



Bader, J., Lehmann, F., Thamsen, L., Will, J., Leser, U. and Kao, O. (2022) Lotaru: Locally Estimating Runtimes of Scientific Workflow Tasks in Heterogeneous Clusters. In: 34th International Conference on Scientific and Statistical Database Management (SSDBM 2022), Copenhagen, Denmark, 6-8 July 2022, ISBN 9781450396677 (doi: [10.1145/3538712.3538739](https://doi.org/10.1145/3538712.3538739)).

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

© The Authors 2022. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in Proceedings of the 34th International Conference on Scientific and Statistical Database Management (SSDBM 2022), Copenhagen, Denmark, 6-8 July 2022, ISBN 9781450396677.

<http://eprints.gla.ac.uk/271579/>

Deposited on: 23 May 2022

Lotaru: Locally Estimating Runtimes of Scientific Workflow Tasks in Heterogeneous Clusters

Jonathan Bader
Technische Universität Berlin
Berlin, Germany
jonathan.bader@tu-berlin.de

Fabian Lehmann
Humboldt-Universität zu Berlin
Berlin, Germany
fabian.lehmann@informatik.hu-berlin.de

Lauritz Thamsen
University of Glasgow
Glasgow, United Kingdom
lauritz.thamsen@glasgow.ac.uk

Jonathan Will
Technische Universität Berlin
Berlin, Germany
will@tu-berlin.de

Ulf Leser
Humboldt-Universität zu Berlin
Berlin, Germany
leser@informatik.hu-berlin.de

Odej Kao
Technische Universität Berlin
Berlin, Germany
odej.kao@tu-berlin.de

ABSTRACT

Many scientific workflow scheduling algorithms need to be informed about task runtimes a-priori to conduct efficient scheduling. In heterogeneous cluster infrastructures, this problem becomes aggravated because these runtimes are required for each task-node pair. Using historical data is often not feasible as logs are typically not retained indefinitely and workloads as well as infrastructure changes. In contrast, online methods, which predict task runtimes on specific nodes while the workflow is running, have to cope with the lack of example runs, especially during the start-up.

In this paper, we present *Lotaru*, a novel online method for locally estimating task runtimes in scientific workflows on heterogeneous clusters. *Lotaru* first profiles all nodes of a cluster with a set of short-running and uniform microbenchmarks. Next, it runs the workflow to be scheduled on the user's local machine with drastically reduced data to determine important task characteristics. Based on these measurements, *Lotaru* learns a Bayesian linear regression model to predict a task's runtime given the input size and finally adjusts the predicted runtime specifically for each task-node pair in the cluster based on the micro-benchmark results. Due to its Bayesian approach, *Lotaru* can also compute robust uncertainty estimates and provides them as an input for advanced scheduling methods.

Our evaluation with five real-world scientific workflows and different datasets shows that *Lotaru* significantly outperforms the baselines in terms of prediction errors for homogeneous and heterogeneous clusters.

CCS CONCEPTS

- **Information systems** → **Information systems applications**;
- **Computer systems organization** → **Distributed architectures**;
- **Software and its engineering** → **Software architectures**.

KEYWORDS

Resource Management, Scientific Workflow, Runtime Estimation, Heterogeneous Cluster Resources, Scheduling

ACM Reference Format:

Jonathan Bader, Fabian Lehmann, Lauritz Thamsen, Jonathan Will, Ulf Leser, and Odej Kao. 2022. Lotaru: Locally Estimating Runtimes of Scientific Workflow Tasks in Heterogeneous Clusters. In *34th International Conference on Scientific and Statistical Database Management (SSDBM 2022)*, July 6–8, 2022, Copenhagen, Denmark. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3538712.3538739>

1 INTRODUCTION

Scientists from many domains, such as bioinformatics, remote sensing, and physics, use scientific workflow management systems to define, compose, and reproducibly execute their data analysis pipelines over large datasets [13, 14, 45]. These workflows are commonly organized as a directed acyclic graph (DAG), consisting of a set of abstract tasks T and a set of directed edges E . While the abstract tasks serve as templates for their physical instances on real datasets, edges describe the flow of data between tasks and thus constrain the order of execution and degree of parallelism of task executions.

Fig. 1 shows an exemplary abstract workflow and a concrete physical execution. The physical representation with two input data samples shows that the physical tasks B1, B2 and C1, C2, respectively, may be executed in parallel as they have no interdependencies. Tasks D, E, F, and G result in only one physical task.

When such workflows are executed over large amounts of data, their runtimes can easily exceed days or even weeks [2, 8, 26, 27]. To effectively use the available cluster resources, many workflow management systems have a scheduling component [6, 12, 24, 31] that determines which tasks are executed when and on which of the available nodes to achieve some optimization goal, such as wallclock-time. To this end, most practical methods today require accurate estimates of the runtime of any task on any node [3, 38]. As these are difficult to obtain, schedulers often resort to user configurations, which, however, are known to be highly error-prone [43, 44]. In practice, the problem is complicated further because clusters available to scientists for their data analysis often consist of heterogeneous hardware [2, 39]. Reasons are, for instance, partially upgraded nodes, hardware replacements over time, or clusters that are intended to serve multiple purposes. In such settings, nodes' basic resource properties differ, like the size and speed of main

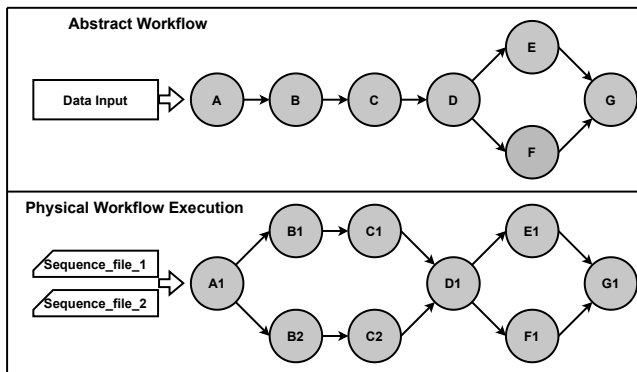


Figure 1: Abstract and physical execution model of a scientific workflow.

memory, number and frequency of cores, network latency and bandwidth, cache sizes, and local I/O throughput [2, 39]. Accordingly, the same physical task will exhibit different runtimes on different nodes. This opens the problem that schedulers need accurate task runtime estimates not only per task but actually per task-node pair. Such information, however, is very often unavailable.

This paper presents Lotaru, a novel method addressing this problem through infrastructure profiling, local workflow executions on downsampled partitions of the entire input, and a Bayesian framework to transfer measured runtimes to specific task-node pairs. It is intended for workflows for embarrassingly parallel problems where the same (sub-)workflows are executed over many inputs or intermediate results. Examples of such problems are abundant in areas like remote sensing (sub workflows analyze many images in parallel [4, 19]), bioinformatics (sub workflows analyze many read sets in parallel [20, 46]), or material science (sub workflows analyze many molecules in parallel [34, 37]). In Figure 1, the sub-workflow is executed once for each of the two input files.

In Lotaru’s first step, the infrastructure profiler analyzes the performance characteristics of a local computer and the nodes in the cluster (called target nodes). Subsequently, one of the foreseen input files is selected and downsampled or sliced into several smaller files, which next are used as input to run the workflow locally and to gather metrics. The workflow is executed twice, once with all samples and once with a subset at reduced CPU speed to identify bottlenecks. Next, Lotaru trains a Bayesian linear regression model for each task based on the collected runtimes. This model is subsequently used to predict task runtimes for arbitrary input sizes. Lastly, these predictions are adjusted to all target nodes using the initial infrastructure profiling measurements.

