

Optimization of Tomographic Reconstruction Workflows on Geographically Distributed Resources

TEKIN BICER,^{a*} DOGA GURSOY,^b RAJKUMAR KETTIMUTHU,^{a,c} FRANCESCO DE CARLO^b AND IAN T. FOSTER^{a,c,d}

^a*Mathematics and Computer Science Division, Argonne National Laboratory, USA.*

E-mail: bicer@anl.gov, ^bAdvanced Photon Source, Argonne National Laboratory,

USA, ^cComputation Institute, University of Chicago and Argonne National Laboratory, USA, and ^dDepartment of Computer Science, University of Chicago USA

Abstract

New technological advancements in light sources enable data acquisitions at unprecedented levels. This emergent trend affects not only the size of the generated data but also the need for larger computational resources. Although beamline scientists and users have access to local computational resources, these are typically limited and can result in extended execution times. Applications such as iterative tomographic reconstruction require high-performance compute clusters for timely analysis of data.

We focus here on time-sensitive analysis and processing of Advanced Photon Source data on geographically distributed resources. We consider two main challenges: (1) modeling of the performance of tomographic reconstruction workflows and (2) transparent execution of these workflows on distributed resources. For the former, we consider three main stages: (1) *data transfer* between storage and computational

resources, (2) *wait/queue* time of reconstruction jobs at compute resources, and (3) *computation* of reconstruction tasks. These performance models let us evaluate and estimate the execution time of any given iterative tomographic reconstruction workflow that runs on geographically distributed resources. For the latter challenge, we build a workflow management system, which can automate the execution of workflows and minimize the user interaction with the underlying infrastructure. Our system utilizes Globus to perform secure and efficient data transfer operations.

The proposed models and the workflow management system are evaluated by using three high-performance computing and two storage resources, all of which are geographically distributed. We created workflows with different computational requirements using two compute-intensive tomographic reconstruction algorithms. Our experimental evaluation shows that the proposed models and system can be used for selecting the optimum resources, which in turn can provide up to 3.13x speedup (on experimented resources). Moreover, the error rates of the models range between 2.1–23.3% (considering workflow execution times), where the accuracy of the model estimations increases with higher computational demands in reconstruction tasks.

1. Introduction

Advances in detector technologies enable increasingly complex experiments and more rapid data acquisition at synchrotron light sources [11]. Current full-field x-ray imaging instruments, including microtomography and high-speed imaging, allow scientists to collect a full 3D or 2D dataset at 1-micron resolution in a fraction of a second (as short as 100 picoseconds). Moreover, the next-generation synchrotron light sources will deliver several orders of magnitude higher temporal resolution at the nanometer range, extending *in situ* and time-dependent studies of phenomena to completely new regimes. Similarly, other instruments based on scanning probe techniques will directly

benefit from technological advancements, reducing data collection time by a factor of 100, which will result in at least two orders of magnitude more data generation [1].

Although data acquisition rates at x-ray light sources are increasing rapidly, the analysis of this data often lags behind, mostly because of the high computational demands of processing and a lack of local compute resources and optimized software tools. These factors significantly hinder the ability to perform near-real time analysis and visualization. Because experimental setups and data acquisition approaches at light sources are typically configuration sensitive and scientists need to make timely decisions in order to collect accurate data, short turnaround times for data analysis are crucial. Therefore, efficient and scalable software tools that are tailored to the requirements of synchrotron light sources will be a necessity, especially considering the increased data acquisition rates that will be delivered by the new technological advancements at light sources.

One widely used compute-intensive application at synchrotron light sources is iterative tomographic image reconstruction [15, 18, 19]. This application enables scientists to observe the internal structure of objects from 3D images and therefore helps in understanding time-dependent phenomena. The processing times of iterative reconstruction applications are sensitive to the input dataset size and configuration. Specifically, the number of iterations, 2D projections, sinograms and voxels determine the computational requirements of the reconstruction tasks. Reconstruction of a typical high-resolution tomographic dataset can take days with a well-built workstation [5].

In this work, we focus on understanding the performance issues with iterative tomographic reconstruction workflows and improving their end-to-end execution times on geographically distributed resources. Particularly, our contributions in this paper are as follows.

- We develop performance models for different stages in tomographic reconstruc-

tion workflows.

- We present a workflow management system that can execute light source data analysis tasks on distributed resources. Our system utilizes **Globus** for efficient and secure management of workflows and data transfers.
- We evaluate and show our system’s accuracy and performance using three high-performance computing (HPC) and two storage resources. During our evaluations, we use two real-world datasets and two iterative reconstruction algorithms.

The remainder of this paper is organized as follows. In Section 2 we introduce our performance models, which are used for estimating the execution times of the workflows. In Section 3 we present our workflow management system, in which the execution of tomographic reconstruction workflows is automated on geographically distributed resources. In Section 4 we evaluate the performance models and workflow management system and use real-world configurations on large-scale HPC resources to assess the accuracy of our model estimations. In Section 5 we present related work, and in Section 6 we summarize our conclusions.

2. Performance Model

We model the performance of a tomographic reconstruction workflow in three stages: (1) input and output data transfer between storage and compute resources, *data transfer*; (2) submission of the reconstruction job to computational resource and its *wait/queue time*; and (3) reconstruction/processing of projection data, *computing*. We assume that the resources used for stages (1) and (3) may be geographically distant from each other.

2.1. Online Estimation of Data Transfer Rate

The tomography dataset are typically stored in local storage resources after their acquisition. Although these resources can provide some level of computational throughput, they are generally limited. Therefore, the data needs to be transferred where sufficient computational resources exist. This transfer process can introduce significant overhead depending on the size of dataset and the status of network.

Since the network is a shared resource, the available bandwidth between resources changes dynamically. Therefore, we model the data transfer rates in an online fashion. Specifically, we measure the bandwidth from storage to compute resources right before the execution of the workflow.

Our online bandwidth estimation method relies on the *bandwidth delay product* (BDP) with *transfer initialization cost*. BDP provides (theoretically) the maximum amount of data that can be in transit before it is acknowledged by the destination resource. It is the product of the theoretical bandwidth between resources and the round-trip time of the packets.

