AnalyzeThis: An Analysis Workflow-Aware Storage System

Hyogi Sim[†], Youngjae Kim[‡], Sudharshan S. Vazhkudai[‡], Devesh Tiwari[‡],
Ali Anwar[†], Ali R. Butta[†], Lavanya Ramakrishnan^{*}

[†]Virginia Tech, [‡]Oak Ridge National Laboratory, ^{*}Lawrence Berkeley Laboratory

{hyogi,ali,butt}@vt.edu, {kimy1,vazhkudaiss,tiwari}@ornl.gov, lramakrishnan@lbl.gov

ABSTRACT

The need for novel data analysis is urgent in the face of a data deluge from modern applications. Traditional approaches to data analysis incur significant data movement costs, moving data back and forth between the storage system and the processor. Emerging Active Flash devices enable processing on the flash, where the data already resides. An array of such Active Flash devices allows us to revisit how analysis workflows interact with storage systems. By seamlessly blending together the flash storage and data analysis, we create an analysis workflow-aware storage system, AnalyzeThis. Our guiding principle is that analysis-awareness be deeply ingrained in each and every layer of the storage, elevating data analyses as first-class citizens, and transforming AnalyzeThis into a potent analytics-aware appliance. We implement the AnalyzeThis storage system atop an emulation platform of the Active Flash array. Our results indicate that AnalyzeThis is viable, expediting workflow execution and minimizing data movement.

Categories and Subject Descriptors

D.4.2 [Operating Systems]: Storage Management—*secondary storage*; D.4.7 [Operating Systems]: Organization and Design—*distributed systems*

Keywords

Data Analytics, Visualization & Storage

1. INTRODUCTION

Data analysis is often considered the fourth paradigm of scientific discovery, complementing theory, experiment, and simulation. Experimental facilities (e.g., Spallation Neutron Source [46]), observational devices (e.g., Sloan Digital Sky Survey [43], Large Synoptic Survey Telescope [24]) and high-performance computing (HPC) simulations of scientific phenomena on clusters (e.g., Titan supercomputer [50] and other Top500 machines [52]) produce hundreds of terabytes

of data that need to be analyzed to glean insights. The data products are often stored in central, shared repositories, supported by networked file systems (NFS) or parallel file systems (PFS) (e.g., Lustre [42] or GPFS [41]). Analyses that operate on these datasets are often I/O-intensive, and involve running a complex workflow job on a smaller cluster. The analysis workflow reads the input data from the central storage, applies a series of analytics kernels, such as statistics, reduction, clustering, feature extraction and legacy application routines, and writes the final, reduced data back to the storage system. We refer to the entire sequence of reading the input data, followed by analysis on a cluster, and writing the output as *Offline data analysis*.

Offline analysis incurs a substantial amount of redundant I/O, as it has to read the inputs from the storage system, and write the reduced results back. Reading back large data for analysis on a cluster exacerbates the I/O bandwidth bottleneck that is already acute in storage systems [19]. This is because, I/O bandwidth has traditionally been lagging behind the compute and memory subsystems, and the data production rates from simulations [51] and experimental facilities are compounding the problem further, creating a *storage wall*. Instead of an offline approach to data analysis, analyzing data *in-situ* on the storage system, where the data resides, can not only minimize data movement, but also expedite the time to solution of the analysis workflow. In this paper, we explore such an approach to data analysis.

To alleviate the I/O bottleneck, network-attached storage systems for clusters are being built with solid-state devices (SSD), resulting in either hybrid SSD/HDD systems or all flash arrays. The lack of mechanical moving parts, coupled with a superior I/O bandwidth and low latency, has made SSDs an attractive choice. We argue that SSDs are not only beneficial for expediting I/O, but also for on-the-fly data analysis. SSDs boast an increasing computational capability on the controllers, which have the potential to execute data analysis kernels in an *in-situ* fashion. In this model, the analysis is conducted near the data, instead of shipping the data to the compute cores of the analysis cluster.

In our prior work on Active Flash [5, 51], we explored the viability of offloading data analysis kernels onto the flash controllers, and analyzed the performance and energy trade-offs of such an offload. We found that Active Flash outperformed offline analysis via a PFS for several analysis tasks. In this paper, we explore how such an active processing element can form the fabric of an entire storage system that is workflow-aware.