We implemented Lotaru¹ and compared it experimentally to three different baselines on a cluster of six different kinds of nodes using five real-life scientific workflows from the popular nf-core² workflow repository [17] with different inputs. Our experiments based on the popular workflow engine Nextflow [14] show that Lotaru significantly outperforms the baselines regarding prediction errors for homogeneous and for heterogeneous clusters.

Lotaru is designed as an online method, i.e., it does not depend on any historic information but performs all measurements and estimations before the start of a specific workflow execution. Notably, this also allows for offline scenarios where the learned models are reused for future executions of the same workflow over different input data. To this end, we also make available 9,317 unique task execution traces from our experiments on six different machines³.

2 RELATED WORK

First, we cover scientific workflow management systems (SWMS) and scheduling strategies that could be used by such systems to show the need for runtime estimates. We discuss approaches focusing on runtime prediction in general and finally focus on cross infrastructure task-runtime prediction. Note that in the following, we focus mostly on the prediction of runtimes. However, estimation methods for the usage of other resources, e.g. main memory or network bandwidth, often apply similar techniques [43, 45].

2.1 Scientific Workflow Management Systems and Resource Manager

SWMS like Pegasus [12], Saasfee [6], and Nextflow [14] use workflow languages to define data analysis workflows in an abstract manner. When executing a workflow on concrete inputs, they derive the physical workflow consisting of physical tasks and steer their execution with the help of a distributed resource manager, such as Slurm [47], Kubernetes [5], or YARN [40]. The resource manager coordinates the target infrastructure and acts as an intermediary between the SWMS and the cluster resources. Current systems often do not perform any advanced form of scheduling because they do not have the capabilities to perform resource predictions. Instead, SWMS send each ready-to-run task to the resource manager together with the user-defined requirements. The resource manager in turn schedules them for execution in a FIFO or fair manner. For example, several YARN distributions use a fair-like scheduler^{4,5}.

2.2 Scheduling Workflow Tasks onto Heterogeneous Clusters

Despite their missing uptake in real systems, the literature on more advanced scheduling algorithms is vast. Scheduling of tasks onto heterogeneous infrastructures can be done in two ways, either statically or dynamically [16, 41]. Static scheduling heuristics like HEFT [38] and HCPPEFT [11] map all tasks to computing resources before the workflow execution. Therefore, these approaches cannot adapt to infrastructure failures or changes in the physical workflow execution plan. In contrast, dynamic scheduling approaches like P-HEFT [3] and FDWS [1] map tasks to infrastructure components at runtime and are therefore more flexible. In particular, they can also be applied when the set of physical tasks depends on intermediate results, in which case the complete physical workflow cannot be fixed before execution [7]. Both approaches have in common that they need at least comprehensive knowledge about execution times of all tasks on all available nodes. However, these values are not

¹github.com/CRC-FONDA/Lotaru

²nf-co.re

³github.com/CRC-FONDA/Lotaru-traces

⁴docs.datafabric.hpe.com/62/AdministratorGuide/Job-Scheduling.html

⁵bdllabs.edureka.co/static/help/topics/admin_fair_scheduler.html

available in advance but must be determined either by asking users for estimates [18, 22, 23], by analyzing historical traces [35, 36, 42], or by using some form of online learning [43, 45]. Lotaru aims to estimate the runtime for all task-node pairs in a cluster to enable the use of existing scheduling methods in real-world systems.

2.3 Task Runtime Prediction Based on Historical Runtime Data

Runtime prediction for scientific workflow tasks based on historical data has been extensively researched. Recent approaches use machine learning for this problem [9, 10, 30, 33].

Nadeem et al. [30] propose the use of neural networks to predict workflow execution times on a grid. Their learning model considers several types of information, like workflow structure, measured task resource requirements in historical traces, and input file sizes. They further study which features influence the runtime the most and which ones can be omitted.

Da Silva et al. [9, 10] estimate runtime, disk space, and memory consumption for tasks in scientific workflows. They use monitoring tools and historical data to apply regression trees for actual predictions. In a pre-processing step, they use density-based clustering to identify data subsets with high correlation. If the data in the selected cluster is correlated, the authors estimate the expected resource usages based on the ratio in the cluster. For uncorrelated data, the authors test a Gamma and a Normal distribution to generate an estimation value. The estimations are updated at workflow runtime once new information becomes available. As we use them as competitors for Lotaru, we will describe them in more detail in Section 4.3.

We call such approaches offline estimators, because they build their initial models on historical data prior to the actual workflow execution. Note that offline estimators are generally not applicable for new workflows with new tasks, as for these, no historical traces for model learning are available. In contrast, Lotaru is designed as an online method which is independent from any historical traces and can be applied out-of-the-box for any kind of workflow on any kind of cluster.

2.4 Cross-Infrastructure Task-Runtime Prediction

Because real clusters are often comprised of nodes with heterogeneous capabilities, all methods that learn a model from past task executions (offline or online) have to consider the challenge that their models must generalize to different nodes. This is particularly important when running a workflow on a completely different cluster, i.e., in a cross-infrastructure setting.

Pham et al. [32] use a two-stage prediction approach to estimate the task execution time in cloud environments. As prediction parameters, they consider input parameters of the workflow, VM specifications, and runtime parameters like memory usage and read/write operations. In the first stage, the authors derive the runtime parameters for the execution on the target VM. The second stage uses the output data from the first stage together with workflow input data and the VM specifications to learn a regression model that predicts the execution time of a given task on a target VM.

Hilman et al. [21] apply an online incremental learning approach using long short-term memory networks (LSTMs) to predict task runtimes. As input features, they consider task characteristics, like the name of the executable, the input data, VM details, and the submission time to encounter for daily trends in resource usage. These features are extended through historical time-series data for CPU usage, memory usage, and disk activities to conduct their incremental prediction method.

Matsunaga and Fortes [28] include system-specific attributes like CPU architecture and the size and speed of the memory to derive more accurate prediction models. Their study compares several machine learning approaches for two well-known bioinformatics applications.

Lotaru does not include explicit hardware characteristics as features for the runtime prediction as they are often difficult to correlate with the runtime of real tasks. Instead, we apply microbenchmarks for obtaining such characteristics in an implicit manner. Furthermore, many of the recent approaches use machine learning methods like neuronal networks or k-nearest neighbors, which are known to require large training data sets to perform well. In contrast, Lotaru uses a Bayesian Linear Regression that already works with few training points and that is also capable of deriving uncertainty estimates for its predictions [25, 29].

3 LOTARU APPROACH

Lotaru is a novel method that aims to predict workflow task runtimes for all nodes in heterogeneous clusters at workflow start-up time and without relying on historical data. Figure 2 provides an overview of our approach. In phase ①, Lotaru uses microbenchmarks to gather performance insights about the target infrastructure and the local developer machine. In the second phase ②, Lotaru selects one of the data inputs and downsamples it to create several small workflow input partitions. Then, the workflow is locally executed two times with the created partitions. For the second execution, only a few of these partitions are used, and the CPU speed is slightly reduced to identify CPU-bound tasks. In the third phase ③, Lotaru uses the collected data points to create a Bayesian linear regression model which predicts the runtime. In the last phase ④, the values from the benchmarks are used to adjust the prediction results for each node in the target infrastructure. The scheduler can then consider the prediction results for all task-node pairs to create a scheduling plan.

One assumption behind Lotaru is that the input consists of multiple input files, which are, at least partly, processed separately. We rely on this property to create local workflows with small inputs for fast measurements of first characteristics. Furthermore, in this work, we assume that Lotaru can downsample or slice an input file further to create a diverse set of local workflow inputs, allowing us to learn the relationship between input size and task runtimes. However, in applications where such a downsampling is not possible, Lotaru could also use different subsets of the input files at the price of a somewhat longer training phase.

