WRENCH: A Framework for Simulating Workflow Management Systems

Henri Casanova, Suraj Pandey, James Oeth, Ryan Tanaka, Frédéric Suter,
Rafael Ferreira da Silva

► To cite this version:

Henri Casanova, Suraj Pandey, James Oeth, Ryan Tanaka, Frédéric Suter, et al.. WRENCH: A Framework for Simulating Workflow Management Systems. WORKS 2018 - 13th Workshop on Workflows in Support of Large-Scale Science, Nov 2018, Dallas, United States. pp.1-12. hal-01948162

HAL Id: hal-01948162

<https://hal.inria.fr/hal-01948162>

Submitted on 7 Dec 2018

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

WRENCH: A Framework for Simulating Workflow Management Systems

Henri Casanova*, Suraj Pandey*, James Oeth[§], Ryan Tanaka*, Frédéric Suter[‡], Rafael Ferreira da Silva[§]

*Information and Computer Sciences, University of Hawaii, Honolulu, HI, USA

[§]Information Sciences Institute, University of Southern California, Marina Del Rey, CA, USA

[‡]IN2P3 Computing Center, CNRS, Villeurbanne, France

{henric,surajp,ryanyt}@hawaii.edu, {rafsilva,oeth}@isi.edu, frederic.suter@cc.in2p3.fr

Abstract—Scientific workflows are used routinely in numerous scientific domains, and Workflow Management Systems (WMSs) have been developed to orchestrate and optimize workflow executions on distributed platforms. WMSs are complex software systems that interact with complex software infrastructures. Most WMS research and development activities rely on empirical experiments conducted with full-fledged software stacks on actual hardware platforms. Such experiments, however, are limited to hardware and software infrastructures at hand and can be labor- and/or time-intensive. As a result, relying solely on real-world experiments impedes WMS research and development. An alternative is to conduct experiments in simulation.

In this work we present WRENCH, a WMS simulation framework, whose objectives are (i) accurate and scalable simulations; and (ii) easy simulation software development. WRENCH achieves its first objective by building on the SimGrid framework. While SimGrid is recognized for the accuracy and scalability of its simulation models, it only provides low-level simulation abstractions and thus large software development efforts are required when implementing simulators of complex systems. WRENCH thus achieves its second objective by providing high-level and directly re-usable simulation abstractions on top of SimGrid. After describing and giving rationales for WRENCH’s software architecture and APIs, we present a case study in which we apply WRENCH to simulate the Pegasus production WMS. We report on ease of implementation, simulation accuracy, and simulation scalability so as to determine to which extent WRENCH achieves its two above objectives. We also draw both qualitative and quantitative comparisons with a previously proposed workflow simulator.

Index Terms—Scientific Workflows, Workflow Management Systems, Simulation, Distributed Computing

I. INTRODUCTION

Scientific workflows have become mainstream in support of research and development activities in numerous scientific domains [1]. Consequently, several Workflow Management Systems (WMSs) have been developed [2]–[7] that allow scientists to execute workflows on distributed platforms that can accommodate executions at various scales. WMSs handle the logistics of workflow executions and make decisions regarding resource selection, data management, and computation scheduling, the goal being to optimize some performance metric (e.g., latency [8], [9], throughput [10], [11], jitter [12], reliability [13]–[15], power consumption [16], [17]). WMSs are complex software systems that interact with complex software infrastructures and can thus employ a wide range of designs and algorithms.

In spite of active WMS development and use in production, which has entailed solving engineering challenges, fundamental questions remain unanswered in terms of system designs and algorithms. Although there are theoretical underpinnings for most of these questions, theoretical results often make assumptions that do not hold with production hardware and software infrastructures. Further, the specifics of the design of a WMS can impose particular constraints on what solutions can be implemented effectively, and these constraints are typically not considered in available theoretical results. Consequently, current research that aims at improving and evolving the state of the art, although sometimes informed by theory, is mostly done via “real-world” experiments: designs and algorithms are implemented, evaluated, and selected based on experiments conducted for a particular WMS implementation with particular workflow configurations on particular platforms. As a corollary, from the WMS user’s perspective, quantifying accurately how a WMS would perform for a particular workflow configuration on a particular platform entails actually executing that workflow on that platform.

Unfortunately, real-world experiments have limited scope, which impedes WMS research and development. This is because they are confined to application and platform configurations available at hand, and thus cover only a small subset of the relevant scenarios that may be encountered in practice. Furthermore, exclusively relying on real-world experiments makes it difficult or even impossible to investigate hypothetical scenarios (e.g., “What if the network had a different topology?”, “What if there were 10 times more compute nodes but they had half as many cores?”). Real-world experiments, especially when large-scale, are often not fully reproducible due to shared networks and compute resources, and due to transient or idiosyncratic behaviors (maintenance schedules, software upgrades, and particular software (mis)configurations). Running real-world experiments is also time-consuming, thus possibly making it difficult to obtain statistically significant numbers of experimental results. Real-world experiments are driven by WMS implementations that often impose constraints on workflow executions. Furthermore, WMSs are typically not monolithic but instead reuse CyberInfrastructure (CI) components that impose their own overheads and constraints on workflow execution. Exploring what lies beyond these constraints via real-world executions,

e.g., for research and development purposes, typically entails unacceptable software (re-)engineering costs. Finally, running real-world experiments can also be labor-intensive. This is due to the need to install and execute many full-featured software stacks, including actual scientific workflow implementations, which is often not deemed worthwhile for “just testing out” ideas.

An alternative to conducting WMS research via real-world experiments is to use *simulation*, i.e., implement a software artifact that models the functional and performance behaviors of software and hardware stacks of interest. Simulation is used in many computer science domains and can address the limitations of real-world experiments outlined above. Several simulation frameworks have been developed that target the parallel and distributed computing domain [18]–[34]. Some simulation frameworks have also been developed specifically for the scientific workflow domain [11], [35]–[40].

We claim that advances in simulation capabilities in the field have made it possible to simulate WMSs that execute large workflows on large-scale platforms accessible via diverse CI services in a way that is accurate (via validated simulation models), scalable (fast execution and low memory footprint), and expressive (ability to describe arbitrary platforms, complex WMSs, and complex software infrastructure). In this work, we build on the existing open-source SimGrid simulation framework [33], [41], which has been one of the drivers of the above advances and whose simulation models have been extensively validated [42]–[46], to develop a WMS simulation framework called WRENCH [47]. More specifically, this work makes the following contributions:

- 1) We justify the need for WRENCH and explain how it improves on the state of the art.
- 2) We describe the high-level simulation abstractions provided by WRENCH that (i) make it straightforward to implement full-fledged simulated versions of complex WMS systems; and (ii) make it possible to instantiate simulation scenarios with only few lines of code.
- 3) Via a case study with the Pegasus [2] production WMS, we evaluate the ease-of-use, accuracy, and scalability of WRENCH, and compare it with a previously proposed simulator, WorkflowSim [35].

This paper is organized as follows. Section II discusses related work. Section III outlines the design of WRENCH and describes how its APIs are used to implement simulators. Section IV presents our case study. Finally, Section V concludes with a brief summary of results and a discussion of future research directions.