An array of such Active Flash devices allows us to rethink

ACM acknowledges that this contribution was authored or co-authored by an employee, or contractor of the national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only. Permission to make digital or hard copies for personal or classroom use is granted. Copies must bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. To copy otherwise, distribute, republish, or post, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SC '15, November 15-20, 2015, Austin, TX, USA

© 2015 ACM. ISBN 978-1-4503-3723-6/15/11...\$15.00

DOI: <http://dx.doi.org/10.1145/2807591.2807622>

the way data analysis workflows interact with storage systems. Traditionally, storage systems and workflow systems have evolved independently of each other, creating a disconnect between the two. By blending the flash storage array and data analysis together in a seamless fashion, we create an analysis workflow-aware storage system, *AnalyzeThis*. Consider the following simple—yet powerful—analogy from day-to-day desktop computing, which explains our vision for *AnalyzeThis*. A *smart folder* on modern operating systems allows us to associate a set of rules that will be implemented on files stored in that folder, e.g., convert all postscript files into pdfs or compress (zip) all files in the folder. A similar idea extrapolated to large-scale data analysis would be: writing data to an analysis-aware storage system automatically *triggers* a sequence of predefined analysis routines to be applied to the data.

Contributions: We propose *AnalyzeThis*, a storage system atop an array of Active Flash devices. Our guiding principle is that analysis-awareness be deeply ingrained within each and every layer of the storage system, thereby elevating the data analysis operations as *first-class citizens*. *AnalyzeThis* realizes workflow-awareness by creating a novel *analysis data object abstraction*, which integrally ties the dataset on the flash device with the analysis sequence to be performed on the dataset, and the lineage of the dataset (Section 3.1). The analysis data object abstraction is overlaid on the Active Flash device, and this entity is referred to as the *Active Flash Element*, *AFE*. We mimic an AFE array using an emulation platform. We explore how scheduling, i.e., both data placement and workflow orchestration, can be performed within the storage, in a manner that minimizes unnecessary data movement between the AFEs, and optimizes workflow performance (Section 3.2). Finally, we design easy-to-use file system interfaces with which the AFE array can be exposed to the user (Section 3.3). The FUSE-based file system layer enables users to read and write data, submit analysis workflow jobs, track and interact with them via a */proc-like* interface, and pose provenance queries to locate intermediate data (Section 3.4). We argue that these concepts bring a fresh perspective to large-scale data analysis. Our results with real-world, complex data analysis workflows on *AnalyzeThis*, built atop an emulation-based AFE prototype, indicate that it is very viable, and can expedite workflows significantly.

1.1 Background on Active Flash

In this section, we present a summary of our prior work, Active Flash, upon which the *AnalyzeThis* storage system is built.

Enabling Trends: First, we highlight the trends that make flash amenable for active processing.

High I/O throughput and internal bandwidth: SSDs offer high I/O throughput and internal bandwidth due to interleaving techniques over multiple channels and flash chips. This bandwidth is likely to increase with devices possessing more channels or flash chips with higher speed interfaces.

Availability of spare cycles on the SSD controller: SSD controllers exhibit idle cycles on many workloads. For example, HPC workloads are bursty, with distinct compute and I/O phases. Typically, a busy short phase of I/O activity is followed by a long phase of computation [7, 19]. Further, the I/O activity recurs periodically (e.g., once every hour), and the total time spent on I/O is usually low

(below 5% [20]). Even some enterprise workloads exhibit idle periods between their I/O bursts [25, 26]. Data ingest from experimental facilities, such as SNS [46] are based on the availability of beam time, and there are several opportunities for idle periods between user experiments, which involve careful calibration of the sample before the beam can be applied to it to collect data. Such workloads expose spare cycles available on the SSD controller, making it a suitable candidate for offloading data analysis tasks.

Multi-core SSD controllers: Recently marketed SSDs are equipped with fairly powerful mobile cores, and even multi-core controllers (e.g. a 4-core 780 MHz controller on the OCZ RevoDrive X2 [32]). Multi-core SSD controllers are likely to become more common place, and hence the available idle time on the SSD controllers will increase as well.