① Local and Target Infrastructure Profiling

We expect the local developer machine to differ from the target cluster and the target nodes themselves to be heterogeneous. To

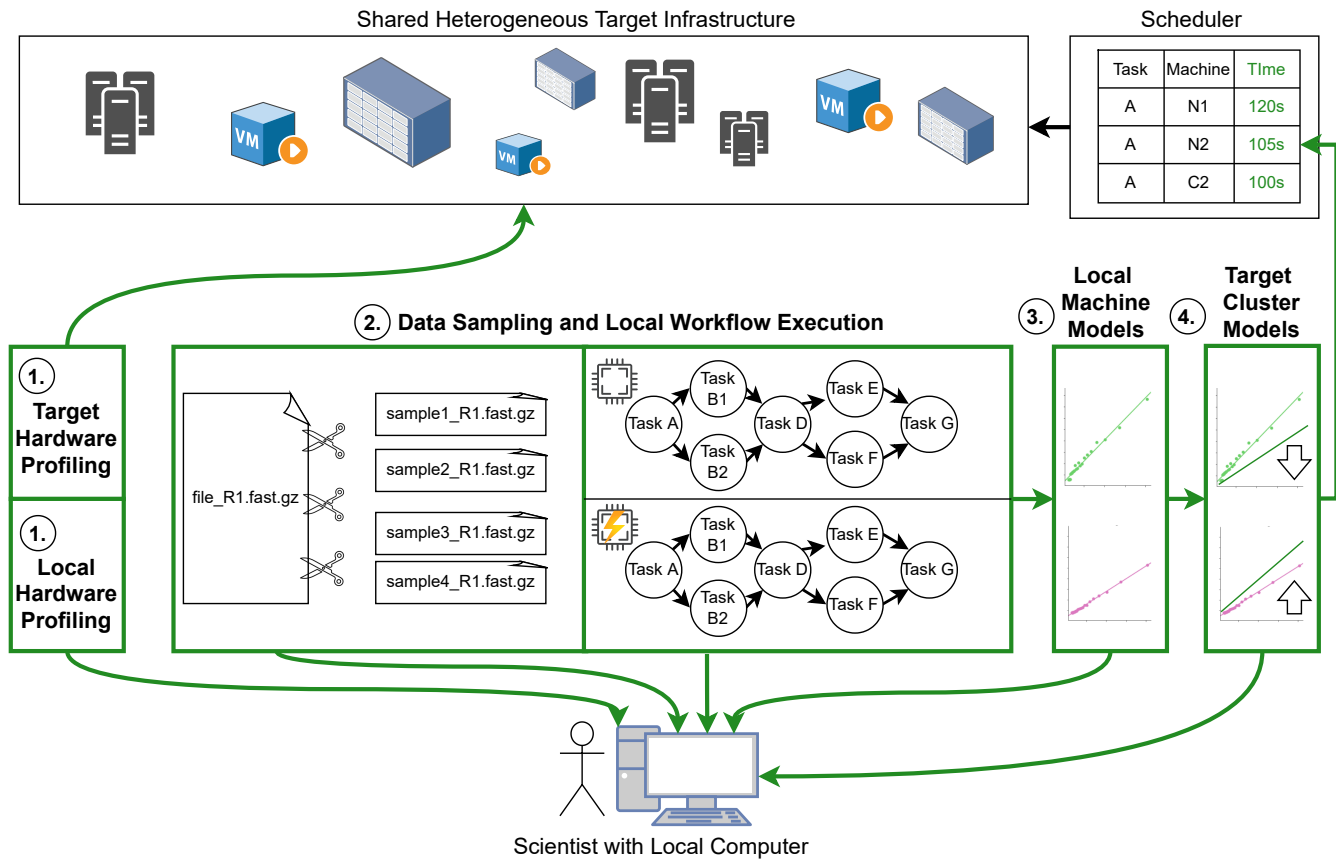


Figure 2: The process of Lotaru's local runtime prediction approach.

measure the differences, we conduct a short profiling phase to gather detailed infrastructure characteristics. Therefore, Lotaru analyzes all target nodes' dynamic performance characteristics like CPU speeds, memory speed, and random and sequential I/O. For this, we use microbenchmarks, which can be executed in parallel and take very short time, typically less than a minute, for each node. This step could be rerun automatically whenever a cluster's resource manager detects hardware changes.

② Data Sampling and Local Workflow Execution

In the next step, Lotaru trains a Bayesian regression model to predict task runtimes based on input size. To this end, it picks one of the input files and downsamples it further to obtain task runtimes for diverse yet small (and hence fast) inputs as input for the learner. For image data used in remote sensing or astronomic workflows, this means dividing a single image into smaller ones keeping the resolution, or decreasing the resolution while leaving the image section the same. In genomics, downsampling means splitting one of the many samples with millions of short sequence reads into multiple smaller partitions.

Next, Lotaru measures local runtimes over multiple different partitions. While using a large set of such partitions covering a large range of data sizes tend to improve the accuracy of the prediction

model, fewer and smaller partitions can be executed faster and lead to quicker but mostly more imprecise runtime estimates.

If, for instance, the data sampling process described before created five partitions, Lotaru runs the workflow with these five input files. This step delivers monitoring data for each task-partition pair but gives no direct insights whether, i.e., a task is CPU-intensive. To identify which hardware resource the task mostly depends on, we decrease the CPU frequency of the local machine by 20% and run the workflow again with the five created partitions. Thereby, we expect CPU-intensive tasks to take around 25% longer.


The sizes of the samples obviously are an important hyperparameter of Lotaru whose effect will be studied in Section 5.1.

③ Local Prediction Model Training

Most existing approaches use the file size on disk as the input or part of the input vector for their predictions or statistical models. We argue that Lotaru should use the uncompressed input data size for compressed files, which scientific workflows frequently use due to the large amount of data. For example, in bioinformatics workflows, the de facto standard file format for storing biological sequences is fastq which is compressed with Gzip. Gzip can compress larger files efficiently, especially when dealing with repetitive data, leading to a

Table 1: Example of a model adjustment for a single task prediction on two target nodes.

Machine	CPU	RAM	I/O	Weight w	Bayesian Prediction	Factor	Final Runtime Prediction
Local	500	20,000	500	0.80	100.00s	1.00	100.00s
N1	400	18,000	300	-	-	1.33	133.00s
N2	520	20,000	500	-	-	0.96	96.00s

**Figure 3: Posterior Prediction for the task FASTQC with an uncompressed input size of ~4,585 MB.**

non-linear file size increase. For example, the file `example.fastq.gz`⁶ has a size of ~2,014 MB and contains 40,517,845 sequences. Splitting the file into two files with the same number of sequences leads to two files with a size of ~1,274 MB each, an increase of 26.46%. Accordingly, Lotaru should not use the compressed file size to predict the runtime.

Using the uncompressed data size and a task’s runtime, Lotaru checks for a linear correlation between both. For this, we use the Pearson correlation coefficient, which is defined as follows:

$$p = \frac{\sum(x_i - \bar{x})(y_i - \bar{y})}{\sqrt{\sum(x_i - \bar{x})^2 \sum(y_i - \bar{y})^2}} \quad (1)$$

where the x values are the actual uncompressed task input data size and y values are the runtime. We define a correlation as significant if p is greater than 0.8. Lotaru uses a Bayesian linear regression to predict the runtime if p is significant. Otherwise, we estimate the median runtime as the expected runtime independent of the concrete input size.

One of the main advantages of using the Bayesian approach is that Lotaru can train it on a small training data set [25, 29], which is especially useful since the local profiling only delivers a few training points for each task. Additionally, while Lotaru can predict the most probable runtime, the Bayesian approach also yields an uncertainty value for this prediction. For example, Lotaru estimates that task A is expected to run 120 seconds. However, it also gives a lower and upper uncertainty at different confidence levels to express that the point estimate probably is not accurate.