II. RELATED WORK

Many simulation frameworks have been developed for parallel and distributed computing research and development. They span domains such as HPC [18]–[21], Grid [22]–[24], Cloud [25]–[27], Peer-to-peer [28], [29], or Volunteer Computing [30]–[32]. Some frameworks have striven to be applicable across some or all of the above domains [33], [34]. Two conflicting concerns are *accuracy* (the ability to

capture the behavior of a real-world system with as little bias as possible) and *scalability* (the ability to simulate large systems with as few CPU cycles and bytes of RAM as possible). The aforementioned simulation frameworks achieve different compromises between these two concerns by using various simulation models. At one extreme are discrete event models that simulate the “microscopic” behavior of hardware/software systems (e.g., by relying on packet-level network simulation for communication [48], on cycle-accurate CPU simulation [49] or emulation for computation). In this case, the scalability challenge can be handled by using Parallel Discrete Event Simulation [50], i.e., the simulation itself is a parallel application that requires a parallel platform whose scale is at least commensurate to that of the simulated platform. At the other extreme are analytical models that capture “macroscopic” behaviors (e.g., transfer times as data sizes divided by bottleneck bandwidths, compute times as numbers of operations divided by compute speeds). While these models are typically more scalable, they must be developed with care so that they are accurate. In previous work, it has been shown that several available simulation frameworks use macroscopic models that can exhibit high inaccuracy [43].

A number of simulators have been developed that target scientific workflows. Some of them are stand-alone simulators [11], [35]–[37]. Others are integrated with a particular WMS to promote more faithful simulation and code reuse [38], [39] or to execute simulations at runtime to guide on-line scheduling decisions made by the WMS [40].

The authors in [39] conduct a critical analysis of the state-of-the-art of workflow simulators. They observe that many of these simulators do not capture the details of underlying infrastructures and/or use naive simulation models. This is the case with custom simulators such as that in [36], [37], [40]. But it is also the case with workflow simulators built on top of generic simulation frameworks that provide convenient user-level abstractions but fail to model the details of the underlying infrastructure, e.g., the simulators in [11], [35], [38], which build on the CloudSim [25] or GroudSim [24] frameworks. These frameworks have been shown to lack in their network modeling capabilities [43]. As a result, some authors readily recognize that their simulators are likely only valid when network effects play a small role in workflow executions (i.e., when workflows are not data-intensive).

To overcome the above limitations, in [39] the authors have improved the network model in GroudSim and also use a separate simulator, DISSECT-CF [27], for simulating cloud infrastructures accurately. Both [39] and [27] acknowledge that the popular SimGrid [33], [41] simulation framework offers compelling capabilities, both in terms of scalability and simulation accuracy. But one of their reasons for not considering SimGrid is that, because it is low-level, using it to implement a simulator of a complex system, such as a WMS and the CI services it uses, would be too labor-intensive. In this work, we address this issue by developing a simulation framework that provides convenient, reusable, high-level abstractions but that builds on SimGrid so as to

benefit from its scalable and accurate simulation models. Furthermore, unlike [38], [39], we do not focus on integration with any specific WMS. The argument in [39] is that stand-alone simulators, such as that in [35], are disconnected from real-world WMSs because they abstract away much of the complexity of these systems. Instead, our proposed framework does capture low-level system details (and simulates them well thanks to SimGrid), but provides high-level enough abstractions to implement faithful simulations of complex WMSs with minimum effort, which we demonstrate via a case study with the Pegasus WMS [2].

Also related to this work is previous research that has not focused on providing simulators or simulation frameworks *per se*, but instead on WMS simulation methodology. In particular, several authors have investigated methods for injecting realistic stochastic noise in simulated WMS executions [35], [51]. These techniques can be adopted by most of the aforementioned frameworks, including the one proposed in this work.

III. WRENCH

A. Objective and Intended Users

WRENCH’s objective is to make it possible to study WMSs in simulation in a way that is accurate (faithful modeling of real-world executions), scalable (low computation and memory footprints on a single computer), and expressive (ability to simulate arbitrary WMS, workflow, and platform scenarios with minimal software engineering effort). WRENCH is not a simulator but a simulation framework that is distributed as a C++ library. It provides high-level reusable abstractions for developing simulated WMS implementations and simulators for the execution of these implementations. There are two categories of WRENCH users:

1. *Users who implement simulated WMSs* – These users are engaged in WMS research and development activities and need an “in simulation” version of their current or intended WMS. Their goals typically include evaluating how their WMS behaves over hypothetical experimental scenarios and comparing competing algorithm and system design options. For these users, WRENCH provides the WRENCH Developer API (described in Section III-D) that eases WMS development by removing the typical difficulties involved when developing, either in real-world or in simulation mode, a system comprised of distributed components that interact both synchronously and asynchronously. To this end, WRENCH makes it possible to implement a WMS as a single thread of control that interacts with simulated CI services via high-level APIs and must react to a small set of asynchronous events.
2. *Users who execute simulated WMSs* – These users simulate how given WMSs behave for particular workflows on particular platforms. Their goals include comparing different WMSs, determining how a given WMS would behave for various workflow configurations, comparing different platform and resource provisioning options, determining performance bottlenecks, engaging in pedagogic activities centered on distributed computing and

Fig. 1: The four layers in the WRENCH architecture from bottom to top: simulation core, simulated core services, simulated WMS implementations, and simulators.

workflow issues, etc. These users can develop simulators via the WRENCH User API (described in Section III-E), which makes it possible to build a full-fledged simulator with only a few lines of code.

Users in the first category above often also belong to the second category. That is, after implementing a simulated WMS these users typically instantiate simulators for several experimental scenarios to evaluate their WMS.

B. Software Architecture Overview

Figure 1 depicts WRENCH’s software architecture. At the bottom layer is the Simulation Core, which simulates low-level software and hardware stacks using the simulation abstractions and models provided by SimGrid (see Section III-C). The next layer implements simulated CI services that are commonly found in current distributed platforms and used by production WMSs. At the time of this writing, WRENCH provides services in 4 categories: *compute services* that provide access to compute resources to execute workflow tasks; *storage services* that provide access to storage resources for storing workflow data; *network monitoring services* that can be queried to determine network distances; and *data registry services* that can be used to track the location of (replicas of) workflow data. Each category includes multiple service implementations, so as to capture specifics of currently available CI services used in production. For instance, in its current version WRENCH provides a “batch-scheduled cluster” compute service, a a “cloud” compute service, and a “bare-metal” compute service. The above layer in the software architecture consists of simulated

WMS, that interact with CI services using the WRENCH Developer API (see Section III-D). These WMS implementations, which can simulate production WMSs or WMS research prototypes, are not included as part of the WRENCH distribution, but implemented as stand-alone projects. One such project is the simulated Pegasus implementation for our case study in Section IV. Finally, the top layer consists of simulators that configure and instantiate particular CI services and particular WMSs on a given simulated hardware platform, that launch the simulation, and that analyze the simulation outcome. These simulators use the WRENCH User API (see Section III-E). Here again, these simulators are not part of WRENCH, but implemented as stand-alone projects.

C. Simulation Core

WRENCH’s simulation core is implemented using SimGrid’s S4U API, which provides all necessary abstractions and models to simulate computation, I/O, and communication activities on arbitrary hardware platform configurations. These platform configurations are defined by XML files that specify network topologies and endpoints, compute resources, and storage resources [52].

At its most fundamental level, SimGrid provides a Concurrent Sequential Processes (CSP) model: a simulation consists of sequential threads of control that consume hardware resources. These threads of control can implement arbitrary code, exchange messages via a simulated network, can perform computation on simulated (multicore) hosts, and can perform I/O on simulated storage devices. In addition, SimGrid provides a virtual machine abstraction that includes a migration feature. Therefore, SimGrid provides all the base abstractions necessary to implement the classes of distributed systems that are relevant to scientific workflow executions. However, these abstractions are low-level and a common criticism of SimGrid is that implementing a simulation of a complex system requires a large software engineering effort. A WMS executing a workflow using several CI services is a complex system, and WRENCH builds on top of SimGrid to provide high-level abstractions so that implementing this complex system is not labor-intensive.