Figure 1 depicts our online bandwidth estimation method. Here, lines 1 and 2 calculate the sample dataset sizes, DS_1 and DS_2 , according to the BDP between resources A and B . Then, the *Generate* function creates these files at resource A . In line 4, the generated sample datasets are transferred from A to B ; and according to the transfer times t_1 and t_2 , the transfer initialization overhead is calculated. We calculate the current bandwidth (without transfer initialization overhead) in line 6. Using the estimated bandwidth and initialization cost, we then derive the transfer time of the real dataset with size RD . Note that DS_1 and DS_2 are typically much more smaller than the original dataset size, hence their transfer costs are minimal.

2.2. Queue Time Prediction

HPC clusters and supercomputers typically use a queue-based batch scheduler to improve resource utilization and provide fair share between users. These schedulers organize user submitted tasks/jobs and assign them to the available resources. Depending on the submitted job's resource requirements and the number of jobs in the queue, a job can wait minutes, hours, or even days for resources. Therefore, the queue time can significantly extend the end-to-end execution time of a workflow.

We estimate the queue time of a computational resource using a queue simulator. Specifically, first we take a snapshot of the target cluster's queue, Q . This snapshot includes information about previously submitted jobs. Once this information is acquired, we generate a synthetic job, j , that represents the resource requirements of our reconstruction task. This job then is enqueued into the Q . We simulate this new Q with a queue simulator using one of the available scheduling algorithms and observe j 's scheduling time.

Depending on the scheduling algorithm, simulated scheduling times can vary significantly. To accommodate the worst-case scenario, we use **conservative backfilling** as our default scheduling algorithm [25]. Since we consider the worst-case scenario, we normalize the estimated queue times with a w_s constant. This constant can be derived empirically by scheduling a test job at a target cluster. Specifically, we simulated the current jobs in the target cluster's queue with a test job, which provides a simulation time t_0 . At the same time, we submitted the same job to the target cluster's queue, which resulted in t_1 scheduling time. The ratio between t_0 and t_1 represents w_s . Note that w_s needs to be calculated only once; therefore, its computation is negligible.

Although we set conservative backfilling as our default scheduling algorithm, other algorithms can also be used. For queue simulations, we used pyss [27], a discrete-event simulator that includes 18 different scheduling algorithms.

2.3. Modeling the Performance of Iterative Reconstruction Algorithms

In general, a tomography dataset is a 3D array that consists of 2D projections or images of an object. Each projection is a set of line integrals associated with total attenuation of x-ray beams while they pass through the object. During reconstruction, each sinogram is mapped to a slice in 3D image. The computational complexity of this mapping function depends mostly on the input sinogram size (i.e., number of projections); size of the 3D image slice (or output slice); number of iterations; and properties of reconstruction algorithm, such as communication and data access patterns.

Here, an output slice is typically composed of a grid of size $col \times col$, in which each cell represents a voxel. The reconstruction of a slice requires iterating over the x-ray attenuation values in the sinogram, calculating backward and forward projections, and updating the output voxels. The computational demand of the reconstruction depends on the number of intersection points between x-ray and voxel grids. We can estimate the number of intersection points that needs to be processed for a projection row as follows.

$$inters(\theta_i) = col^2 * (|\sin(\theta_i)| + |\cos(\theta_i)|) \quad (1)$$

Here, θ_i denotes the rotation of the projection, and col^2 is the number of voxels in the output slice.

Although $inters(\theta_i)$ can quantify the computational demands of a given projection row (with θ_i rotation), hardware-specific information needs to be incorporated for more precise estimations—especially if the goal is to estimate *execution time*. We capture hardware-specific information by introducing a cost function, $cost(\theta_i)$, and a constant parameter, t_d .

We use $cost(\theta_i)$ to incorporate overheads due to the data access pattern during reconstruction. For instance, if the rotation of the projection is 0, that is, $\theta = 0$, then the reconstruction algorithm operates on a continuous array of elements (voxels

that are in the same row). This type of processing results in high cache utilization and therefore minimum overhead. If the projection rotation is 90 degrees (or $\theta_i = 90$), however, then reconstruction performs strided accesses to the voxels, which in turn introduce low cache utilization and high overhead. We use Eq. 2 to capture the overhead due to the data access pattern for a given rotation.

$$cost(\theta_i) = 1 + (2 * |sin(\theta_i)|) \quad (2)$$

Equations 1 and 2 can be used for computing the data access overhead and computational load of reconstructing a sinogram; however, the execution time estimation also depends on the number of instructions that can be executed per second by the running CPU. Therefore, we use the t_d constant to capture the required time (in seconds) for processing a single intersection point and its corresponding unit time hardware overhead in a compute node.

Assuming that a tomography dataset is composed of n_s sinograms and that each sinogram (S) consists of projection rows with θ_i , namely, $S = \{\theta_1, \theta_2, \dots, \theta_{|S|}\}$, then we can estimate the reconstruction time of an input tomography dataset using Eq. 3.

$$n_s * \sum_{i=1}^{|S|} inters(\theta_i) * cost(\theta_i) * t_d \quad (3)$$

Note that we can derive all variables and functions in Eqs. 1 and 2 from the input and output datasets. In other words, these parameters and functions can be determined before computation. However, the value of t_d requires reconstruction of a sample sinogram (training data). Fortunately, this process needs to be performed only once for the target hardware (compute node) and typically takes negligible time.

Since our performance model focuses on the computational demand and processing time of a single intersection point, we can estimate the reconstruction times even with large-scale compute resources where a sinogram is being reconstructed by more than one processing unit, such as multicore architectures.

3. Workflow Optimization and System Design

Execution of geographically distributed workflows requires automation tools to ease synchronization and interaction with resources. In this section, we present the components of our system and discuss its execution.

3.1. System Components

Our system consists of four main components: (1) data transfer, (2) metascheduler, (3) performance models, and (4) workflow execution engine. The interactions between these components are shown in Fig. 2 and explained in the following subsections.