As illustration, Figure 3 shows the prediction for the task FASTQC in the ChIPseq-1 workflow with an input data size of ~4,585 MB (see Section 4.2 for details on the data and the workflow). The prediction error for the estimated mean value is around 0.30%. The value is in a confidence interval of 23.23% uncertainty. With a confidence of 50% uncertainty, we would consider the runtime to be between 99.4s and 100.7s. The scheduler can consider this uncertainty and plan with it, which would not be possible for frequentist approaches.

In contrast to these classical frequentist approaches, we try to find a posterior distribution for our model parameters. Specifically,

our model computes a posterior distribution depending on the input values as shown in Formula 2, we assume that ϵ_i and y_i are normally distributed with $\epsilon_i \sim N(0, \sigma^2)$ and $y_i \sim N(b^T x_i; \sigma^2)$

$$y_i = x_i^T b + \epsilon_i; \quad (2)$$

Our Bayesian approach now tries to maximize the posterior

$$\max_b P(b|y_i); \quad (3)$$

Through the rules of Bayes and a short equation transformation, we get to

$$\max_b P(y_i|b)P(b); \quad (4)$$

where the first term, $P(y_i|b)$, the likelihood, can be computed with our previous assumption about the distribution of y_i . For the second term, the posterior, $P(b)$, we have to assume a distribution. We decided to set the prior to a Gaussian distribution, which results in an L2-regressor for our Bayesian regression.

In our prediction model, x_i is a scalar, i.e., the uncompressed task data input size, and y_i is the runtime.

④ Model Adjustment for Target Infrastructure

With the Bayesian Linear Regression model created, Lotaru can now predict the tasks’ runtime for nodes with the same hardware as the developer machine. However, we aim to estimate runtimes for all different kinds of nodes in the cluster. Therefore, we take the monitoring data from the infrastructure profiling and the local workflow runs. For each abstract task, we compare the runtimes for the task-sample pairs which occurred in the two local workflow executions, the normal one and the one with reduced CPU speed. Lotaru determines the deviation of a task’s pair’s runtime as $dev = \frac{time_{new} - time_{old}}{time_{old}}$, where *old* refers to the normal execution and *new* to the execution with reduced CPU speed.

Since each task at least has to read the input file and to write the output file, I/O capabilities are essential for every task and are, therefore, included in the following adjustment step. Another important factor is CPU speed, whereas in prior experiments we found that memory speed only has a negligible impact. Thus, we decided to exclude the memory speed from our adjustment step. Accordingly, we define the impact of the CPU and the I/O on a task’s execution time through the following weighting:

$$w = \max\left(0; \min\left(1; \frac{median_{dev}}{(freq_{old}/freq_{new}) - 1}\right)\right) \quad (5)$$

Then, Lotaru sets the runtime factor for each task as following:

$$f_t = w * \frac{cpu_{local}}{cpu_{target}} + (1 - w) * \frac{io_{local}}{io_{target}} \quad (6)$$

⁶s3://nf-core-awsmegetests/eager/ENA_Data_Fish/ERR1943601_1.fastq.gz

Table 1 gives an example for a single task and a prediction on two different target nodes, namely N1 and N2. The values with the grey background are from the profiling step, and the CPU-weight w for the factor is 0.8. For one out of the various inputs, e.g., with a size of 10GB, a runtime of 100 seconds is predicted on the local machine with the Bayesian regression model. We can now adjust the estimated runtime to the target nodes by calculating the factor with the monitoring values. The results are depicted in yellow.

4 EXPERIMENTAL SETUP

This section presents our prototype implementation and the environment where we run our experiments. The source code to reproduce the evaluation is available online⁷.

4.1 Prototype Implementation

Infrastructure Profiler: The infrastructure profiler uses *sysbench*⁸ as a first measuring tool and *LINPACK* [15] as a second microbenchmark since both measure different CPU characteristics and are well-known. With *sysbench*, we run a benchmark that verifies prime numbers with a limit of ten seconds and a maximum verification prime number of 20,000. Additionally, *LINPACK* measures the Floating Point Operations Per Second (FLOPS) using a default array size of 200x200. The memory speed is tested through *sysbench*, setting the block size buffer to one megabyte and the total memory size to 100 gigabytes.

Since we run *sysbench* on computers that differ in the number of CPU cores, we decided to always set the number of benchmarked CPU threads to one. Therefore, we differentiate between a node’s overall and single-core performance because the resource manager assigns a fixed number of cores to workflow tasks. Therewith, we have a better comparability, and Lotaru avoids that nodes with more but weaker CPU cores score higher. Otherwise, for example, a node with four very powerful cores would score less than a node with eight slower cores.

Lotaru tests the I/O performance of the local machine and target cluster by using *fiio*⁹. We conduct a benchmark only for sequential read-write and avoid measuring random read-write characteristics since this access pattern is scarce in scientific data analysis. Typically big input files are read sequentially. Further, we omit memory use and the page cache, so these components do not influence the I/O measurements.

Nowadays, hardware is tailored to achieve many points in popular benchmarks like *sysbench* and *fiio*. Note, however, that our goal is not to benchmark the actual performance of the hardware but to compare the performance of different nodes for adapting runtime estimates.

Data Sampling and Local Runtime Prediction: Lotaru relies on downsampling of workflow input data to obtain measurements for task executions quickly. Such downsampling obviously must take the actual nature of the data into account and therefore must be implemented specifically for different input data. As all our evaluation workflows run on genome sequencing data, we implemented

downsampling for genome data in the fastq format using the open-source software *fastqsplitter*¹⁰ to split the inputs in partitions. Note that Lotaru features an interface to support downsampling or slicing files in arbitrary domains.

For managing the CPU frequency, we use the userspace tool *cpupower*¹¹. As a workflow management system, we choose *Nextflow* [14], which gathers task runtime metrics already by default. We extended the monitoring interface of *Nextflow* to collect additional data, like compressed and uncompressed input size of tasks and the overall workflow input size.

4.2 Cluster Setup and Evaluation Workflows

We evaluate our approach on six different machines: a local machine, two machines from a heterogeneous commodity cluster, and three virtual machines in the Google Cloud Platform (GCP). Specifications can be found in Table 2 together with results of our microbenchmarks. The local machine consists of an Intel Xeon E3-1230 V2 CPU (four cores, eight threads, 3.30 GHz base frequency), 16GB memory, and an HDD. The two machines from the commodity cluster, A1 and A2, have two Intel Xeon X5355 2.66 GHz each, 32GB of memory and different hard drives. Since the *LINPACK* benchmark failed on A1 and A2, due to the age of the machines, the values are not included in the table and only the *sysbench* score is used for the factor. We use N1, N2, and C2 instances as (heterogeneous) nodes in the cluster. While the N1 machines are based on Intel Broadwell with a base clock of 2.0 GHz, the N2 machines use Intel Cascade Lake CPUs with a base clock of 2.8 GHz. The C2 machines are compute-optimized and based on Intel Cascade Lake with a turbo clock of up to 3.8 GHz¹².

We selected five real-world bioinformatics workflows from the nf-core repository¹³ and ran each of them with two different data sets to evaluate our approach. Table 3 gives an overview of these workflows and data inputs. Each of the five workflows performs different types of sequence analysis: The Eager workflow [46] analyzes ancient genomic data, the ChIPseq workflow¹⁴ is used to analyze Chromatin Immunoprecipitation sequencing (ChIP-seq) data, and the Methyleseq workflow¹⁵ is used for analyzing Bisulfite sequencing data in epigenomics. Atacseq¹⁶ analyzes ATAC-seq data and Bacass¹⁷ is a workflow for simple bacterial assembly and annotation. The workflow runtimes in Table 3 were obtained by running the workflows with one of various data inputs on the local machine. Note, typically, workflows run hundreds or thousands of inputs, resulting in much higher workflow runtimes for real inputs, even on large-scale clusters.