We have selected SimGrid for WRENCH for the following reasons. SimGrid has been used successfully in many distributed computing domains (cluster, peer-to-peer, grid, cloud, volunteer computing, etc.), and thus can be used to simulate WMSs that execute over a wide range of platforms. SimGrid is open source and freely available, has been stable for many years, is actively developed, has a sizable user community, and has provided simulation results for over 350 research publications since its inception. SimGrid has also been the object of many invalidation and validation studies [42]–[46], and its simulation models have been shown to provide compelling advantages over other simulation frameworks in terms of both accuracy and scalability [33]. Finally, most SimGrid simulations can be executed in minutes on a standard laptop computer, making it possible to perform large numbers of simulations quickly with minimal compute resource expenses.

Algorithm 1 Blueprint for a WMS execution

```

1: procedure MAIN(workflow)
2:   Obtain list of available services
3:   Gather static information about the services
4:   while workflow execution has not completed/failed do
5:     Gather dynamic service/resource information
6:     Submit pilot jobs if needed
7:     Make data/computation scheduling decisions
8:     Interact with services to enact decisions
9:     Wait for and react to the next event
10:  end while
11:  return
12: end procedure

```

To the best of our knowledge, among comparable available simulation frameworks (as reviewed in Section II), SimGrid is the only one to offer all the above desirable characteristics.

D. WRENCH Developer API

With the Developer API, a WMS is implemented as a single thread of control that executes according to the pseudo-code blueprint shown in Algorithm 1. Given a workflow to execute, a WMS first gathers information about all the CI services it can use to execute the workflow (lines 2-3). Examples of such information include the number of compute nodes provided by a compute service, the number of cores per node and the speed of these cores, the amount of storage space available in a storage service, the list of hosts monitored by a network monitoring service, etc. Then, the WMS iterates until the workflow execution is complete or has failed (line 4). At each iteration it gathers dynamic information about available services and resources if needed (line 5). Example of such information include currently available capacities at compute or storage services, current network distances between pairs of hosts, etc. Then, if desired, the WMS can submit *pilot jobs* [53] to compute services that support them, if any (line 6). Based on resource information and on the current state of the workflow, the WMS can then make whatever scheduling decisions it sees fit (line 7). It then enacts these decisions by interacting with appropriate services. For instance, it could decide to submit a “job” to a compute service to execute a ready task on some number of cores at some compute service and copy all produced files to some storage service, or it could decide to just copy a file between storage services and then update a data location service to keep track of the location of this new file replica. It is the responsibility of the developer to implement all decision-making algorithms employed by the WMS. At the end of the iteration, the WMS simply waits for a workflow execution event to which it can react if need be. Most common events are job completions/failures and data transfer completions/failures.

The WRENCH Developer API provides a rich set of methods to process analyze the workflow and to interact with CI services to execute the workflow. These methods were designed based on current and envisioned capabilities of current state-of-the-art WMSs. We refer the reader to the WRENCH Web site [47] for more information on how to use this API

and for the full API documentation. The key objective of this API is to make it straightforward to implement a complex system, namely a full-fledged WMS that interact with diverse CI services. We achieve this objective by providing simple solutions and abstractions to handle well-known challenges when implementing a complex distributed system (whether in the real world or in simulation), as explained hereafter.

SimGrid provides simple point-to-point communication between threads of control via a mailbox abstraction. One of the recognized strengths of SimGrid is that it employs highly accurate and yet scalable network simulation models. However, unlike some of its competitors, it does not provide any higher-level simulation abstractions meaning that distributed systems must be implemented essentially from scratch, with many message-based interactions. All message-based communication is abstracted away by WRENCH, and although the simulated CI services exchange many messages with the WMS and among themselves, the WRENCH Developer API only exposes higher-level interaction with services (“run this job”, “move this data”) and only requires that the WMS handle a few events. The WMS developer thus completely avoids the need to send and receive (and thus orchestrate) network messages.

Another challenge when developing a system like a WMS is the need to handle asynchronous interactions. While some service interactions can be synchronous (e.g., “are you up?”, “tell me your current load”), most need to be asynchronous so that the WMS retains control. The typical solution is to maintain sets of request handles and/or to use multiple threads of control. To free the WMS developer from these responsibilities, WRENCH provides already implemented “managers” that can be used out-of-the-box to take care of asynchronicity. A WMS can instantiate such managers, which are independent threads of control. Each manager transparently interacts with CI services, maintains a database of pending requests, provides a simple API to check on the status of these requests, and automatically generates workflow execution events. For instance, a WMS can instantiate a “job manager” through which it will create and submit jobs to compute services. It can at any time check on the status of a job, and the job manager interacts directly (and asynchronously) with compute services so as to generate “job done” or “job failed” events to which the WMS can react. In our experience developing simulators from scratch using SimGrid, the implementation of asynchronous interactions with simulated processes is a non-trivial development effort, both in terms of amount of code to write and difficulty to write this code correctly. We posit that this is one of the reasons why some users have preferred using simulation frameworks that provide higher-level abstractions than SimGrid but offer less attractive accuracy and/or scalability features. WRENCH provides such higher-level abstractions to the WMS developers, and as a result implementing a WMS with WRENCH can be straightforward.

Finally, one of the challenges when developing a WMS is failure handling. It is expected that compute, storage, and network resources, as well as the CI services that use them, can fail through the execution of the WMS. SimGrid has the

capability to simulate arbitrary failures via availability traces. Furthermore, failures can occur due to the WMS implementation itself, e.g., if it fails to check that the operations it attempts are actually valid, if concurrent operations initiated by the WMS work at cross purposes. WRENCH abstracts away all these failures as C++ exceptions that can be caught by the WMS implementation, or caught by a manager and passed to the WMS as workflow execution events. Regardless, each failure exposes a failure cause, which encodes a detailed description of the failure. For instance, after initiating a file copy from a storage service to another storage service, a “file copy failed” event sent to the WMS would include a failure cause that could specify that when trying to copy file x from storage service y to storage service z , storage service z did not have sufficient storage space. Other example failure causes could be that a network error occurred when storage service y attempted to receive a message from storage service z , or that service z was down. All CI services implemented in WRENCH simulate well-defined failure behaviors, and failure handling capabilities afforded to simulated WMSs can actually allow more sophisticated failure tolerance strategies than currently done or possible in real-world implementations. But more importantly, the amount of code that needs to be written for failure handling in a simulated WMS is minimal.

Given the above, WRENCH makes it possible to implement a simulated WMS with very little code and effort. The example WMS implementation provided with the WRENCH distribution, which is simple but functional, is under 200 lines of C++ (once comments have been removed). See more discussion of the effort needed to implement a WMS with WRENCH in the context of our Pegasus case study (Section IV).

E. WRENCH User API

With the User API one can quickly build a simulator, which typically follows these steps:

1. Instantiate a platform based on a SimGrid XML platform description file;
2. Create one or more workflows;
3. Instantiate services on the platform;
4. Instantiate one or more WMSs telling each what services are at its disposal and what workflow it should execute starting at what time;
5. Launch the simulation; and
6. Process the simulation outcome.