3.1.1. Data Transfer: The data transfer component is responsible for initiating data transfer requests during workflow execution. It also monitors the status of transfer and triggers actions according to state of the transmission. For instance, if the data transfer is successful, this component informs the workflow execution engine for the next stage. Otherwise, it dispatches error messages for failure handling.

The metadata information about data transfers and resources is defined in the *resource configuration and metadata* file. This information includes the *type* and (end-point) *address* of the resource, default *input/output directories* for the workflow management system, and the theoretically *available bandwidth*. Here, a resource is a Globus end-point, and the type of resource is either storage or compute. If the resource type is compute, it can provide both computational and storage resources, whereas storage resources can be used only for storing data. Input/output directories are the default folders that are being used by the workflow management system. Typically, the input directory defines where input data is stored (for both storage and compute resources), and the output directory is where output data is stored after the completion of data analysis at the compute resource. The available bandwidth is used for estimating data

transfer rates (using 2.1).

This component also provides an interface for interacting with Globus. The interface can be used for setting the configuration parameters of Globus, including the number of parallel streams, concurrency, and time-out parameters.

3.1.2. Metascheduler for Multisite Environments: Our system utilizes geographically distributed resources to minimize the end-to-end execution time. This process requires interacting with various resources and depends on taking into consideration operations such as resource allocations, preparation, and scheduling of jobs at target system.

The metascheduler component creates an abstraction between the resource managers of the compute clusters and the workflow execution engine. Therefore, jobs can be defined in the workflow execution engine (job description) and materialized to real job requests (scripts) in the metascheduler. During execution, these scripts are transferred to the target compute cluster and submitted to its job queue. A job description consists of job's resource requirements, such as the number of requested processors and allocation time. This information is used for generating a suitable job script for the target computer's resource manager, for example, slurm or torque.

Similar to the data transfer component, information required for job submission is specified in the resource configuration and metadata file. The information includes the type of resource manager and the project name in which the allocation request is going to be charged. Information about the number of nodes, processors, and scheduling algorithm of the target resource is also stored in the resource configuration file so that the performance and queue time estimation can be computed.

3.1.3. Performance Models and Estimations: Our workflow management system uses the models and simulation techniques in Section 2 to quantify the performance of the

resources. Further, it helps in estimating the execution times of different stages. This information can be used for selecting the most efficient resources and minimizing the end-to-end execution time of workflows.

As mentioned earlier, we consider three main stages during the execution of a workflow: data transfer, computation, and queue/wait times. The performance models component evaluates these stages according to given set of resources and user-provided models. Specifically, for the data transfer stage, this component first prepares scripts that generate sample datasets on the target resource. It next initiates the transfer of these sample datasets between resources (using Globus). It then collects the required performance information and parameters and calculates the estimated data transfer time by using a performance model. The computation stage requires using input data size information (i.e., projection data), t_d value of target compute clusters, number of nodes/processes, and application-specific information such as the number of iterations. This information is then used for setting the parameters in the computational model, and execution times for different compute resources are estimated. For the queue/wait time stage, the system first collects the current states of the queues in the target compute clusters. This state information includes all jobs that are currently running or waiting in the clusters' queues. The performance models component simulates each of these queues with the user-specified reconstruction job and estimates the queue times.

Note that the performance models component extensively uses Globus and the resource configuration file for estimations. It also provides an interface in which different light source application models can easily be integrated.

Our system can be used for execution and modeling of large workflows that execute on many geographically distributed resources.

3.1.4. Workflow Execution Engine: The workflow execution engine is the intermediate layer that orchestrates the other components. Users of the workflow management system interact with this layer to initiate workflow executions and performance estimations.

3.2. Execution of Workflow Management System

Figure 3 shows the execution flow of our system, which consists of two main processes (daemons) that manage the execution of workflows: **master daemon** (MD), and **worker daemon** (WD). The user interacts with the MD process and initiates the workflow execution. MD is responsible for running the previously defined components, namely, the data transfer, metascheduler, performance models, and workflow execution engine.

Our workflow management system can interact with any end point that is enabled in Globus. In Fig. 3, we show four different end points: two are compute end points, one is a cloud end point, and one is a storage endpoint. Since storage resources can be managed easily with Globus, WD is executed on compute resources only.

Figure 3 illustrates a sample workflow scenario. Initially, the user interacts with MD and (1) initiates the workflow execution. MD then establishes a connection with Globus and (2) starts transferring data (which needs to be processed) from the storage to the compute end-point. While data is being transferred, MD (3) creates command scripts and initiates another transfer request from MD to the compute resource. After the transmission, WD starts executing the retrieved scripts. The scripts consist of different bash commands, such as job submission, sample data generation (for data transfer estimations), or queue status commands (for queue time estimation). In our workflow example, the transferred script has job submission commands; hence WD enqueues the job to the cluster’s queue. This job is then executed by the cluster’s resource manager. Recall that the metascheduler component creates job submission commands

depending on the target cluster’s environment. Therefore the scripts can transparently be generated by **MD**. Once **WD** finishes executing all the commands in the script, it informs **MD** with a status file. **MD** then initiates another data transfer that (4) sends the processed data from the compute to the data end-point.

4. Experimental Results

In this section, we present the experimental evaluation of our performance models and workflow management system. In particular, we compare and analyze the *estimated* and *real* end-to-end execution times of tomographic reconstruction workflows that run on geographically distributed resources.

4.1. Experimental Setup

In our experiments, we used three compute clusters—**Mira**, **Stampede**, and **Gordon**—for data processing. Table 1 summarizes the technical details of the compute clusters. We also used two storage resources—**Chameleon**¹ and **Petrel**²—for data storage. **Petrel** is located at Argonne National Laboratory and can provide up to 1.7 PB of storage space. **Chameleon** is a *cloud* resource located at the Texas Advanced Computing Center (TACC) and has 1.5 PB storage space.