Active Flash: In our prior work on Active Flash [5, 51], we presented an approach to perform in-situ data analysis on SSDs. We presented detailed performance and energy models for Active Flash and offline analysis via PFS, and studied their provisioning cost, performance, and energy consumption. Our modeling and simulation results indicated that Active Flash is better than the offline approach in helping to reduce both data movement, and energy consumption, while also improving the overall application performance. Interestingly, our results suggest that Active Flash can even help defray part of the capital expenditure of procuring flash devices through energy savings. We also studied hybrid analysis, involving processing on both flash and host cores, and explored when it might be suitable to offload analysis to flash. Next, our simulation of I/O-compute trade-offs demonstrated that internal scheduling may be used to allow Active Flash to perform data analysis without impact on I/O performance. To this end, we explored several internal scheduling strategies within the flash translation layer (FTL) such as analyzing while data written to the flash is still in the controller’s DRAM, analyzing only during idle times (when there is no I/O due to data ingest), and combining idle time analysis with the scheduling of garbage collection (GC) to preempt GC interrupting an ongoing data analysis due to the lack of available free pages. Finally, we have demonstrated the feasibility of Active Flash through the construction of a prototype, based on the OpenSSD development platform, extending the OpenSSD FTL with data analysis functions. We have explored the offloading of several data analysis kernels, such as edge detection, finding local extrema, heartbeat detection, data compression, statistics, pattern matching, transpose, PCA, Rabin fingerprinting and k-means clustering, and found Active Flash to be very viable and cost-effective for such data analysis.

2. ANALYZETHIS STORAGE SYSTEM

2.1 Goals

In this section, we discuss our key design principles.

Analysis-awareness: Our main objective is to introduce analysis-aware semantics into the storage system. There is an urgent need to analyze the data in-situ, on the storage component, where the data already resides.

Reduce Data Movement: It is expected that in future, exascale data centers, the cost of data movement will rival that of the computation itself [16]. Thus, we need to minimize data movement in analysis workflows as well as across the AFEs within the storage.

Figure 1: *AnalyzeThis* overview. Figure shows analysis-awareness at each and every layer of *AnalyzeThis*.

Capture Lineage: There is a need to track provenance and intermediate data products generated by the analysis steps on the distributed AFEs. The intermediate data can serve as starting points for future workflows.

Easy-to-use File System Interface: The workflow orchestration across the AFEs needs to be masqueraded behind an easy-to-use, familiar interface. Users should be able to easily submit workflow to the storage system, monitor and track them, query the storage system for intermediate data products of interest and discover them.

2.2 Overview

We envision *AnalyzeThis* as a smart, analytics pipeline-aware storage system atop an array of Active Flash devices (Figure 1). The analysis workflow job is submitted to the *AnalyzeThis* storage system. As the input data to be processed becomes available on *AnalyzeThis* (from experiments, observations or simulations) the workflow that the user has submitted is applied to it. The final processed data, or any of the intermediate data is stored in *AnalyzeThis*, and may be retained therein, transferred to other repositories that may be available to the user (e.g., archive, PFS), or removed based on lifetime metadata attributes that the user may have associated with the dataset. Thematic to the design of *AnalyzeThis* is that analysis-awareness be deeply embedded within each layer of the storage system. In the future, we expect that such analysis-aware semantics will be adopted into existing PFS and NFS storage. Below is a bottom-up description of the system.

Active Flash Array: At the lowest level is the Active Flash array that is composed of discrete Active Flash devices, capable of running individual analysis kernels. We envision an array of such devices that are connected via SATA, PCIe or NVMe.

Analysis Object Abstraction: On top of the Active Flash array, we propose to create a new data model, the *analysis object abstraction* that encapsulates the data collection, the analysis workflow to be performed on the data, and the lineage of how the data was derived. We argue that such a rich data model makes analysis a *first-class citizen* within the storage system by integrally tying together the data and the processing to be performed (or was performed)

on the data. The analysis abstraction, coupled with the Active Flash device (capable of processing), is referred to as the *Active Flash Element*, “AFE.”

Workflow Scheduling Layer: The goal of this layer is to mimic how users interact with batch computing systems and integrate similar semantics into the storage system. Such a strategy would be a concrete step towards bridging the gap between storage and analysis workflows. Users typically submit a workflow, e.g., a PBS [15] or a DAGMAN [49] script, to a cluster’s batch scheduler, which creates a dependency graph and dispatches the tasks onto the compute nodes based on a policy. Similarly, we propose a *Workflow Scheduler* that determines both data placement and scheduling analysis computation across the AFEs in a manner that optimizes both end-to-end workflow performance and data movement costs.