The above steps can be implemented with only a few lines of C++. An example WRENCH simulator is shown in Figure 2, which uses a WMS implementation (called `SomeWMS`) that has already been developed using the WRENCH Developer API (see previous section). After initializing the simulation (lines 5-6), the simulator instantiates a platform (line 8) and a workflow (line 10-11). A workflow is defined as a set of computation tasks and data files, with control and data dependencies between tasks. Each task can also have a priority, which can then be taken into account by a WMS for scheduling purposes. Although the workflow can be defined purely programmatically, in this example the workflow is imported from

```

1  #include <math.h>
2  #include <wrench.h>
3  int main(int argc, char **argv) {
4      // Declare and initialize a simulation
5      wrench::Simulation simulation;
6      simulation.init(&argc, argv);
7      // Instantiate a platform
8      simulation.instantiatePlatform("my_platform.xml");
9      // Instantiate a workflow
10     wrench::Workflow workflow;
11     workflow.loadFromDAX("my_workflow.dax", "1000Gf");
12     // Instantiate a storage service
13     auto storage_service = simulation.add(
14         new wrench::SimpleStorageService("storage_host", pow(2,50)));
15     // Instantiate a sompute service (a batch-scheduled 4-node cluster that uses the
16     // EASY backfilling algorithm and is subject to a background load)
17     auto batch_service = simulation.add(
18         new wrench::BatchService("batch_login", {"node1", "node2", "node3", "node4"}, pow(2,40),
19             {{wrench::BatchServiceProperty::SIMULATED_WORKLOAD_TRACE_FILE, "load.swf"},
20              {wrench::BatchServiceProperty::BATCH_SCHEDULING_ALGORITHM, "easy_bf"}}));
21     // Instantiate a compute service (a 4-host cloud platform that does not support pilot jobs)
22     auto cloud_service = simulation.add(
23         new wrench::CloudService("cloud_gateway", {"host1", "host2", "host3", "host4"}, pow(2,42),
24             {{wrench::CloudServiceProperty::SUPPORTS_PILOT_JOBS, "false"}}));
25     // Instantiate a data registry service
26     auto data_registry_service = simulation.add(new wrench::FileRegistryService("my_desktop"));
27     // Instantiate a network monitoring service
28     auto network_monitoring_service =
29         simulation.add(new wrench::NetworkProximityService(
30             "my_desktop", {"my_desktop", "batch_login", "cloud_gateway"},
31             {{wrench::NetworkProximityServiceProperty::NETWORK_PROXIMITY_SERVICE_TYPE,
32              "vivaldi"}}));
33     // Stage a workflow input file at the storage service
34     simulation.stageFile(workflow.getFileByID("input_file"), storage_service);
35     // Instantiate a WMS...
36     auto wms = simulation.add(
37         new wrench::SomeWMS({batch_service, cloud_service}, {storage_service,
38             {network_monitoring_service}, {data_registry_service}, "my_desktop"));
39     // ... and assign the workflow to it, to be executed one hour in
40     wms->addWorkflow(&workflow, 3600);
41     // Launch the simulation
42     simulation.launch();
43     // Retrieve task completion events
44     auto trace = simulation.getOutput().getTrace<wrench::SimulationTimestampTaskCompletion>();
45     // Determine the completion time of the last task that completed
46     double completion_time = trace[trace.size()-1]->getContent()->getDate();
47 }

```

Fig. 2: Example fully functional WRENCH simulator. Try-catch clauses are omitted.

a workflow description file in the DAX format [54]. At line 13 the simulator creates a storage service with 1PiB capacity accessible on host `storage_host`. This and other hostnames are specified in the XML platform description file. At line 17 the simulator creates a compute service that corresponds to a 4-node batch-scheduled cluster. The physical characteristics of the compute nodes (`node[1-4]`) are specified in the platform description file. This compute service has a 1TiB scratch storage space. Its behavior is customized by passing a couple of property-value pairs to its constructor. It will be subject to a background load as defined by a trace in the standard SWF format [55], and its batch queue will be managed using the EASY Backfilling scheduling algorithm [56]. The simulator then creates a second compute service (line 22), which is a 4-host cloud service, customized so that it does not support pilot jobs. Two helper services are instantiated, a data registry service so that the WMS can keep track of file locations

(line 26) and a network monitoring service that uses the Vivaldi algorithm [57] to measure network distances between the two hosts from which the compute services are accessed (`batch_login` and `cloud_gateway`) and the `my_host` host, which is the host that runs these helper services and the WMS (line 28). At line 34, the simulator specifies that the workflow data file `input_file` is initially available at the storage service. It then instantiates the WMS and passes to it all available services (line 36), and assigns the workflow to it (line 39). The crucial call is at line 41, where the simulation is launched and the simulator hands off control to WRENCH. When this call returns the workflow has either completed or failed. Assuming it has completed, the simulator then retrieves the ordered set of task completion events (line 43) and performs some (in this example, trivial) mining of these events (line 45).

For brevity, the example in Figure 2 omits `try/catch`

clauses. Also, note that although the simulator uses the `new` operator to instantiate WRENCH objects, the simulation object takes ownership of these objects (using unique or shared pointers), so that there is no memory deallocation onus placed on the user. This example showcases only the most fundamental features of the WRENCH User API, and we refer the reader to the WRENCH Web site [47] for more detailed information on how to use this API and for the full API documentation. In the future this API will come with Python binding so that users can implement simulators in Python.

IV. CASE STUDY: SIMULATING A PRODUCTION WMS

In this section, we present a WRENCH-based simulator of a state-of-the-art WMS, Pegasus [2], as a case study for evaluation and validation purposes.

Pegasus is being used in production to execute workflows for dozens of high-profile applications in a wide range of scientific domains [2]. Pegasus provides the necessary abstractions for scientists to create workflows and allows for transparent execution of these workflows on a range of compute platforms including clusters, clouds, and national cyberinfrastructures. During execution, Pegasus translates an abstract resource-independent workflow into an executable workflow, determining the specific executables, data, and computational resources required for the execution. Workflow execution with Pegasus includes data management, monitoring, and failure handling, and is managed by HTCondor DAGMan [58]. Individual workflow tasks are managed by a workload management framework, HTCondor [59], which supervises task executions on local and remote resources.

A. Implementing Pegasus with WRENCH

Since Pegasus relies on HTCondor, first we have implemented the HTCondor services as simulated core CI services, which together form a new Compute Service that exposes the WRENCH Developer API. This makes HTCondor available to any WMS implementation that is to be simulated using WRENCH, and will be included in the next WRENCH release as part of the growing set of simulated core CI services provided by WRENCH.

HTCondor is composed of six main service daemons (`startd`, `starter`, `schedd`, `shadow`, `negotiator`, and `collector`). In addition, each host on which one or more of these daemons is spawned must also run a `master` daemon, which controls the execution of all other daemons (including initialization and completion). The bottom part of Figure 3 depicts the components of our simulated HTCondor implementation, where daemons are shown in red-bordered boxes. In our simulator we implement the 3 fundamental HTCondor services, implemented as particular sets of daemons, as depicted in the bottom part of Figure 3 in borderless white boxes. The *Job Execution Service* consists of a `startd` daemon, which adds the host on which it is running to the HTCondor pool, and of a `starter` daemon, which manages task executions on this host. The *Central Manager Service* consists of a `collector` daemon, which collects information

Fig. 3: Overview of the WRENCH Pegasus simulation components, including components for DAGMan and HTCondor frameworks. Red boxes denote Pegasus services developed with WRENCH’s Developer API, and white boxes denote WRENCH reused components.

about all other daemons, and of a `negotiator` daemon, which performs task/resource matchmaking. The *Job Submission Service* consists of a `schedd` daemon, which maintains a queue of tasks, and of several instances of a `shadow` daemon, each of which corresponds to a task submitted to the Condor pool for execution.