We evaluated our computation models using two iterative tomographic reconstruction algorithms: *simultaneous iterative reconstruction technique* (SIRT) and *accelerated penalized maximum likelihood* (APMLR). APMLR is computationally more demanding than SIRT because it incorporates information of neighboring voxels during reconstruction. For data transfer operations and estimations, we used GlobusOnline. During the data transfer estimations, we set the sample data size replication to 16 and the communication overhead multiplier to 2. For job queue simulations, we used a *con-*

¹ <https://chameleoncloud.org>

² <http://petrel.alcf.anl.gov>

servative backfilling algorithm and set its weight in the performance models to 0.04.

We used two real-world datasets, **Seed** and **Hornby**, collected from Advanced Photon Source (APS) beamlines. The **Seed** dataset was acquired from a seed of *Arabidopsis thaliana*, a flowering plant [19]; it consists of 720 projections, 512 sinograms, and 501 columns (i.e., 720x2048x501 single-precision floating-point numbers). **Hornby** is an x-ray microtomography data from a shale sample [22]; it includes 360 projections, 2,048 sinograms, and 1,024 columns. During our experiments, we varied the number of projections and sinograms of the datasets in order to introduce varying computational demands for reconstruction.

We assumed that a workflow consists of four stages: (1) transferring the projections dataset (input) from storage to the compute resource; (2) submitting the reconstruction job to the queue and wait; (3) reconstructing the tomographic dataset; and (4) transferring the 3D image (output) from the compute resource to storage. During our experiments, we categorized these stages as *data transfer*, *queue*, and *computation*.

4.2. Evaluation of Performance Models on Geographically Distributed Resources

First, we evaluated the estimated and real execution times of the **Seed** and **Hornby** datasets using the SIRT reconstruction algorithm. During our evaluation, we varied the experimental configurations and created workflows with different performance characteristics.

Figure 4 shows the experimental results for the **Hornby** dataset. For all experiments, **Real** refers to the actual end-to-end execution times of the workflows, whereas **Estimated** refers to the estimated execution times after running the performance models. Different dataset sizes were used in order to introduce varying computational demands; therefore, the accuracy of the performance model estimations can be observed. Specifically, we generated several versions of the **Hornby** dataset and set

their dimensions to $P \times 128 \times 2048$, in which P refers to the number of projections in the tomographic dataset. We also used different configurations for the storage and compute resources, so that varying data transfer and computational costs were introduced. In particular, each workflow, W , was configured as $S_0 - C - S_1$, where C refers to *compute resource*, and S_0 and S_1 are *source* and *destination storage resources*, respectively. Since $C = \{Gordon, Mira, Stampede\}$ in each experiment, we present C with $*$ in Figs. (a)-(d). For instance, in Fig. 4(d), the Hornby dataset's dimensions are $180 \times 128 \times 2048$ ($P=180$), and the dataset is initially stored on Chameleon. For each workflow in (d), first the dataset is transferred from Chameleon to one of the computational resources (C). Then, its reconstruction job is submitted to the compute cluster's queue. After the reconstruction, the output dataset (3D image) is transferred to Petrel. For this set of experiments, we used 128, 32, and 32 compute nodes for Mira, Stampede, and Gordon, respectively.

Considering the estimated and real execution times in Fig. 4, we see that the error rates range from 1% to 18.2%, with an average of 9.8%. The highest error rate is observed in the $P=180$ and $C=Stampede$ configuration. As mentioned before, we consider (and model) three phases in a workflow: data transfer, queue, and computation. Among these, our computational model provides the most accurate estimation, mainly because of the explicit allocation of compute nodes. On the other hand, the data transfer and job queue are shared between users, thus introducing noise and performance fluctuations depending on the utilization of the network and compute resources. Because smaller datasets require less computation, errors in data transfer and queue time estimations become more visible in the overall estimation. Specifically, for the $P=180$ and $C=Stampede$ configuration, the real and estimated computation times are 137 and 136 seconds, respectively. However, the total real time for the queue and data transfer is 197 seconds, whereas the total estimated time is 258 seconds.

Among the compute resources, Stampede provides the highest computational throughput per node. When we analyze the execution times in Fig. 4, however, we see that in many cases configurations with Stampede perform poorly. The main reasons are the longer job queues and data transfer times at Stampede. Note that our performance models can provide estimations that can help in selecting the most efficient workflow configurations. If we compare the real execution times in each figure, the speedups in selecting the best configuration vs. the worst range between 1.41 and 1.95, which indicate that our models can have a significant effect on the end-to-end execution time.

In Fig. 5, we repeat the previous experiments with the Seed dataset. The dimensions of the Seed dataset are $Px512x501$. Since this dataset is smaller than Hornby's, its computational demands are lower. Therefore, we observe larger error rates. Specifically, considering the real and estimated times, we see that the average error rate is 9.5%, where the rates range between 1% and 23.3%. The largest error rate is observed with the $P=90$ and $C=Stampede$ configuration, in which the error rate of the computation time is 6.2%, and that of the queue and data transfer is 28.1%. Most of the remaining configurations' error rates are lower than 14.5%.

Similar to the Hornby experiments, if we compare the estimated and real execution times, we see that our models can help select the best configuration that can provide the shortest end-to-end execution time. Considering the configurations with the best and worst execution times, the speedups of selecting the right configuration range between 1.89 and 3.13.

Next, in Fig. 6, we focus on the workflow execution times of the Hornby dataset, considering different phases and varying the number of compute nodes. The dimensions of the dataset are configured to $180x2048x2048$; that is, all the sinograms of Hornby are reconstructed. The computations are performed on Stampede, and both the source

and target storage resources are set to Chameleon.

Figure 6(a) shows the total execution times. The error rates of the estimated execution times range from 3% to 14.4%. Similar to the previous set of experiments, the estimated times successfully reflect the real execution times. If we focus on the computation times, we observe almost linear scalability; however, the data transfer and queue times do not follow the same trend, which can also be seen in (b). For instance, while the queue time of the 32-node configuration takes 25.3% of the execution time, the queue times of the $N=128$ and $N=256$ configurations are negligible.