A File System Interface: We tie the above components together into a cohesive system for the user by employing a FUSE-based file system interface with limited functionality (“anFS”). anFS supports a namespace, reads and writes to the AFE array, directory creation, internal data movement between the AFEs, a “/proc-like” infrastructure, and the ability to pose provenance queries to search for intermediate analysis data products. Similar to how /proc is a control and information center for the OS kernel, presenting runtime system information on memory, mounted devices and hardware, /mnt/anFS/.analyzethis/, allows users to submit workflow jobs, track and interact with them, get status information, e.g., load about the AFEs.

Together, these constructs provide a very potent in-situ data analytics-aware storage appliance.

3. DESIGN AND IMPLEMENTATION

Figure 2(a) presents the architecture of *AnalyzeThis*. The *AnalyzeThis* appliance exposes a FUSE file system, “anFS,” to the users that can be mounted via an NFS protocol. Users submit analysis workflows and write data objects to *AnalyzeThis* via anFS that is mounted on their desktop computer or an analysis cluster. Thereafter, users can monitor and query the status of the jobs, and search for intermediate data products of branches of the workflow. At the backend, the *AnalyzeThis* appliance comprises of one or more *storage servers* to which multiple Active Flash devices are connected. The storage server and the Active Flash devices run several services, and collectively help realize the analysis-aware storage appliance. Each Active Flash device runs the software service that overlays an analysis object abstraction atop, transforming it into an AFE, as well as other services required to handshake with the storage server. The storage server runs services such as those required for making the AFEs available as a file system (namespace management and data protocols), distributed workflow orchestration, distributed provenance capture, and interfacing with the AFEs (device management). These services are implemented behind a FUSE layer on the storage server. Together, in a distributed fashion, they achieve workflow awareness.

Central to our design is the seamless integration of workflow scheduling and the file system. To this end, behind anFS is a *workflow scheduler* that constructs a directed acyclic graph (DAG) from the analysis workflow job. The scheduler produces multiple mini DAGs based on task dependencies and optimization strategies, e.g., to minimize data movement. The mini DAGs comprise of a series of tasks, which

(a) **AnalyzeThis Architecture:** The figure shows how the desktop clients interact with AnalyzeThis by mounting the anfs client via the NFS protocol. AnalyzeThis comprises of multiple Active Flash devices, connected to one or more storage servers. Together, they run several services such as AFE, anFS servers, namespace, workflow, provenance and device managers.

(b) **anFS Architecture and Data and Control Paths:** The Workflow and device managers handle the active file operations or the control path. The Namespace manager, along with ExoFS, exposes the AFE as a file system to the FUSE layer for traditional data operations. The FUSE layer of anFS ties together the control and data paths into a user-level file system.

Figure 2: AnalyzeThis architecture and components.

the storage server maintains in a lightweight database. The scheduler dispatches the mini DAGs for execution on the AFEs; AFEs form the bottom-most layer of our system, and are capable of running analysis kernels on the device controllers. We use an *analysis object abstraction*, which encapsulates all necessary components of an task, including analysis kernels, input and output datasets, and the lineage information of all the objects therein. The analysis kernels for a given workflow is assumed to be stored as a platform-dependent binary executable object (.so format), compiled for specific devices as needed, which can run on the AFEs.

3.1 Analysis Encapsulation

We introduce analysis awareness in the Active Flash array by building on our prior work on Active Flash that has demonstrated how to run an analysis kernel on the flash controller [5, 51]. Our goal here is to study how to overlay an *analysis object abstraction* atop the Active Flash device, both of which together form the AFE. The construction of an AFE involves interactions with the flash hardware to expose features that higher-level layers can exploit, communication protocol with the storage server and flash device, and the necessary infrastructure for analysis object semantics. An array of AFEs serve as building blocks for AnalyzeThis.