Given the simulated HTCondor implementation above, we then implemented the simulated Pegasus WMS, including the DAGMan workflow engine, using the WRENCH Developer API. This implementation instantiates all services and parses the workflow description file, the platform description file, and a Pegasus-specific configuration file. DAGMan orchestrates the workflow execution (e.g., a task is marked as ready for execution once all its parent tasks have successfully completed), and monitors the status of tasks submitted to the HTCondor pool using a pull model, i.e., task status is fetched from the pool at regular time intervals. The top part of Figure 3 depicts the components of our simulated Pegasus implementation (each shown in a red box).

By leveraging WRENCH’s high-level simulation abstractions, implementing HTCondor as a reusable core WRENCH service using the Developer API required only 613 lines of code. Similarly, implementing a simulated version of Pegasus, including DAGMan, was done with only 666 lines of code (127 of which are merely parsing simulation configuration files). These numbers include both header and source files, but exclude comments. We argue that the above corresponds to minor simulation software development efforts when considering the complexity of the system being simulated.

Service implementations in WRENCH are all parameterizable. For instance, as services use message-based communications it is possible to specify all message payloads in bytes (e.g., for control messages). Other parameters encompass various overheads, either in seconds or in computation volumes (e.g., task startup overhead on a compute service). In WRENCH, all service implementations come with default

Experimental Scenario Workflow	Platform	Avg. Makespan Error (%)	Task Submissions		Tasks completions	
			<i>p</i> -value	distance	<i>p</i> -value	distance
1000Genome	ExoGENI	1.10 $\pm$ 0.28	0.06 $\pm$ 0.01	0.21 $\pm$ 0.04	0.72 $\pm$ 0.06	0.12 $\pm$ 0.01
Montage-1.5	AWS-t2.xlarge	4.25 $\pm$ 1.16	0.08 $\pm$ 0.01	0.16 $\pm$ 0.03	0.12 $\pm$ 0.05	0.21 $\pm$ 0.02
Montage-2.0	AWS-m5.xlarge	3.37 $\pm$ 0.46	0.11 $\pm$ 0.03	0.06 $\pm$ 0.02	0.10 $\pm$ 0.01	0.11 $\pm$ 0.01

TABLE I: Average simulated makespan error (%), and *p*-values and Kolmogorov-Smirnov (KS) distances for task submission and completion dates, computed for 5 runs of each of our 3 experimental scenarios.

values for all these parameters, but it is possible to pick custom values upon service instantiation. The process of picking parameter values so as to match a specific real-world system is referred to as *simulation calibration*. We calibrated our simulator by measuring delays observed in event traces of real-world executions for workflows on hardware/software infrastructures (see Section IV-B).

The simulator code, details on the simulation calibration procedure, and experimental scenarios used in the rest of this section are all publicly available online [60].

B. Experimental Scenarios

We consider experimental scenarios defined by particular workflow instances to be executed on particular platforms. Due to the lack of publicly available detailed workflow execution traces (i.e., execution logs that include data sizes for all files, all execution delays, etc.), we have performed real workflow executions with Pegasus and collected raw, time-stamped event traces from these executions. These traces form the ground truth to which we can compare simulated executions. We consider these workflow applications:

- *1000Genome* [61]: A data-intensive workflow that identifies mutational overlaps using data from the 1000 genomes project in order to provide a null distribution for rigorous statistical evaluation of potential disease-related mutations. We consider a 1000Genome instance that comprises 71 tasks.
- *Montage* [2]: A compute-intensive astronomy workflow for generating custom mosaics of the sky. For this experiment, we ran Montage for processing 1.5 and 2.0 square degrees mosaic 2MASS. We thus refer to each configuration as Montage-1.5 and Montage-2.0, respectively. Montage-1.5, resp. Montage-2.0, comprises 573, resp. 1,240, tasks.

We use these platforms, deploying on each a submit node (which runs Pegasus, DAGMan, and HTCondor’s job submission and central manager services), four worker nodes (4 cores per node / shared file system), and a data node in the WAN:

- *ExoGENI*: A widely distributed networked infrastructure-as-a-service testbed representative of a “bare metal” platform. Each worker node is a 4-core 2.0GHz processor with 12GiB of RAM. The bandwidth between the data node and the submit node was $\sim$ 0.40 Gbps, and the bandwidth between the submit and worker nodes was $\sim$ 1.00 Gbps.
- *AWS*: Amazon’s cloud platform, on which we use two types of virtual machine instances: `t2.xlarge` and

`m5.xlarge`. The bandwidth between the data node and the submit node was $\sim$ 0.44 Gbps, and the bandwidth between the submit and worker nodes on these instances were $\sim$ 0.74 Gbps and $\sim$ 1.24 Gbps, respectively.

C. Simulation Accuracy

To evaluate the accuracy of our simulator, we consider 3 particular experimental scenarios: 1000Genome on ExoGENI, Montage-1.5 on AWS-t2.xlarge, and Montage-2.0 on AWS-m5.xlarge. Each execution is repeated 5 times and the overall workflow execution times, or makespans, are recorded.

The third column in Table I shows average relative differences between actual and simulated makespans. We see that simulated makespans are close to actual makespans across the board (average relative error is below 5%). One of the key advantages of building WRENCH on top of SimGrid is that WRENCH simulators benefit from the high-accuracy network models in SimGrid, e.g., these models capture many features of the TCP protocol. And indeed, when comparing real-world and simulated executions we observe average relative error below 3% for data movement operations. The many processes involved in a workflow execution with Pegasus interact by exchanging (typically small) control messages. Our simulator simulates these interactions. For instance, each time an output file is produced by a task a data registry service is contacted so that a new entry can be added to its database of file replicas, which incurs some overhead due to a message exchange. When comparing real-world to simulated executions we observe average relative simulation error below 1% for these data registration overheads.

To draw comparisons with a state-of-the-art simulator, we repeated the above simulations using WorkflowSim [35]. WorkflowSim does not provide a detailed simulated HTCondor implementation, does not offer the same simulation calibration capabilities as WRENCH, and is built on top of the CloudSim simulation framework [25]. Nevertheless, we have painstakingly calibrated our WorkflowSim simulator so that it models the hardware and software infrastructures of our experimental scenarios as closely as possible. For each of the 3 experimental scenarios, we find that the relative average makespan percentage error is 12.09 ± 2.84 , 26.87 ± 6.26 , and 13.32 ± 1.12 , respectively, i.e., from 4x up to 11x larger than the error values obtained with our WRENCH-based simulator. The reasons for the discrepancies between WorkflowSim and real-world results are twofold. First, WorkflowSim uses the simplistic network models in CloudSim (see discussion in Section II) and thus suffers from simulation bias w.r.t. data

transfer times. Second, WorkflowSim does not capture all the relevant details of the system and its execution. By contrast, implementing a fully detailed simulator with WRENCH can be done in a few hundred lines of code.

In our experiments we also record the submission and completion dates of each task, thus obtaining empirical cumulative density functions (ECDFs) of these times, for both real-world executions and simulated executions. To further validate the accuracy of our simulation results we apply Kolmogorov-Smirnov goodness of fit tests (KS tests) with null hypotheses (H_0) that the real-world and simulation samples are drawn from the same distributions. The two-sample KS test results in a miss if the null hypothesis (two-sided alternative hypothesis) is rejected at 5% significance level (p -value ≤ 0.05). Each test for which the null hypothesis is not rejected (p -value > 0.05), indicates that the simulated execution statistically matches the real-world execution. Table I shows p -value and KS test distance for both task submission times and task completion times. The null hypothesis is not rejected, and we thus conclude that simulated workflow task executions statistically match real-world executions well. These conclusions are confirmed by visually comparing ECDFs. For instance, Figure 4 shows real-world and simulated ECDFs for sample runs of Montage-2.0 on AWS-m5.xlarge, with task submission, resp. completion, date ECDFs on the left-hand, resp. right-hand, side. We observe that the simulated ECDFs (“wrench”) track the real-world ECDFs (“pegasus”) closely. We repeated these simulations using WorkflowSim, and found that the null hypothesis is rejected for all 3 simulation scenarios. This is confirmed visually in Figure 4, where the ECDFs obtained from the WorkflowSim simulation (“workflowsim”) are far from the real-world ECDFs.