If we analyze the data transfer times, we observe similar results for all the workflows. Because the output dataset size is the same for all configurations, the data transfer takes similar times. Specifically, the data transfer times range from 374 to 404 seconds for the real configurations and 431 to 560 seconds for the estimated configurations. Because Stampede and Chameleon are located at the same site (TACC), the data transfer is more isolated, and the network can provide better performance.

5. Related Work

In recent years, data analysis problems and workflows at synchrotron light sources have gained immense popularity, primarily because of the technological advances in scientific instruments and unprecedented data generation rates at beamlines.

Several activities focus on data management and analysis of light sources [3, 21, 26]. CAMERA is an interdisciplinary project at LBNL [14]. The main goal of the CAMERA project is to investigate problems of DOE user facilities, including synchrotron light sources, and develop fundamental new mathematical solutions. Similarly, the Computational Science Initiative at Brookhaven National Laboratory provides workflow systems to define data-processing tasks for NSLS-II facility users [9].

Tomography is a widely used imaging technique at synchrotron light sources [7, 15,

18, 19, 24, 28]. Iterative tomographic reconstruction typically provides higher-quality images than the filtered-back projection method does, because of the better reduction in noise and in missing wedge artifacts. However, iterative approaches typically require more computational throughput [2, 4, 29, 31]. Accelerators, such as GPUs, have been used to improve performance of iterative algorithms [6, 8, 23, 32]; but memory limitations of these accelerators can introduce significant overhead due to the data movement from host to device memory. Gursoy *et al.*'s TomoPy [20] framework provides a complete set of tools for the portable analysis of tomography datasets, which can run on workstations. Our workflow management system uses highly parallel versions of TomoPy reconstruction algorithms [5] to decrease the execution time of reconstruction tasks from days to minutes on high-performance resources. Both tools use data exchange model [10] for organization and portability of tomography data.

Scientific workflows have been extensively researched by different scientific communities [16, 17, 30, 33]. Deelman *et al.*'s Pegasus workflow management system [12] maps abstract workflow descriptions to geographically distributed execution environments. In this work, we analyze and integrate performance models for light source data analysis workflows. Deslippe *et al.* developed the SPOT framework [13] that specifically focuses on end-to-end analysis of light source data. They evaluate the execution of their system at ALS beamlines and the NERSC supercomputer. In contrast, our work uses Globus for data management and workflow execution, and therefore it can utilize every end point that is accessible by the user.

6. Conclusion

We analyzed and modeled the performance of light source data analysis applications. Our models estimate the computation, data transfer, and task/job wait performance and times. Our models are tailored to iterative tomographic image reconstruction,

in which large computational throughput is desired for timely execution. We implemented these models in a workflow management system that executes reconstruction tasks on geographically distributed resources. We evaluated our models and system using two tomographic datasets (from APS beamlines) and two compute-intensive tomographic reconstruction algorithms. Moreover, we used three HPC and two storage resources, all of which are distributed. Our experimental results show that our system can estimate reconstruction workflows within 2.1% to 23.3% error rates. The models provide better accuracy if the task's computation dominates the overall execution time.

Acknowledgements: This material is based upon work supported by the U.S. Department of Energy, Office of Science, under contract number DE-AC02-06CH11357.

References

- [1] Early Science at the Upgraded Advanced Photon Source. Technical report, Argonne National Laboratory, Advanced Photon Source, 2015.
- [2] J. Agulleiro and J.-J. Fernandez. Fast tomographic reconstruction on multicore computers. *Bioinformatics*, 27(4):582–583, 2011.
- [3] M. Basham, J. Filik, M. T. Wharmby, P. C. Y. Chang, B. El Kassaby, M. Gerring, J. Aishima, K. Levik, B. C. A. Pulford, I. Sikharulidze, D. Sneddon, M. Webber, S. S. Dhesi, F. Maccherozzi, O. Svensson, S. Brockhauser, G. Nray, and A. W. Ashton. Data Analysis WorkbeNch (DAWN). *Journal of Synchrotron Radiation*, 22(3):853–858, 2015.
- [4] M. Beister, D. Kolditz, and W. A. Kalender. Iterative reconstruction methods in x-ray CT. *Physica Medica*, 28(2):94–108, 2012.
- [5] T. Bicer, D. Gursoy, R. Kettimuthu, F. De Carlo, G. Agrawal, and I. T. Foster. Rapid tomographic image reconstruction via large-scale parallelization. In *Euro-Par 2015: Parallel Processing*, pages 289–302. Springer Berlin Heidelberg, 2015.
- [6] F. Brun, S. Pacile, A. Accardo, G. Kourousias, D. Dreossi, L. Mancini, G. Tromba, and R. Pugliese. Enhanced and flexible software tools for x-ray computed tomography at the Italian Synchrotron Radiation Facility Elettra. *Fundam. Inform.*, 141:233–243, 2015.
- [7] R.-C. Chen, D. Dreossi, L. Mancini, R. Menk, L. Rigon, T.-Q. Xiao, and R. Longo. *PITRE*: software for phase-sensitive X-ray image processing and tomography reconstruction. *Journal of Synchrotron Radiation*, 19(5):836–845, Sep 2012.