The first step to this end is to devise a richer construct than just files. Data formats, e.g., HDF [45, 10, 34, 53] NetCDF [30, 21], and NeXus [31], offer many desirable features such as access needs (parallel I/O, random I/O, partial I/O, etc.), portability, processing, efficient storage and self-describing behavior. However, we also need a way to tie the datasets with the analysis lifecycle in order to support future data-intensive analysis. To address this, we extend the

concept of a basic data storage unit from traditional file(s) to an analysis object abstraction that includes a file(s) plus a sequence of analyses that operate on them plus the lineage of how the file(s) were derived. Such an abstraction can be created at a data collection-level, which may contain thousands of files, e.g., climate community. The analysis data abstraction would at least have either the analysis sequence or the lineage of analysis tools (used to create the data) associated with the dataset during its lifetime on AnalyzeThis. The elegance of integrating data and operations is that one can even use this feature to record *data management* activities as part of the dataset and not just analyses. For example, we could potentially annotate the dataset with a *lifetime* attribute that tells AnalyzeThis which datasets (final or intermediate data of analysis) to retain and for how long. The analysis object abstraction transforms the dataset into an encapsulation that is more than just a pointer to a byte stream; it is now an entity that lends itself to analysis.

3.1.1 Extending OSD Implementation for AFE

We realize the analysis object abstraction using the object storage device (OSD) protocol. The OSD protocol provides a foundation to build on, by supporting storage server to AFE communication and by enabling an object container-based view of the underlying storage. However, it does not support analysis-awareness specifically. We use an open source implementation of the OSD T10 standard, Linux open-osd target [33], and extend it further with new features to implement the AFEs. We refer to our version of the OSD implementation as “AFE-OSD” (Figure 2(a)). Our extensions are as follows: (i) *Mini Workflow* supports the execution of entire branches of an analysis workflow that are handed down by the higher-level Workflow Scheduler on the storage server; (ii) *Task Tracker* tracks the status of running tasks on the AFE; (iii) *AFE Status* checks the internal status of the AFEs (e.g., load on the controller, capacity, wear-out), and makes them available to the higher-level Workflow Scheduler on the storage server to enable informed scheduling decisions; (iv) *Lineage Container* captures the lineage of the executed tasks; and (v) *Lightweight Database Infrastructure* supports the above components by cataloging the necessary information and their associations.

Mini Workflow Engine: The AFE-OSD initiator on the storage server submits the mini DAG to the AFE-OSD target. The mini DAG represents a self-contained branch of the workflow that can be processed on an AFE independently. Each mini DAG is composed of a set of tasks. A task is represented by an analysis kernel, and a series of inputs and outputs. The storage server dispatches the mini DAGs to the AFEs using a series of ANALYZE_THIS execution commands, with metadata on the tasks. To handle the tasks on the AFEs, we have implemented a new *analysis task collection* primitive in the AFE-OSD, which is an encapsulation to integrally tie together the analysis kernel, its inputs and outputs (*Task Collection* and the Task Collection Attribute Page are represented in the bottom half of Figure 3). Once the AFE receives the execution command, it will create an analysis task collection, and insert the task into a FIFO task queue that it maintains internally. As we noted earlier, inputs and outputs can comprise of thousands of files. To capture this notion, we create a *linked collection* encapsulation for input and output datasets (using an existing *Linked* collection primitive), which denotes that a set of

files are linked together and belong to a particular dataset.

Kernel Executer: The kernel executer is a multi-threaded entity that checks the task queue and dispatches the tasks to the AFE controller. We have only used one core from the multi-core controller, but the design allows for the use of many cores. We rely on the ability of the Active Flash component of the AFE to run the kernel on the controller core, which has been studied in our prior work [5, 51]. Active Flash locates the predefined entry point (`start_kernel`) from the analysis kernel code (`.so` file), and begins the execution. In our prior work on Active Flash [5] (summarized in Section 1.1), we have explored FTL scheduling techniques to coordinate regular flash I/O, active computation and garbage collection, which can be used by the kernel executer.

Task Tracker: The Task Collection and the Task Collection Attribute page provide a way to track the execution status of a task and its execution history, i.e., run time. Each task collection has a unique task id. The storage server can check the status of a running task by reading an attribute page of its task collection using the `get_attribute` command and the task id. The workflow scheduler on the storage server also queries the AFE for the execution history of analysis kernels, to get an estimate of run times that are then used in scheduling algorithms, e.g., *Minimum Wait* (in Section 3.2).