4.3 Baselines

We compare the accuracy of Lotaru’s runtime predictions with three baselines: a Naive Approach (NA), Online-M [9], and Online-P [10].

⁷github.com/CRC-FONDA/Lotaru

⁸github.com/akopytov/sysbench

⁹github.com/axboe/fio

¹⁰github.com/LUMC/fastqsplitter

¹¹linux.die.net/man/1/cpupower

¹²cloud.google.com/compute/docs/machine-types

¹³github.com/nf-core

¹⁴github.com/nf-core/chipseq

¹⁵github.com/nf-core/methyleseq

¹⁶github.com/nf-core/atacseq

¹⁷github.com/nf-core/bacass

Table 2: The results from applying the infrastructure profiling on the six different nodes.

Machine	# CPUs	Memory	Storage	CPU events/s	LINPACK	RAM score	read IOPs	write IOPS
Local	8	16 GB	HDD	458	3,959,800	18,700	414	415
A1	2 x 4	32 GB	HDD	223	-	11,000	306	301
A2	2 x 4	32 GB	HDD	223	-	11,000	341	336
N1	8	16 GB	HDD	369	3,620,426	13,400	481	483
N2	8	16 GB	HDD	468	4,045,289	17,000	481	483
C2	8	32 GB	HDD	523	4,602,096	18,900	481	483

Table 3: The workflows used in our experiments with their input data and key characteristics.

Workflow	# Abstract Task Definitions	Sample	Size	Uncompressed Size	Workflow Runtime One Input
Eager	13	1	1.52 GB	8.33 GB	148 min
		2	4.34 GB	25.71 GB	211 min
Methylseq	8	1	3.61 GB	17.03 GB	90 min
		2	4.75 GB	22.50 GB	117 min
Chipseq	14	1	1.33 GB	4.81 GB	140 min
		2	8.71 GB	32.98 GB	948 min
Atacseq	14	1	3.26 GB	14.09 GB	184 min
		2	2.40 GB	11.81 GB	104 min
Bacass	5	1	1.23 GB	3.64 GB	237 min
		2	1.45 GB	4.35 GB	253 min

The Naive Approach estimates the ratio $r_t = \frac{run_q}{d_q}$ for each training data tuple q (uncompressed input size d_q , runtime run_q) and takes the mean \bar{r}_t for task t over these ratios. It then uses this mean ratio to predict the runtime of a task t with uncompressed target input size d_t , using $\bar{r}_t * d_t$. Online-P and Online-M use density-based clustering to identify high-density areas. Then, a cluster is determined according to the I/O read value of the estimated task. Since the clustering is not possible with the sparse data from the local executions, we take the data point closest to the task being estimated. Then, a pearson correlation between all input and output parameters is calculated. If the data correlates, the ratio between output and input parameter is computed and used for the prediction. If the data is uncorrelated, Online-M directly estimates the mean, while Online-P first tries to sample from a Normal or Gamma distribution. Both approaches monitor the workflow execution and can update the estimates as more information becomes available. However, this is not implemented since we focus on the out-of-the-box prediction accuracy without historical data.

5 EVALUATION RESULTS

We run three types of experiments. First, we test different downsampling combinations and sizes to evaluate their impact on Lotaru’s prediction errors. In the second experimental scenario, we train the task models on a similar machine as the target environment to get an unbiased view of the prediction models’ capabilities. The third scenario evaluates the setup in the heterogeneous target cluster. Here, we have to adjust the runtime predictions from our local machine to the different nodes in the target cluster using the adjustment factor.

For our evaluation, we introduce the median prediction error (MPE). This metric is calculated for every workflow and aggregates

the prediction error of the tasks inside the workflow. To this end, we compute the prediction error for a single task as:

$$err_t = abs \left(\frac{predicted_runtime - actual_runtime}{actual_runtime} \right), \quad (7)$$

5.1 Impact of the Downsampling on Prediction Accuracy

Lotaru takes one of the many workflow inputs and downsamples or slices this input for local workflow execution to gather training data. This is a crucial step in our approach since the number and sizes of the chosen partitions highly influence the prediction error and the local workflow runtime.

Therefore, we first want to evaluate how many samples Lotaru should create from one original input and which sizes in relation to it are necessary to achieve good prediction results. The experiment results can be generalized for genomic workflows. Other domains, such as remote-sensing or material science, need a separate analysis. Consequently, we designed our experiment the following:

For each pair of workflow and input, we cut one of the original input files with a size of X into ten partitions for the Eager, Methylseq, Atacseq, and Bacass workflow and 16 partitions for the Chipseq workflow. The size of the first partition, s_1 , is set to $s_1 = \frac{X}{2}$, and $s_n = \frac{s_{n-1}}{2}$, so that s_1 has half, s_2 a quarter of the original size, and so on. Then, we apply our prediction model to all possible partition combinations. Therefore, ten possible input partitions, result in $\sum_{k=2}^{10} \frac{10!}{k!(10-k)!} = 1,013$ combinations for each task and prediction method.

Due to the high number of evaluated workflows with two different datasets used for each workflow, we decided to highlight

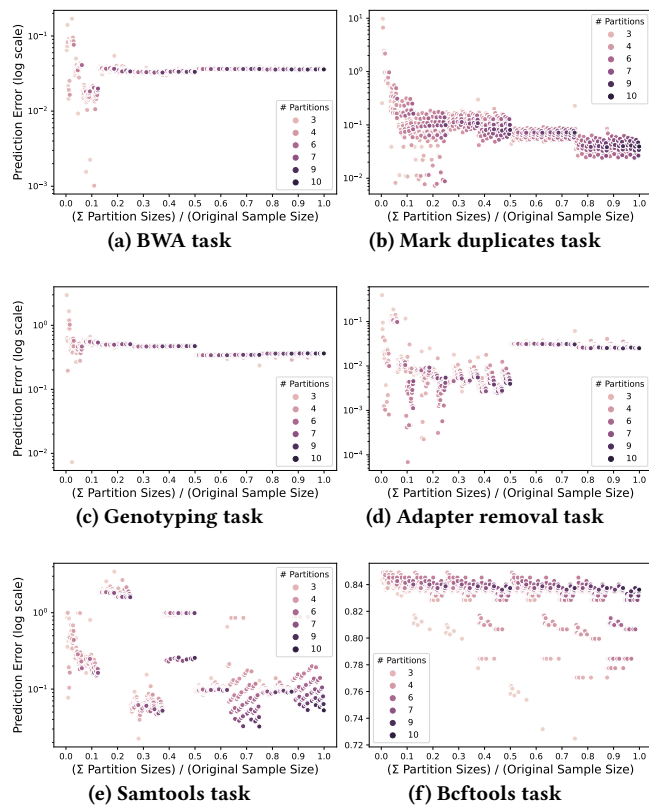


Figure 4: Relationship between the number and the cumulated size of the downsampled partitions and the prediction error for Lotaru and various tasks from the Eager-1 workflow.

the Eager-1 workflow as the example workflow but made all traces available online¹⁸. Therefore, Figure 4 shows Lotaru’s prediction error for six representative tasks from the Eager-1 workflow, depending on the number of the downsampled partitions and their cumulative input size. First, one can see that if the cumulative input size is below 10% of the original input size, the predictions yield a high variance which is reflected in the prediction error. Further, once an input combination reaches this threshold, the number of partitions does not affect prediction error significantly, as long as there are at least three partitions.