- [8] S. Chilingaryan, A. Kopmann, A. Mirone, and T. dos Santos Rolo. A GPU-based architecture for real-time data assessment at synchrotron experiments. In *Real Time Conference (RT), 2010 17th IEEE-NPSS*, pages 1–8, May 2010.
- [9] Computational Science Initiative, Brookhaven National Laboratory. <https://www.bnl.gov/compisci/c3d/programs/NSLS.php>. [Online accessed November-2015].
- [10] F. De Carlo, D. Grsoy, F. Marone, M. Rivers, D. Y. Parkinson, F. Khan, N. Schwarz, D. J. Vine, S. Vogt, S.-C. Gleber, S. Narayanan, M. Newville, T. Lanzilotti, Y. Sun, Y. P. Hong, and C. Jacobsen. Scientific data exchange: a schema for HDF5-based storage of raw and analyzed data. *Journal of Synchrotron Radiation*, 21(6):1224–1230, 2014.
- [11] M. D. de Jonge, C. G. Ryan, and C. J. Jacobsen. X-ray nanoprobes and diffraction-limited storage rings: opportunities and challenges of fluorescence tomography of biological specimens. *Journal of Synchrotron Radiation*, 21(5):1031–1047, Sep 2014.
- [12] E. Deelman, K. Vahi, G. Juve, M. Rynge, S. Callaghan, P. J. Maechling, R. Mayani, W. Chen, R. F. da Silva, M. Livny, and K. Wenger. Pegasus, a workflow management system for science automation. *Future Generation Computer Systems*, 46:17 – 35, 2015.
- [13] J. Deslippe, A. Essiari, S. J. Patton, T. Samak, C. E. Tull, A. Hexemer, D. Kumar, D. Parkinson, and P. Stewart. Workflow management for real-time analysis of lightsource experiments. In *Proceedings of the 9th Workshop on Workflows in Support of Large-Scale Science*, WORKS '14, pages 31–40, Piscataway, NJ, USA, 2014. IEEE Press.
- [14] J. Donatelli, M. Haranczyk, A. Hexemer, H. Krishnan, X. Li, L. Lin, F. Maia, S. Marchesini, D. Parkinson, T. Perciano, D. Shapiro, D. Ushizima, C. Yang, and J. Sethian. Camera: The center for advanced mathematics for energy research applications. *Synchrotron Radiation News*, 28(2):4–9, 2015.
- [15] D. J. Duke, A. B. Swantek, N. M. Sovis, F. Z. Tilocco, C. F. Powell, A. L. Kastengren, D. Gürsoy, and T. Biçer. Time-resolved x-ray tomography of gasoline direct injection sprays. *SAE International Journal of Engines*, 9(2015-01-1873), 2015.
- [16] J. Frey, T. Tannenbaum, M. Livny, I. Foster, and S. Tuecke. Condor-g: a computation management agent for multi-institutional grids. In *High Performance Distributed Computing, 2001. Proceedings. 10th IEEE International Symposium on*, pages 55–63, 2001.
- [17] J. Goecks, A. Nekrutenko, J. Taylor, et al. Galaxy: a comprehensive approach for supporting accessible, reproducible, and transparent computational research in the life sciences. 2010.
- [18] D. Gürsoy, T. Biçer, J. D. Almer, R. Kettimuthu, S. R. Stock, and F. De Carlo. Maximum a posteriori estimation of crystallographic phases in x-ray diffraction tomography. *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 373(2043):20140392, 2015.
- [19] D. Gürsoy, T. Biçer, A. Lanzilotti, M. G. Newville, and F. De Carlo. Hyperspectral image reconstruction for x-ray fluorescence tomography. *Optics Express*, 23(7):9014–9023, 2015.
- [20] D. Gursoy, F. De Carlo, X. Xiao, and C. Jacobsen. TomoPy: a framework for the analysis of synchrotron tomographic data. *Journal of Synchrotron Radiation*, 21(5):1188–1193, 2014.
- [21] Y. P. Hong, S. Chen, and C. Jacobsen. A new workflow for x-ray fluorescence tomography: MAPSToTomoPy. volume 9592, pages 95920W–95920W–8, 2015.

- [22] W. Kanitpanyacharoen, D. Y. Parkinson, F. De Carlo, F. Marone, M. Stampanoni, R. Mokso, A. MacDowell, and H.-R. Wenk. A comparative study of x-ray tomographic microscopy on shales at different synchrotron facilities: ALS, APS and SLS. *Journal of synchrotron radiation*, 20(1):172–180, 2013.
- [23] A. Mirone, E. Brun, E. Gouillart, P. Tafforeau, and J. Kieffer. The pyhst2 hybrid distributed code for high speed tomographic reconstruction with iterative reconstruction and a priori knowledge capabilities. *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*, 324:41 – 48, 2014. 1st International Conference on Tomography of Materials and Structures.
- [24] K. Mohan, S. Venkatakrishnan, J. Gibbs, E. Gulsoy, X. Xiao, M. De Graef, P. Voorhees, and C. Bouman. TIMBIR: A method for time-space reconstruction from interlaced views. *Computational Imaging, IEEE Transactions on*, PP(99):1–1, 2015.
- [25] A. W. Mu’alem and D. G. Feitelson. Utilization, predictability, workloads, and user runtime estimates in scheduling the ibm sp2 with backfilling. *IEEE Transactions on Parallel and Distributed Systems*, 12(6):529–543, Jun 2001.
- [26] S. Patton, T. Samak, C. Tull, and C. Mackenzie. Spade: Decentralized orchestration of data movement and warehousing for physics experiments. In *Integrated Network Management (IM), 2015 IFIP/IEEE International Symposium on*, pages 1014–1019, May 2015.
- [27] PYSS: Python Scheduling Simulator. <https://code.google.com/p/pyss>. [Online accessed November-2015].
- [28] J. Qi and R. M. Leahy. Iterative reconstruction techniques in emission computed tomography. *Physics in Medicine and Biology*, 51(15):R541, 2006.
- [29] E. Y. Sidky, C.-M. Kao, and X. Pan. Accurate image reconstruction from few-views and limited-angle data in divergent-beam CT. *Journal of X-ray Science and Technology*, 14(2):119–139, 2006.
- [30] I. Taylor, M. Shields, I. Wang, and A. Harrison. The Triana workflow environment: Architecture and applications. In I. Taylor, E. Deelman, D. Gannon, and M. Shields, editors, *Workflows for e-Science*, pages 320–339. Springer, New York, Secaucus, NJ, USA.
- [31] J. Treibig, G. Hager, H. G. Hofmann, J. Hornegger, and G. Wellein. Pushing the limits for medical image reconstruction on recent standard multicore processors. *International Journal of High Performance Computing Applications*, 2012.
- [32] M. Vogelgesang, S. Chilingaryan, T. d. Santos, and A. Kopmann. Ufo: A scalable gpu-based image processing framework for on-line monitoring. In *High Performance Computing and Communication 2012 IEEE 9th International Conference on Embedded Software and Systems (HPCC-ICESS), 2012 IEEE 14th International Conference on*, pages 824–829, June 2012.
- [33] 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*, page gkt328, 2013.