Although KS tests and ECDFs visual inspections validate that the WRENCH-simulated ECDFs match the real-world ECDFs statistically, these results do not distinguish between individual tasks. In fact, there are some discrepancies between real-world and simulated schedules. For instance, Figure 5 shows Gantt charts corresponding to the workflow executions shown in Figure 4, with the real-world execution on the left-hand side (“pegasus”) and the simulated execution on the right-hand side (“wrench”). Tasks executions are shown on the vertical axis, each shown as a line segment along the horizontal time axis, spanning the time between the task’s start time and the task’s finish time. Different task types, i.e., different executables, are shown with different colors. In this workflow, all tasks of the same type are independent and have the same priority. We see that the shapes of the yellow regions, for example, vary between the two executions. These variations are explained by implementation-dependent behaviors of the workflow scheduler. In many instances throughout workflow execution several ready tasks can be selected for execution, e.g., sets of independent tasks in the same level of the workflow. When the number of available compute resources, n , is smaller than the number of ready tasks, the scheduler picks n ready tasks for immediate execution. In most WMSs, these tasks are picked as whatever first n tasks

Fig. 4: Empirical cumulative distribution function of task submit times (left) and task completion times (right) for sample real-world (“pegasus”) and simulated (“wrench” and “workflowsim”) executions of Montage-2.0 on AWS-m5.xlarge.

Fig. 5: Task execution Gantt chart for sample real-world (“pegasus”) and simulated (“wrench”) executions of the Montage-2.0 workflow on the AWS-m5.xlarge platform.

are returned when iterating over data structures in which task objects are stored. Building a perfectly faithful simulation of a WMS would thus entail implementing/using the exact same data structures as that in the actual implementation. This could be labor intensive or perhaps not even possible depending on which data structures, languages, and/or libraries are used in that implementation. In the context of this Pegasus case study, the production implementation of the DAGMan scheduler uses a custom priority list implementation to store ready tasks, while our simulation version of it stores workflow tasks in a `std::map` data structure indexed by task string IDs. Consequently, when the real-world scheduler picks the first n ready tasks it typically picks different tasks than those picked by its simulated implementation. This is the cause the discrepancies seen in Figure 5.

D. Simulation Scalability

Table II shows average simulated makespans and simulation execution times for our 3 experimental scenarios. Simulations are executed on a single core of a MacBook Pro 3.5 GHz Intel

Experimental Workflow	Scenario Platform	Avg. Workflow Makespan (s)	Avg. Simulation Time (s)
1000Genome	ExoGENI	761.0 $\pm$ 7.93	0.3 $\pm$ 0.01
Montage-1.5	AWS-t2.xlarge	1,784.0 $\pm$ 137.67	8.3 $\pm$ 0.09
Montage-2.0	AWS-m5.xlarge	2,911.8 $\pm$ 48.80	28.1 $\pm$ 0.52

TABLE II: Simulated workflow makespans and simulation times averaged over 5 runs of each of our 3 experimental scenarios.

Fig. 6: Average simulation time (in seconds, left vertical axis) and memory usage (maximum resident set size, right vertical axis) in MiB vs. workflow size.

Core i7 with 16GiB of RAM. For these scenarios, simulation times are more than 100x and up to 2500x shorter than real-world workflow executions. This is because SimGrid simulates computation and communication operations as delays computed based on computation and communication volumes using simulation models with low computational complexity.

To evaluate the scalability of our simulator, we use a workflow generator [62] to generate representative randomized configurations of the Montage workflow with from 1,000 up to 10,000 tasks. We generate 5 workflow instances for each number of tasks, and simulate the execution of these generated workflow instances on 128 cores (AWS-m5.xlarge with 32 4-core nodes). Figure 6 shows simulation time (left vertical axis) and maximum resident set size (right vertical axis) vs. the number of tasks in the workflow. Each sample point is the average over the 5 workflow instances (error bars are shown as well). As expected, both simulation time and memory footprint increase as workflows become larger. The memory footprint grows linearly with the number of tasks (simply due to the need to store more task objects). The simulation time grows faster initially, but then linearly beyond 7,000 tasks. We conclude that the simulation scales well, making it possible to simulate very large 10,000-task Montage configurations in under 40 minutes on a standard laptop computer.

Figure 6 also includes results obtained with WorkflowSim. We find that WorkflowSim has a larger memory footprint than our WRENCH-based simulator (by a factor ~ 1.48 for 10,000-task workflows). However, WorkflowSim is faster than our WRENCH-based simulator (by a factor ~ 1.81 for 10,000-task workflows), with roughly similar trends. The reason why

WorkflowSim is faster is because it simply does not simulate many aspects of the execution. The downside, as seen in the previous section, is that its simulation results are inaccurate.

V. CONCLUSION

In this paper we have presented WRENCH, a simulation framework for building simulators of Workflow Management Systems. WRENCH implements high-level simulation abstractions on top of the SimGrid simulation framework, so as to make it possible to build simulators that are accurate, that can run scalably on a single computer, and that can be implemented with minimal software development effort. Via a case study for the Pegasus production WMS we have demonstrated that WRENCH achieves these objectives, and that it favorably compares to a recently proposed workflow simulator. The main finding is that with WRENCH one can implement an accurate and scalable simulator of a complex real-world system with a few hundred lines of code. WRENCH is open source and welcomes contributors. WRENCH is already being used for several research and education projects, and Version 1.1 was released in August 2018. We refer the reader to <http://wrench-project.org> for software, documentation, and links to related projects.

A short-term development direction is to use WRENCH to simulate the execution of current production WMSs (as was done for Pegasus in Section IV). Although we have designed WRENCH with knowledge of these WMSs and with the intent of making their implementations with WRENCH feasible, we expect that WRENCH APIs and abstractions will evolve once we set out to realize these implementations. Another development direction is the implementation of more CI service abstractions in WRENCH, e.g., a Hadoop Compute Service, specific distributed cloud Storage Services. From a research perspective, a future direction is that of automated simulation calibration. As seen in our Pegasus case study, even when using validated simulation models, the values of a number of simulation parameters must be carefully chosen in order to obtain accurate simulation results. This issue is not confined to WRENCH, but is faced by all distributed system simulators. In our case study we have calibrated these parameters manually by analyzing and comparing simulated and real-world execution event traces. While, to the best of our knowledge, this is the typical practice, what is truly needed is an automated calibration method. Ideally, this method would process a (small) number of (not too large) real-world execution traces for “training scenarios”, and compute a valid and robust set of calibration parameter values. An important research question will then be to understand to which extent these automatically computed calibrations can be composed and extrapolated to scenarios beyond the training scenarios.

Acknowledgments. This work is funded by NSF contracts #1642369 and #1642335, “SI2-SSE: WRENCH: A Simulation Workbench for Scientific Workflow Users, Developers, and Researchers”, and by CNRS under grant #PICS07239. We thank Martin Quinson, Arnaud Legrand, and Pierre-François Dutot for their valuable help.