Out of the 13 abstract tasks from the Eager workflow, eleven tasks follow this pattern, the samtools tasks shows no relation, Figure 4e, and the bcftools task is predicted according to the median, Figure 4f. Additionally, our observations show that increasing the cumulative size can further reduce the prediction error but also increase the local runtime. Examples for this are the tasks from the Figures 4b and 4c, while the other figures do not follow this pattern. The results from our other workflow executions underline our observations.

At the same time, the cumulative input data size for the local workflow execution determines the runtime on the developer’s

¹⁸github.com/CRC-FONDA/Lotaru-traces

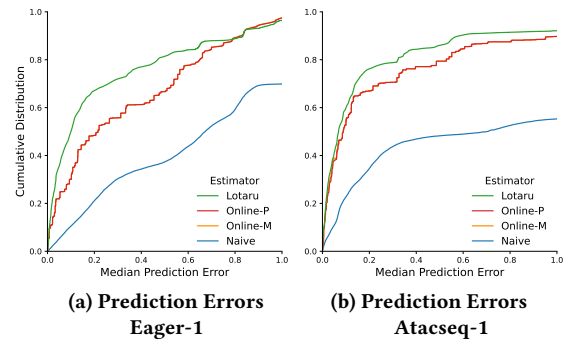


Figure 5: Cumulative Distribution of the prediction error from the four approaches for two workflows. Red line similar to the orange line.

computer. Most tasks have a linear relationship between input data size s and runtime r . Thus, we can extrapolate the runtime r_s from one original input sample.

As examples, Figure 5 shows the cumulative distribution of the prediction error for all tasks of the complete Eager-1 and Atacseq-1 workflow, and compares Lotaru to the three other approaches. The Atacseq-1 workflow is a setup where the baselines Online-P and Online-M show a similar error distribution for around 38% of all combinations compared to Lotaru. The plot line for Online-M is mostly covered by the plot line of Online-P as both perform similarly. The main difference between both approaches is handling uncorrelated relationships between input data size and task runtime. Here, Online-M estimates a runtime according to the median, whereas Online-P considers a statistical distribution. However, our data shows that nearly all tasks have correlated data, leading to similar results.

In Figure 5a, Lotaru has an $MPE \leq 10.00\%$ for 50% of all combinations, while for Online-M and Online-P 50% of all combinations have an $MPE \leq 21.60\%$. The naive approach has the highest prediction error with more than 30.12% of all combinations having an $MPE > 1$.

In comparison to this, Figure 5b exhibits slight changes between Online-M and Online-P. Further, it shows a high slope at the beginning for Online-P, Online-M, and Lotaru. In this region, all three approaches score similarly, however, higher prediction errors occur less frequently for Lotaru.

Regardless of our results, in the following experiments, we select all ten partitions for the prediction models to provide a fair comparison between Lotaru and the baselines.

5.2 Predictions for a Homogeneous Cluster

In our next experiment, we investigate the prediction performance without having to cope with heterogeneity in the cluster. Therefore, we assume that the local machine is similar to the target nodes and that no model adjustment is needed. Figure 6 shows results for each workflow separately. One can see that Lotaru outperforms the three baselines, achieving an MPE over all workflow tasks of 5.70%, while the best performing baseline, Online-P, has an MPE of 10.34%. The largest difference can be observed for the Eager-2 workflow,




Figure 6: Prediction Errors for a homogeneous cluster where no model adjustments are required.

Table 4: Median difference between the actual factor and the factor calculated through Lotaru for Eager-1.

Node	A1	A2	N1	N2	C2
Median Factor Difference	0.15	0.14	0.17	0.06	0.03

where Lotaru achieves an MPE of 9.54% while Online-P results in an MPE of 19.40%. An exception is the Atacseq-1 workflow, where Online-M and Online-P can achieve a lower MPE of 4.27% compared to Lotaru’s 6.03%, however, our max error is 55.00% lower. In three out of five workflows, Lotaru achieves a lower max error, while for two workflows similar max error values are achieved.

5.3 Model Adjustment

A critical point in Lotaru’s approach is the mapping of the local predictions to the heterogeneous target nodes. With the model adjustment, we want to adapt the predicted runtime for hardware differences on the target machines. First, Table 4 compares the differences between the actual runtime factor and the factor Lotaru calculated.

The term difference refers to the absolute difference between actual factor and calculated factor. One can see that for C2 and N2 the difference is the lowest. This is expected since our infrastructure profiling showed that these nodes are the closest to the local node regarding performance characteristics. In contrast, a difference of 0.14 or 0.15 for A1 and A2 seems high, however, both machines have a significant difference in the actual hardware characteristics

and thus, only score half of the CPU events/s and much lower I/O values.

For Table 5, we continue to take the Eager-1 workflow as an example. The table compares the estimated adjustment factor with the actual factor between the local machine and the C2 machine for all 13 tasks in the Eager-1 workflow. One can see that for only three out of 13 tasks, the difference is greater than 0.05, while the median difference between our estimated factor and the actual ratio is 0.03.

The task fastqc from Table 5 is a common bioinformatics task that spots potential problems in sequences data and occurs in all five workflows. Therefore, we chose the task fastqc for our analysis across all machines.

The actually calculated factor for fastqc on C2 shows a difference of only 0.03 and N2 shows an even lower difference of 0.02. Similarly as in Table 4, again, the outlier is node A1. Node A1 yields a difference of 0.31, whereby one has to consider that the actual factor of 2.37 is much higher than for the other machines. Therefore, an absolute difference of 0.31 for an actual factor of 2.37 corresponds to a relative difference of 13.08%.


Concluding, our used adjustment strategy is able to accurately reflect hardware differences.

5.4 Predictions for a Heterogeneous Cluster

In the third experiment, we use our local machine and predict the runtimes for all target nodes A1, A2, N1, N2, and C2 for all tasks in all five evaluation workflows. Table 6 gives an overview of the median prediction errors from all four prediction approaches. Over

Table 5: Comparison of the calculated adjustment factor and the actual runtime factor between the local and the C2 machine for all tasks of the Eager-1 workflow.

Task	bwa	bcftools_stats	samtools_f_a_f	damageprofiler	preseq	genotyping_hc	fastqc_a_c
Actual Factor	0.87	0.79	0.88	0.82	0.84	0.83	0.92
Calculated Factor	0.87	0.87	0.87	0.87	0.87	0.87	0.87
Task	fastqc	samtools_flag	samtools_filter	markduplicates	qualimap	adapter_rem	
Actual Factor	0.90	0.84	0.65	0.86	0.76	0.83	
Calculated Factor	0.87	0.87	0.88	0.88	0.86	0.87	

**Figure 7: Prediction Errors for target node C2 with the influence of our model adjustment, logarithmic scale.****Table 6: MPE for all approaches on all machines over the five experiment workflows.**

Node	A1	A2	N1	N2	C2
Naive	53.11%	52.65%	58.53%	73.01%	83.10%
Online-M	41.82%	39.96%	20.21%	18.40%	30.58%
Online-P	41.82%	39.91%	20.20%	18.40%	30.43%
Lotaru	21.71%	19.91%	14.19%	13.80%	14.62%

all workflow tasks and the five target nodes, Lotaru achieves a median prediction error of 15.99% compared to 30.90% for Online-P, which constitutes a prediction error reduction of 48.25%. The differences in prediction errors are thus much more considerable for the more realistic case of heterogeneous clusters than for a homogeneous cluster.

Further, one can see that for N2, the node closest to the local machine regarding profiled performance characteristics, the estimation error of Lotaru and Online-P is the lowest for all target machines. The error increases for both approaches once the machines differ more, i.e., the error for A1 and A2 is much higher than for N1 or N2. However, while the differences regarding the prediction error between Lotaru and Online-P for N1 and N2 are rather small, they increase for A1 and A2.