Require: $RTT_{A \rightarrow B}$: Round trip time from resource A to B
 $BW_{A \rightarrow B}$: Outgoing bandwidth from resource A to B c : Sample
 DS : Dataset replication constant
 m : Initialization overhead replication constant ($m > 1$)
 RD : Size of the real dataset that is to be transferred

```

function EstimateDataTransferTime( $RD, RTT_{A \rightarrow B}, BW_{A \rightarrow B}, c, m$ )
1:    $DS_1 := RTT_{A \rightarrow B} * BW_{A \rightarrow B} * c;$ 
2:    $DS_2 := DS_1 * m;$ 
3:    $Generate(DS_1, DS_2, A);$ 
4:    $t_1 := Transfer(DS_1, A, B); t_2 := Transfer(DS_2, A, B);$ 
5:    $Overhead := (m * t_1 - t_2) / (m - 1);$ 
6:    $EST_{BW} := DS_2 / (t_2 - Overhead);$ 
7:   return  $RD / EST_{BW} + Overhead;$ 

```

Figure 1: Algorithm for *online bandwidth estimation*

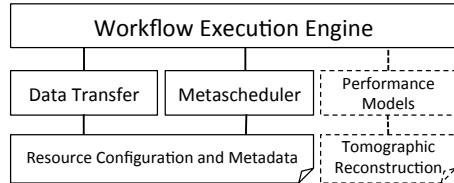


Figure 2: Components of workflow execution engine.

Cluster	Peak Perf.	# Nodes	# Cores	Memory
Mira ³	10 PFLOPS	49,125	786,432	768 TiB
Stampede ⁴	9.6 PFLOPS	6,400	522,080	270 TiB
Gordon ⁵	341 TFLOPS	1,024	16,384	64 TiB

³ Located at ANL: <http://www.alcf.anl.gov/mira>

⁴ Located at TACC: <https://www.tacc.utexas.edu/stampede>

⁵ Located at SDSC: <http://www.sdsc.edu/services/hpc>

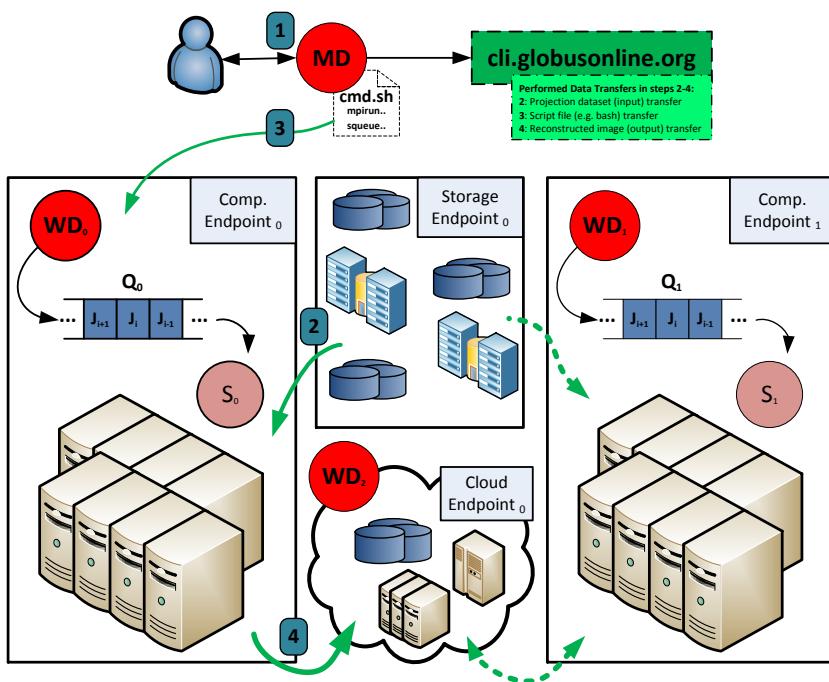


Figure 3: Execution of workflow management system.