REFERENCES

- [1] I. J. Taylor, E. Deelman, D. B. Gannon, and M. Shields, *Workflows for e-Science: scientific workflows for grids*. Springer Publishing Company, Incorporated, 2014.
- [2] E. Deelman, K. Vahi, G. Juve, M. Rynge, S. Callaghan, P. J. Maechling, R. Mayani, W. Chen, R. Ferreira da Silva, M. Livny, and K. Wenger, "Pegasus: a Workflow Management System for Science Automation," *Future Generation Computer Systems*, vol. 46, pp. 17–35, 2015.
- [3] T. Fahringer, R. Prodan, R. Duan, J. Hofer, F. Nadeem, F. Nerieri, S. Podlipnig, J. Qin, M. Siddiqui, H.-L. Truong *et al.*, "Askalon: A development and grid computing environment for scientific workflows," in *Workflows for e-Science*. Springer, 2007, pp. 450–471.
- [4] M. Wilde, M. Hategan, J. M. Wozniak, B. Clifford, D. S. Katz, and I. Foster, "Swift: A language for distributed parallel scripting," *Parallel Computing*, vol. 37, no. 9, pp. 633–652, 2011.
- [5] K. Wolstencroft, R. Haines, D. Fellows, A. Williams, D. Withers, S. Owen, S. Soiland-Reyes, I. Dunlop, A. Nenadic, P. Fisher *et al.*, "The taverna workflow suite: designing and executing workflows of web services on the desktop, web or in the cloud," *Nucleic acids research*, p. gkt328, 2013.
- [6] I. Altintas, C. Berkley, E. Jaeger, M. Jones, B. Ludascher, and S. Mock, "Kepler: an extensible system for design and execution of scientific workflows," in *Scientific and Statistical Database Management, 2004. Proceedings. 16th International Conference on*. IEEE, 2004, pp. 423–424.
- [7] M. Albrecht, P. Donnelly, P. Bui, and D. Thain, "Makeflow: A portable abstraction for data intensive computing on clusters, clouds, and grids," in *1st ACM SIGMOD Workshop on Scalable Workflow Execution Engines and Technologies*. ACM, 2012, p. 1.
- [8] N. Vydyanathan, U. V. Catalyurek, T. M. Kurc, P. Sadayappan, and J. H. Saltz, "Toward optimizing latency under throughput constraints for application workflows on clusters," in *Euro-Par 2007 Parallel Processing*. Springer, 2007, pp. 173–183.
- [9] A. Benoit, V. Rehn-Sonigo, and Y. Robert, "Optimizing latency and reliability of pipeline workflow applications," in *Parallel and Distributed Processing, 2008. IPDPS 2008. IEEE International Symposium on*. IEEE, 2008, pp. 1–10.
- [10] Y. Gu and Q. Wu, "Maximizing workflow throughput for streaming applications in distributed environments," in *Computer Communications and Networks (ICCCN), 2010 Proceedings of 19th International Conference on*. IEEE, 2010, pp. 1–6.
- [11] M. Malawski, G. Juve, W. Deelman, and J. Nabrzyski, "Algorithms for cost- and deadline-constrained provisioning for scientific workflow ensembles in IaaS clouds," *Future Generation Computer Systems*, vol. 48, pp. 1–18, 2015.
- [12] J. Chen and Y. Yang, "Temporal dependency-based checkpoint selection for dynamic verification of temporal constraints in scientific workflow systems," *ACM Transactions on Software Engineering and Methodology (TOSEM)*, vol. 20, no. 3, p. 9, 2011.
- [13] G. Kandaswamy, A. Mandal, D. Reed *et al.*, "Fault tolerance and recovery of scientific workflows on computational grids," in *Cluster Computing and the Grid, 2008. CCGRID'08. 8th IEEE International Symposium on*. IEEE, 2008, pp. 777–782.
- [14] R. Ferreira da Silva, T. Glatard, and F. Desprez, "Self-healing of workflow activity incidents on distributed computing infrastructures," *Future Generation Computer Systems*, vol. 29, no. 8, pp. 2284–2294, 2013.
- [15] W. Chen, R. Ferreira da Silva, E. Deelman, and T. Fahringer, "Dynamic and fault-tolerant clustering for scientific workflows," *IEEE Transactions on Cloud Computing*, vol. 4, no. 1, pp. 49–62, 2016.
- [16] H. M. Fard, R. Prodan, J. J. D. Barrionuevo, and T. Fahringer, "A multi-objective approach for workflow scheduling in heterogeneous environments," in *Proceedings of the 2012 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (ccgrid 2012)*. IEEE Computer Society, 2012, pp. 300–309.
- [17] I. Pietri, M. Malawski, G. Juve, E. Deelman, J. Nabrzyski, and R. Sakellariou, "Energy-constrained provisioning for scientific workflow ensembles," in *Cloud and Green Computing (CGC), 2013 Third International Conference on*. IEEE, 2013, pp. 34–41.
- [18] M. Tikir, M. Laurenzano, L. Carrington, and A. Snaveley, "PSINS: An Open Source Event Tracer and Execution Simulator for MPI Applications," in *Proc. of the 15th Intl. Euro-Par Conf. on Parallel Processing*, ser. LNCS, no. 5704. Springer, Aug. 2009, pp. 135–148.
- [19] T. Hoefer, T. Schneider, and A. Lumsdaine, "LogGOPSIm - Simulating Large-Scale Applications in the LogGOPS Model," in *Proc. of the ACM Workshop on Large-Scale System and Application Performance*, Jun. 2010, pp. 597–604.
- [20] G. Zheng, G. Kakulapati, and L. Kalé, "BigSim: A Parallel Simulator for Performance Prediction of Extremely Large Parallel Machines," in *Proc. of the 18th Intl. Parallel and Distributed Processing Symposium (IPDPS)*, Apr. 2004.
- [21] R. Bagrodia, E. Deelman, and T. Phan, "Parallel Simulation of Large-Scale Parallel Applications," *IJHPCA*, vol. 15, no. 1, pp. 3–12, 2001.
- [22] W. H. Bell, D. G. Cameron, A. P. Millar, L. Capozza, K. Stockinger, and F. Zini, "OptorSim - A Grid Simulator for Studying Dynamic Data Replication Strategies," *IJHPCA*, vol. 17, no. 4, pp. 403–416, 2003.
- [23] R. Buyya and M. Murshed, "GridSim: A Toolkit for the Modeling and Simulation of Distributed Resource Management and Scheduling for Grid Computing," *Concurrency and Computation: Practice and Experience*, vol. 14, no. 13-15, pp. 1175–1220, Dec. 2002.
- [24] S. Ostermann, R. Prodan, and T. Fahringer, "Dynamic Cloud Provisioning for Scientific Grid Workflows," in *Proc. of the 11th ACM/IEEE Intl. Conf. on Grid Computing (Grid)*, 2010, pp. 97–104.
- [25] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. F. De Rose, and R. Buyya, "CloudSim: A Toolkit for Modeling and Simulation of Cloud Computing Environments and Evaluation of Resource Provisioning Algorithms," *Software: Practice and Experience*, vol. 41, no. 1, pp. 23–50, Jan. 2011.
- [26] A. Nez, J. Viquez-Poletti, A. Caminero, J. Carretero, and I. M. Llorente, "Design of a New Cloud Computing Simulation Platform," in *Proc. of the 11th Intl. Conf. on Computational Science and its Applications*, June 2011, pp. 582–593.
- [27] G. Kecskemeti, "DISSECT-CF: A simulator to foster energy-aware scheduling in infrastructure clouds," *Simulation Modelling Practice and Theory*, vol. 58, no. 2, pp. 188–218, 2015.
- [28] A. Montresor and M. Jelasity, "PeerSim: A Scalable P2P Simulator," in *Proc. of the 9th Intl. Conf. on Peer-to-Peer*, Sep. 2009, pp. 99–100.
- [29] I. Baumgart, B. Heep, and S. Krause, "OverSim: A Flexible Overlay Network Simulation Framework," in *Proc. of the 10th IEEE Global Internet Symposium*. IEEE, May 2007, pp. 79–84.
- [30] M. Tauber, A. Kerstens, T. Estrada, D. Flores, and P. J. Teller, "SimBA: A Discrete Event Simulator for Performance Prediction of Volunteer Computing Projects," in *Proc. of the 21st Intl. Workshop on Principles of Advanced and Distributed Simulation*, 2007, pp. 189–197.
- [31] T. Estrada, M. Tauber, K. Reed, and D. P. Anderson, "EmBOINC: An Emulator for Performance Analysis of BOINC Projects," in *Proc. of the Workshop on Large-Scale and Volatile Desktop Grids (PCGrid)*, 2009.
- [32] D. Kondo, "SimBOINC: A Simulator for Desktop Grids and Volunteer Computing Systems," Available at <http://simboinc.gforge.inria.fr/>, 2007.
- [33] H. Casanova, A. Giersch, A. Legrand, M. Quinson, and F. Suter, "Versatile, Scalable, and Accurate Simulation of Distributed Applications and Platforms," *Journal of Parallel and Distributed Computing*, vol. 74, no. 10, pp. 2899–2917, 2014.
- [34] C. D. Carothers, D. Bauer, and S. Pearce, "ROSS: A High-Performance, Low Memory, Modular Time Warp System," in *Proc. of the 14th ACM/IEEE/SCS Workshop of Parallel on Distributed Simulation*, 2000, pp. 53–60.
- [35] W. Chen and E. Deelman, "WorkflowSim: A Toolkit for Simulating Scientific Workflows in Distributed Environments," in *Proc. of the 8th IEEE Intl. Conf. on E-Science*, 2012, pp. 1–8.
- [36] A. Hiraes-Carbajal, A. Tchernykh, T. Rblitz, and R. Yahyapour, "A Grid simulation framework to study advance scheduling strategies for complex workflow applications," in *In Proc. of IEEE Intl. Symp. on Parallel Distributed Processing Workshops (IPDPSW)*, 2010.
- [37] M.-H. Tsai, K.-C. Lai, H.-Y. Chang, K. Fu Chen, and K.-C. Huang, "Pewss: A platform of extensible workflow simulation service for workflow scheduling research," *Software: Practice and Experience*, vol. 48, no. 4, pp. 796–819, 2017.
- [38] S. Ostermann, K. Plankensteiner, D. Bodner, G. Kraler, and R. Prodan, "Integration of an Event-Based Simulation Framework into a Scientific Workflow Execution Environment for Grids and Clouds," in *In proc. of the 4th ServiceWave European Conference*, 2011, pp. 1–13.
- [39] G. Kecskemeti, S. Ostermann, and R. Prodan, "Fostering Energy-Awareness in Simulations Behind Scientific Workflow Management Systems," in *Proc. of the 7th IEEE/ACM Intl. Conf. on Utility and Cloud Computing*, 2014, pp. 29–38.