Figure 7 gives more insights into the prediction errors for the example target node C2. Again, Lotaru outperforms all three baselines. The best one, Online-P, achieves an MPE of 30.43% compared to 14.62% for Lotaru. Lotaru’s 75th percentile is always below the mean of the other predictors and two times below the median. Regarding the 90th and 95th percentile, Lotaru consistently achieves the lowest prediction error. Additionally, our predictions yield a lower standard deviation for all workflows. Figure 7 shows that for four out of the five workflows, Lotaru’s minimum prediction error

is lower than the prediction error of any baseline. In three out of five workflows the minimum prediction error is even three times lower. The maximum errors are always below the maximum values of any baseline.

6 CONCLUSION

This paper presented Lotaru, a system that estimates the runtime of scientific workflow tasks on the developer’s machine. To this end, Lotaru profiles the target infrastructure with microbenchmarks, reduces the input data to execute the workflow locally quickly, and estimates the runtime with a Bayesian regression based on the gathered data points. We presented an implementation of our approach, while the code and the experimental data have been made available online. Further, our introduced interface is easily extendable to other domains with other input data like images in remote sensing or astrophysics.

Lotaru works for all workflows that consist of multiple input files, which are, at least partly, processed separately and can be downsampled. In contexts where such a downsampling is not possible, Lotaru could also use different subsets of the input files at the price of a somewhat longer training phase. Further, since most big data analysis tasks, such as scans, show a linear input-runtime relation, Lotaru also assumes this linear relation.

Our evaluation with five real-world bioinformatics workflows on different data inputs shows that Lotaru’s local estimation approach achieves low prediction errors and outperforms classical runtime prediction baselines. For predictions where the target node is equal to the machine where the local profiling ran, Lotaru achieves a median prediction error over all workflows of 5.70%, while the best performing baseline achieves 10.34%. Our comparison for the runtime predictions on five target nodes that differ from the local machine shows a median error over all workflows of 15.99% compared to 30.90% for the best competitor, a reduction of the prediction error of 48.25%.

In the future, we plan to extend our implementation to different domains and adopt existing schedulers to consider our prediction’s confidence and uncertainty values.

ACKNOWLEDGMENTS

Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) as FONDA (Project 414984028, SFB 1404).

REFERENCES

- [1] Hamid Arabnejad and Jorge Barbosa. 2012. Fairness resource sharing for dynamic workflow scheduling on heterogeneous systems. In *2012 IEEE 10th International Symposium on Parallel and Distributed Processing with Applications*. IEEE.
- [2] Jonathan Bader, Lauritz Thamsen, Svetlana Kulagina, Jonathan Will, Henning Meyerhenke, and Odej Kao. 2021. Tarema: Adaptive Resource Allocation for Scalable Scientific Workflows in Heterogeneous Clusters. In *BigData*. IEEE.
- [3] Jorge G Barbosa and Belmiro Moreira. 2011. Dynamic scheduling of a batch of parallel task jobs on heterogeneous clusters. *Parallel computing* 37, 8 (2011).
- [4] G Bruce Berriman, Ewa Deelman, John C Good, Joseph C Jacob, Daniel S Katz, Carl Kesselman, Anastasia C Laity, Thomas A Prince, Gurmeet Singh, and Mei-Hu Su. 2004. Montage: a grid-enabled engine for delivering custom science-grade mosaics on demand. In *Optimizing scientific return for astronomy through information technologies*, Vol. 5493. SPIE, 221–232.
- [5] Brendan Burns, Brian Grant, David Oppenheimer, Eric Brewer, and John Wilkes. 2016. Borg, omega, and kubernetes. *Queue* 14, 1 (2016).
- [6] Marc Bux, Jörgen Brandt, Carsten Lipka, Kamal Hakimzadeh, Jim Dowling, and Ulf Leser. 2015. SAASFEE: scalable scientific workflow execution engine. *Proceedings of the VLDB Endowment* 8, 12 (2015).
- [7] Marc Bux, Jörgen Brandt, Carl Witt, Jim Dowling, and Ulf Leser. 2017. Hi-way: Execution of scientific workflows on hadoop yarn. In *20th International Conference on Extending Database Technology, EDBT 2017, 21 March 2017 through 24 March 2017*. OpenProceedings.org, 668–679.
- [8] Rafael Ferreira da Silva, Henri Casanova, Anne-Cécile Orgerie, Ryan Tanaka, Ewa Deelman, and Frédéric Suter. 2020. Characterizing, modeling, and accurately simulating power and energy consumption of i/o-intensive scientific workflows. *Journal of computational science* 44 (2020).
- [9] Rafael Ferreira Da Silva, Gideon Juve, Ewa Deelman, Tristan Glatard, Frédéric Desprez, Douglas Thain, Benjamin Tovar, and Miron Livny. 2013. Toward fine-grained online task characteristics estimation in scientific workflows. In *Proceedings of the 8th Workshop on Workflows in Support of Large-Scale Science*. 58–67.
- [10] Rafael Ferreira Da Silva, Gideon Juve, Mats Rynge, Ewa Deelman, and Miron Livny. 2015. Online task resource consumption prediction for scientific workflows. *Parallel Processing Letters* 25, 03 (2015).
- [11] Yanyan Dai and Xiangli Zhang. 2014. A synthesized heuristic task scheduling algorithm. *The Scientific World Journal* 2014 (2014).
- [12] Ewa Deelman, Karan Vahi, Gideon Juve, Mats Rynge, Scott Callaghan, Philip J Maechling, Rajiv Mayani, Weiwei Chen, Rafael Ferreira Da Silva, Miron Livny, et al. 2015. Pegasus, a workflow management system for science automation. *Future Generation Computer Systems* 46 (2015).
- [13] Ewa Deelman, Karan Vahi, Mats Rynge, Rajiv Mayani, Rafael Ferreira da Silva, George Papadimitriou, and Miron Livny. 2019. The evolution of the pegasus workflow management software. *Computing in Science & Engineering* 21, 4 (2019).
- [14] Paolo Di Tommaso, Maria Chatzou, Evan W Floden, Pablo Prieto Barja, Emilio Palumbo, and Cedric Notredame. 2017. Nextflow enables reproducible computational workflows. *Nature biotechnology* 35, 4 (2017).
- [15] Jack J Dongarra, Piotr Luszczek, and Antoine Petit. 2003. The LINPACK benchmark: past, present and future. *Concurrency and Computation: practice and experience* 15, 9 (2003), 803–820.
- [16] Kalka Dubey, Mohit Kumar, and SC Sharma. 2018. Modified HEFT algorithm for task scheduling in cloud environment. *Procedia Computer Science* 125 (2018).
- [17] Philip A Ewels, Alexander Peltzer, Sven Fillinger, Harshil Patel, Johannes Alneberg, Andreas Wilm, Maxime Ulysse Garcia, Paolo Di Tommaso, and Sven Nahnsen. 2020. The nf-core framework for community-curated bioinformatics pipelines. *Nature biotechnology* 38, 3 (2020), 276–278.
- [18] Dror G Feitelson. 2015. *Workload modeling for computer systems performance evaluation*. Cambridge University Press.
- [19] David Frantz. 2019. FORCE—Landsat+ Sentinel-2 analysis ready data and beyond. *Remote Sensing* 11, 9 (2019), 1124.
- [20] Maxime Garcia, Szilveszter Juhos, Malin Larsson, Pall I Olason, Marcel Martin, Jesper Eisfeldt, Sebastian DiLorenzo, Johanna Sandgren, Teresita Diaz De Ståhl, Philip Ewels, et al. 2020. Sarek: A portable workflow for whole-genome sequencing analysis of germline and somatic variants. *F1000Research* 9 (2020).
- [21] Muhammad Hafizhuddin Hilman, Maria Alejandra Rodriguez, and Rajkumar Buyya. 2018. Task runtime prediction in scientific workflows using an online incremental learning approach. In *2018 IEEE/ACM 11th International Conference on Utility and Cloud Computing (UCC)*. IEEE, 93–102.
- [22] Adán Hiraes-Carbajal, Andrei Tchernykh, Ramin Yahyapour, José Luis González-García, Thomas Röblitz, and Juan Manuel Ramírez-Alcaraz. 2012. Multiple workflow scheduling strategies with user run time estimates on a grid. *Journal of Grid Computing* 10, 2 (2012), 325–346.
- [23] Alexey Ilyushkin and Dick Epema. 2018. The impact of task runtime estimate accuracy on scheduling workloads of workflows. In *2018 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGRID)*. IEEE, 331–341.
- [24] Johannes Köster and Sven Rahmann. 2012. Snakemake—a scalable bioinformatics workflow engine. *Bioinformatics* 28, 19 (2012), 2520–2522.
- [25] Sik-Yum Lee and Xin-Yuan Song. 2004. Evaluation of the Bayesian and maximum likelihood approaches in analyzing structural equation models with small sample sizes. *Multivariate behavioral research* 39, 4 (2004), 653–686.
- [26] Fabian Lehmann, David Frantz, Sören Becker, Ulf Leser, and Patrick Hostert. 2021. FORCE on Nextflow: Scalable Analysis of Earth Observation data on Commodity Clusters. In *Proceedings of the CIKM 2021 Workshops*. Online.
- [27] Philip Maechling, Ewa Deelman, Li Zhao, Robert Graves, Gaurang Mehta, Nitin Gupta, John Mehlinger, Carl Kesselman, Scott Callaghan, David Okaya, et al. 2007. SCEC CyberShake workflows—automating probabilistic seismic hazard analysis calculations. In *Workflows for e-Science*. Springer.
- [28] Andréa Matsunaga and José AB Fortes. 2010. On the use of machine learning to predict the time and resources consumed by applications. In *2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*. IEEE, 495–504.
- [29] Daniel McNeish. 2016. On using Bayesian methods to address small sample problems. *Structural Equation Modeling: A Multidisciplinary Journal* 23, 5 (2016).
- [30] Farrukh Nadeem, Danial Alghazzawi, Abdulfattah Mashat, Khalid Fakeeh, Abdullah Almalaise, and Hani Hagraas. 2017. Modeling and predicting execution time of scientific workflows in the grid using radial basis function neural network. *Cluster Computing* 20, 3 (2017), 2805–2819.