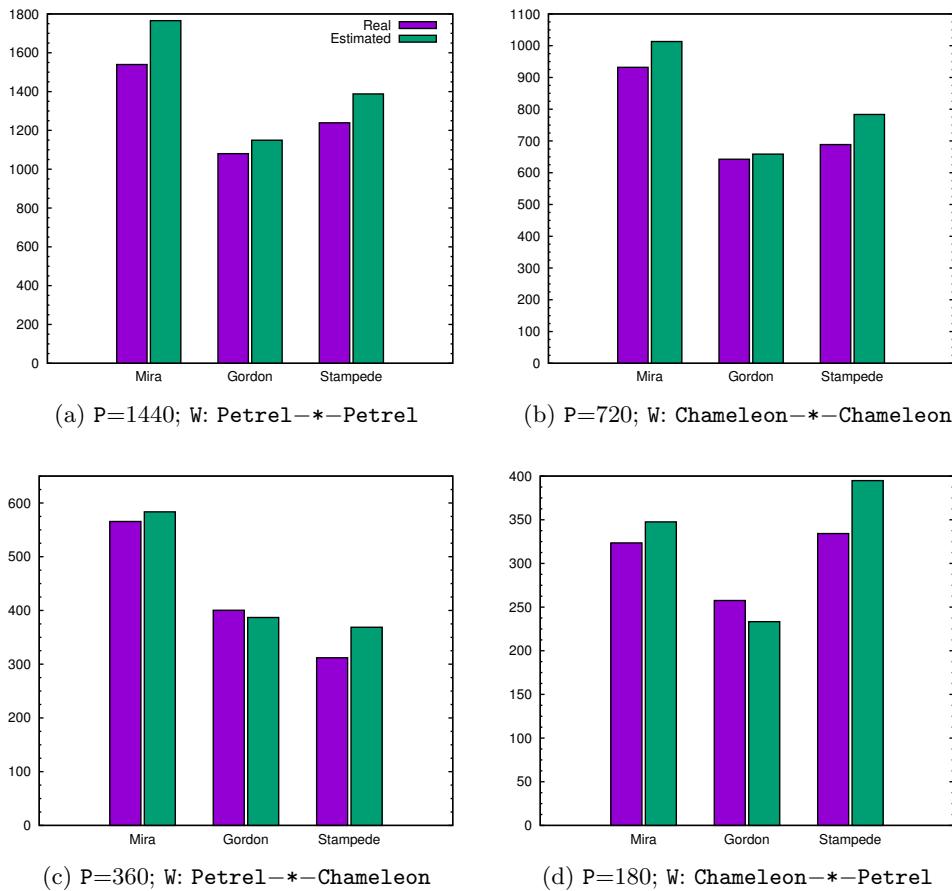


Figure 4: Real and estimated execution times (in seconds) of Hornby image reconstruction workflow using geographically distributed compute and storage resources.

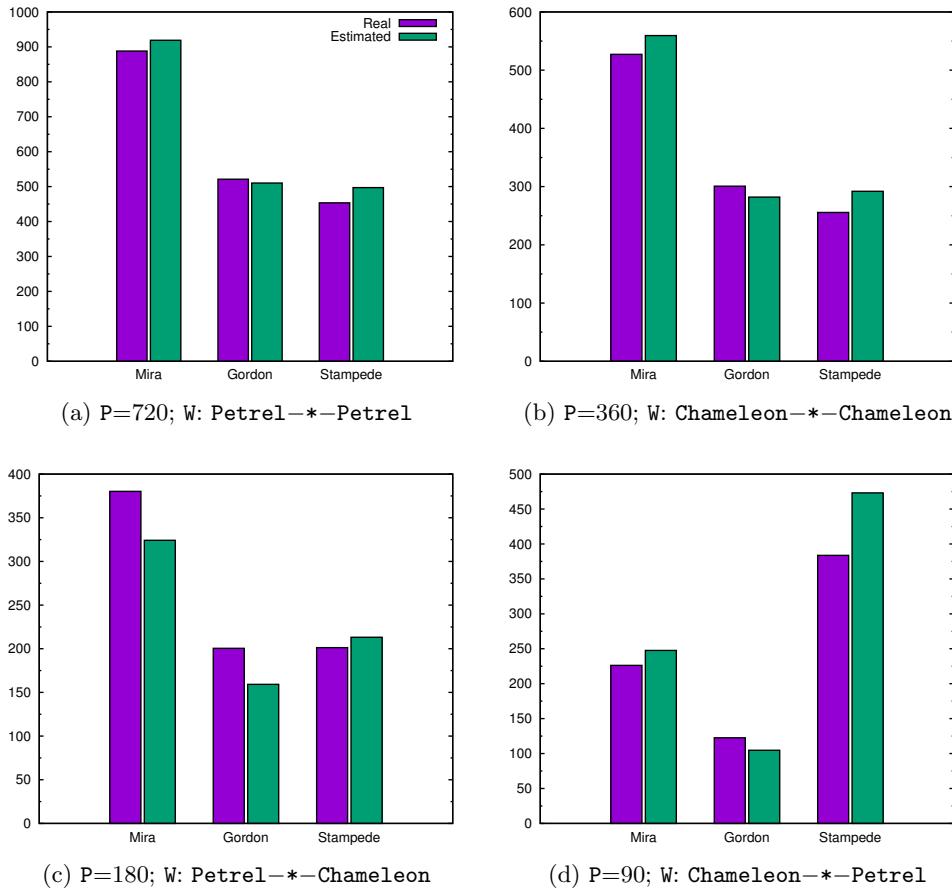


Figure 5: Real and estimated execution times (in seconds) of Seed image reconstruction workflow using geographically distributed compute and storage resources.

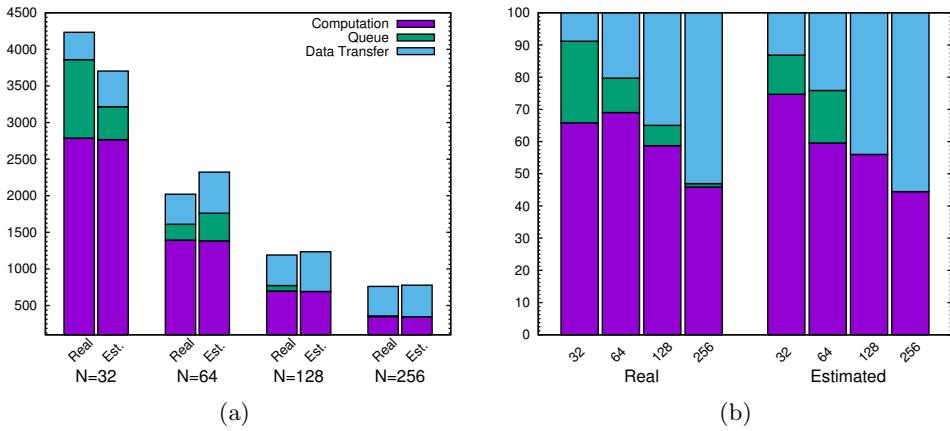


Figure 6: (a) compares real and estimated execution times (in seconds) considering different phases for Hornby dataset with varying number of compute nodes at Stampede. The reconstruction algorithm is APMLR. The dimensions of the dataset are 180x2048x2048. Workflow of the reconstruction is configured as W: Chameleon—Stampede—Chameleon. (b) shows the distribution (%) of phases w.r.t. total execution time grouped in real and estimated configurations.

The submitted manuscript has been created by UChicago Argonne, LLC, Operator of Argonne National Laboratory ("Argonne"). Argonne, a U.S. Department of Energy Office of Science laboratory, is operated under Contract No. DE-AC02-06CH11357. The U.S. Government retains for itself, and others acting on its behalf, a paid-up nonexclusive, irrevocable worldwide license in said article to reproduce, prepare derivative works, distribute copies to the public, and perform publicly and display publicly, by or on behalf of the Government.