment (XSEDE).
- [6] M. Nardelli, S. Nastic, S. Dustdar, M. Villari, and R. Ranjan, "Osmotic flow: Osmotic computing+ iot workflow," *IEEE Cloud Computing*, vol. 4, no. 2, pp. 68–75, 2017.
- [7] R. Sumbaly, J. Kreps, and S. Shah. "The big data ecosystem at linkedin," in *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*. ACM, 2013, pp. 1125–1134.
- [8] M. Beauchemin. (2014) Apache Airflow Project.
- [9] M. Kotliar *et al.*, "CWL-Airflow: a lightweight pipeline manager supporting common workflow language," *bioRxiv*, 2018.
- [10] E. Deelman, K. Vahi *et al.*, "Pegasus: a workflow management system for science automation," *Future Generation Computer Systems*, 2015.
- [11] R. L. Graham *et al.*, "Optimization and approximation in deterministic sequencing and scheduling: a survey," in *Annals of discrete mathematics*. Elsevier, 1979.
- [12] B. Ludäscher, I. Altintas, C. Berkley *et al.*, "Scientific workflow management and the kepler system," *Concurrency and Computation: Practice and Experience*, vol. 18, no. 10, pp. 1039–1065, 2006.
- [13] M. Albrecht *et al.*, "Makeflow: A portable abstraction for data intensive computing on clusters, clouds, and grids," in *Proceedings of the 1st ACM SIGMOD Workshop on Scalable Workflow Execution Engines and Technologies*. ACM, 2012, p. 1.
- [14] T. Oinn, M. Addis, J. Ferris *et al.*, "Taverna: a tool for the composition and enactment of bioinformatics workflows," *Bioinformatics*, vol. 20, no. 17, pp. 3045–3054, 2004.
- [15] Y. N. Babuji *et al.*, "Parsl: Scalable parallel scripting in python." in *IWSG*, 2018.
- [16] J. Dean and S. Ghemawat, "Mapreduce: a flexible data processing tool," *Communications of the ACM*, vol. 53, no. 1, pp. 72–77, 2010.
- [17] V. K. Vavilapalli, A. C. Murthy, C. Douglas *et al.*, "Apache hadoop yarn: Yet another resource negotiator," in *Proceedings of the 4th annual Symposium on Cloud Computing*. ACM, 2013, p. 5.
- [18] H. Pathak, M. Rathi, and A. Parekh, "Introduction to real-time processing in apache apex," *Int. J. Res. Advent Technol.*, p. 19, 2016.
- [19] P. Di Tommaso, M. Chatzou, E. W. Floden, P. P. Barja, E. Palumbo, and C. Notredame, "Nextflow enables reproducible computational workflows," *Nature biotechnology*, vol. 35, no. 4, p. 316, 2017.
- [20] Pachyderm, Inc. (2017) Pachyderm. [Online]. Available: <https://www.pachyderm.io/>
- [21] G. Bosilca *et al.*, "Dague: A generic distributed dag engine for high performance computing," *Parallel Computing*, 2012.
- [22] R. Sethi, "Scheduling graphs on two processors," *SIAM Journal on Computing*, pp. 73–82, 1976.
- [23] D. Bernstein, "Containers and cloud: From lxc to docker to kubernetes," *IEEE Cloud Computing*, vol. 1, no. 3, pp. 81–84, 2014.
- [24] J. A. Novella *et al.*, "Container-based bioinformatics with pachyderm," *Bioinformatics*, vol. 35, no. 5, pp. 839–846, 2018.
- [25] L. Torvalds and J. Hamano. (2005) Git: Fast version control system.
- [26] The 1000 Genomes Project Consortium *et al.*, "A global reference for human genetic variation," *Nature*, 09 2015.
- [27] E. Lyons *et al.*, "Toward a dynamic network-centric distributed cloud platform for scientific workflows: A case study for adaptive weather sensing," in *15th eScience Conference*, 2019.
- [28] D. Thain, T. Tannenbaum, and M. Livny, "Distributed computing in practice: the condor experience," *Concurrency and computation: practice and experience*, vol. 17, no. 2-4, pp. 323–356, 2005.