- [40] J. Cao, S. Jarvis, S. Saini, and G. Nudd, "GridFlow: Workflow Management for Grid Computing," in *Proc. of the 3rd IEEE/ACM Intl. Symp. on Cluster Computing and the Grid (CCGrid)*, 2003, pp. 198–205.
- [41] "The SimGrid Project," Available at <http://simgrid.org/>, 2018.
- [42] P. Bedaride, A. Degomme, S. Genaud, A. Legrand, G. Markomanolis, M. Quinson, M. Stillwell, F. Suter, and B. Videau, "Toward Better Simulation of MPI Applications on Ethernet/TCP Networks," in *Prod. of the 4th Intl. Workshop on Performance Modeling, Benchmarking and Simulation of High Performance Computer Systems*, 2013.
- [43] P. Velho, L. Mello Schnorr, H. Casanova, and A. Legrand, "On the Validity of Flow-level TCP Network Models for Grid and Cloud Simulations," *ACM Transactions on Modeling and Computer Simulation*, vol. 23, no. 4, 2013.
- [44] P. Velho and A. Legrand, "Accuracy Study and Improvement of Network Simulation in the SimGrid Framework," in *Proc. of the 2nd Intl. Conf. on Simulation Tools and Techniques*, 2009.
- [45] K. Fujiwara and H. Casanova, "Speed and Accuracy of Network Simulation in the SimGrid Framework," in *Proc. of the 1st Intl. Workshop on Network Simulation Tools*, 2007.
- [46] A. Lèbre, A. Legrand, F. Suter, and P. Veyre, "Adding Storage Simulation Capacities to the SimGrid Toolkit: Concepts, Models, and API," in *Proc. of the 8th IEEE Intl. Symp. on Cluster Computing and the Grid*, 2015.
- [47] "The WRENCH Project," <http://wrench-project.org>, 2018.
- [48] "The ns-3 Network Simulator," Available at <http://www.nsnam.org>.
- [49] E. León, R. Riesen, A. Maccabe, and P. Bridges, "Instruction-Level Simulation of a Cluster at Scale," in *Proc. of the Intl. Conf. for High Performance Computing and Communications (SC)*, Nov. 2009.
- [50] R. Fujimoto, "Parallel Discrete Event Simulation," *Commun. ACM*, vol. 33, no. 10, pp. 30–53, 1990.
- [51] R. Matha, S. Ristov, and R. Prodan, "Simulation of a workflow execution as a real Cloud by adding noise," *Simulation Modelling Practice and Theory*, vol. 79, pp. 37–53, 2017.
- [52] L. Bobelin, A. Legrand, D. A. G. Márquez, P. Navarro, M. Quinson, F. Suter, and C. Thiery, "Scalable Multi-Purpose Network Representation for Large Scale Distributed System Simulation," in *Proceedings of the 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid)*, Ottawa, Canada, May 2012, pp. 220–227.
- [53] M. Turilli, M. Santcroos, and S. Jha, "A Comprehensive Perspective on Pilot-Job Systems," *ACM Comput. Surv.*, vol. 51, no. 2, pp. 43:1–43:32, 2018.
- [54] "Pegasus' DAX Workflow Description Format," https://pegasus.isi.edu/documentation/creating_workflows.php, 2018.
- [55] "The Standard Workload Format," <http://www.cs.huji.ac.il/labs/parallel/workload/swf.html>, 2018.
- [56] D. Lifka, "The ANL/IBM SP Scheduling System," in *Proc. of the 1st Workshop on Job Scheduling Strategies for Parallel Processing, LCNS*, vol. 949, 1995, pp. 295–303.
- [57] F. Dabek, R. Cox, F. Kaashoek, and R. Morris, "Vivaldi: A Decentralized Network Coordinate System," in *Proc. of SIGCOMM*, 2004.
- [58] J. Frey, "Condor dagman: Handling inter-job dependencies," *University of Wisconsin, Dept. of Computer Science, Tech. Rep.*, 2002.
- [59] 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.
- [60] "The WRENCH Pegasus Simulator," <https://github.com/wrench-project/pegasus>, 2018.
- [61] R. Ferreira da Silva, R. Filgueira, E. Deelman, E. Pairo-Castineira, I. M. Overton, and M. Atkinson, "Using simple pid controllers to prevent and mitigate faults in scientific workflows," in *11th Workflows in Support of Large-Scale Science*, ser. WORKS'16, 2016, pp. 15–24.
- [62] R. Ferreira da Silva, W. Chen, G. Juve, K. Vahi, and E. Deelman, "Community resources for enabling and evaluating research on scientific workflows," in *10th IEEE International Conference on e-Science*, ser. eScience'14, 2014, pp. 177–184.