- [31] Tom Oinn, Matthew Addis, Justin Ferris, Darren Marvin, Martin Senger, Mark Greenwood, Tim Carver, Kevin Glover, Matthew R Pocock, Anil Wipat, et al. 2004. Taverna: a tool for the composition and enactment of bioinformatics workflows. *Bioinformatics* 20, 17 (2004), 3045–3054.
- [32] Thanh-Phuong Pham, Juan J Durillo, and Thomas Fahringer. 2017. Predicting workflow task execution time in the cloud using a two-stage machine learning approach. *IEEE Transactions on Cloud Computing* 8, 1 (2017), 256–268.
- [33] Seyed Masoud Sadjadi, Shu Shimizu, Javier Figueroa, Raju Rangaswami, Javier Delgado, Hector Duran, and Xabriel J Collazo-Mojica. 2008. A modeling approach for estimating execution time of long-running scientific applications. In *2008 IEEE International Symposium on Parallel and Distributed Processing*. IEEE, 1–8.
- [34] Joerg Schaarschmidt, Jie Yuan, Timo Strunk, Ivan Kondov, Sebastiaan P Huber, Giovanni Pizzi, Leonid Kahle, Felix T Bölle, Ivano E Castelli, Tejs Vegge, et al. 2021. Workflow Engineering in Materials Design within the BATTERY 2030+ Project. *Advanced Energy Materials* (2021), 2102638.
- [35] Dominik Scheinert, Alireza Alamgiralem, Jonathan Bader, Jonathan Will, Thorsten Wittkopp, and Lauritz Thamsen. 2021. On the Potential of Execution Traces for Batch Processing Workload Optimization in Public Clouds. In *2021 IEEE International Conference on Big Data (Big Data)*. IEEE, 3113–3118.
- [36] Dominik Scheinert, Lauritz Thamsen, Houkun Zhu, Jonathan Will, Alexander Acker, Thorsten Wittkopp, and Odej Kao. 2021. Bellamy: Reusing performance models for distributed dataflow jobs across contexts. In *2021 IEEE International Conference on Cluster Computing (CLUSTER)*. IEEE, 261–270.
- [37] Helge S Stein and John M Gregoire. 2019. Progress and prospects for accelerating materials science with automated and autonomous workflows. *Chemical Science* 10, 42 (2019), 9640–9649.
- [38] Haluk Topcuoglu, Salim Hariri, and Min-you Wu. 2002. Performance-effective and low-complexity task scheduling for heterogeneous computing. *IEEE transactions on parallel and distributed systems* 13, 3 (2002).
- [39] Dave Turner, Dan Andresen, Kyle Hutson, and Adam Tygart. 2018. Application Performance on the Newest Processors and GPUs. In *Proceedings of the Practice and Experience on Advanced Research Computing*. ACM.
- [40] Vinod Kumar Vavilapalli, Arun C Murthy, Chris Douglas, Sharad Agarwal, Mahadev Konar, Robert Evans, Thomas Graves, Jason Lowe, Hitesh Shah, Siddharth Seth, et al. 2013. Apache hadoop yarn: Yet another resource negotiator. In *Proceedings of the 4th annual Symposium on Cloud Computing*. 1–16.
- [41] Guan Wang, Yuxin Wang, Hui Liu, and He Guo. 2016. HSIP: A novel task scheduling algorithm for heterogeneous computing. *Scientific Programming* 2016 (2016).
- [42] Jonathan Will, Lauritz Thamsen, Dominik Scheinert, Jonathan Bader, and Odej Kao. 2021. C3O: Collaborative Cluster Configuration Optimization for Distributed Data Processing in Public Clouds. In *2021 IEEE International Conference on Cloud Engineering (IC2E)*. IEEE, 43–52.
- [43] Carl Witt, Marc Bux, Wladislaw Gusew, and Ulf Leser. 2019. Predictive performance modeling for distributed batch processing using black box monitoring and machine learning. *Information Systems* 82 (2019), 33–52.
- [44] Carl Witt, Jakob van Santen, and Ulf Leser. 2019. Learning low-wastage memory allocations for scientific workflows at iccube. In *2019 International Conference on High Performance Computing & Simulation (HPCS)*. IEEE, 233–240.
- [45] Carl Witt, Dennis Wagner, and Ulf Leser. 2019. Feedback-Based Resource Allocation for Batch Scheduling of Scientific Workflows. In *2019 HPCS*. IEEE.
- [46] James A Fellows Yates, Theseas C Lamnidis, Maxime Borry, Aida Andrades Valtueña, Zandra Fagernäs, Stephen Clayton, Maxime U Garcia, Judith Neukamm, and Alexander Peltzer. 2021. Reproducible, portable, and efficient ancient genome reconstruction with nf-core/eager. *PeerJ* 9 (2021), e10947.
- [47] Andy B Yoo, Morris A Jette, and Mark Grondona. 2003. Slurm: Simple linux utility for resource management. In *Workshop on Job Scheduling Strategies for Parallel Processing*. Springer.