AFE Status: The storage server can use an AFE’s hardware status for better task scheduling. To this end, we have created a *status object* to expose the internal information to the storage server. The status object includes the AFE device details such as wear-out for the flash, resource usage for controller, the AFE task queue details, and Garbage Collection status. The AFEs are configured to periodically update a local status object, which can then be retrieved by the storage server as needed. Thus, the storage server can check the status of the device by sending a `get_attribute` command on the status object using its object id.

Lineage Container: Lineage information of tasks and data objects are maintained in an AFE’s internal database. The lineage container helps answer provenance queries (more details in Section 3.4).

Lightweight Database Infrastructure: We use a lightweight database infrastructure (Figure 3), using SQLite [48], to implement analysis-aware semantics into the storage system. One approach is to have the storage server maintain all of the analysis workflow, data, and analysis-semantics. However, such a centralized approach is not resilient in the face of storage server crashes. Instead, we implement a decentralized approach (refer to Figure 2(a)), wherein the storage server (the FUSE layer) and the AFEs (the AFE-OSD Target) maintain relevant information and linkages to collectively achieve the analysis abstraction.

The storage server database table (*anFS_wf*) maintains high-level information about the analysis workflow, e.g., mini DAGs (job ID), the associated tasks (task collection ID), and the AFE ID on which to run the task. For example, in Figure 3 user Alice runs a job (id = 1) and executes an analysis task, *kmeans* (CID = 0x10), on AFE (id = 0). The local AFE database tables store detailed metadata on all objects, namely mini DAGs, task collections, input and output datasets, their attributes and associations.

Each AFE manages three tables. The *Object table* is used to identify the type of an object, e.g., whether it is a task

Figure 3: Analysis object abstraction implemented using database engine. CID: task collection id, OID: object id, Attr. offset: an offset in the attribute page.

collection, or a data object. For each object, it maintains object identifiers, and object types. The *Dataset TaskCollection Membership table*, *tcid_oid*, manages the membership of data objects to task collections. Multiple data objects can belong to a task collection (e.g., multiple inputs to a task) or a data object can be a member of multiple task collections (e.g., a given dataset is input to multiple tasks). The *Attribute table* manages all the attributes of data objects and task collections (e.g., those represented in the Task Collection Attribute Page). Each attribute (or record) in the attribute table is defined using a data object or task collection id, page number, and attribute number inside the Attribute Page. Given this metadata, the storage server can query information on the tasks and their associated datasets. For example, given a task collection of 0x10 and an index into the attribute page, 1 (to refer to input datasets), the attribute table points to a value of 0x15, which can be reconciled with the *tcid_oid* table to obtain the input datasets 0x1002 and 0x1003.

3.2 Workflow Engine

We have built a *workflow engine* within AnalyzeThis, to orchestrate both the placement of data objects as well as the scheduling of analysis tasks across the AFEs. The scheduler is implemented in the FUSE file system (anFS). Once a user submits the analysis workflow script via anFS, it distinguishes this request from a normal I/O request. This is accomplished by treating the “write” request coming through the special file (`.submit`) as a job submission instead of normal write operation by anFS. The script is delivered to the scheduler which parses it to build a directed acyclic graph (DAG), schedules the tasks, and sends the execution requests to the AFEs via the AFE-OSD protocol. The vertices in the DAG represent the analysis kernels, and inputs and outputs represent incoming and outgoing edges. The sched-

uler decides which branches (mini DAGs) will run on which AFEs based on several heuristics. While the mapping of a mini DAG to AFE is determined apriori by the scheduler, the tasks are not dispatched until the analysis job’s inputs are written to AnalyzeThis. This is akin to the *smart folder* concept discussed in Section 1. The analysis sequence is first registered with AnalyzeThis, and once the input data is available the tasks are executed.

Workflow Description and DAG: In our implementation, we have chosen to represent a job script using *Lib-config* [22], a widely used library for processing structured configuration files. Listing 1 shows an example of a job that finds the maximum value in each input file. Each tasklet is represented by the input and output object lists, and a kernel object that operates on the input objects. Any data dependencies among the tasklets are detected by the scheduler via a two-pass process. In the first pass, the scheduler examines each task in the script and inserts the task record (key: output file, value: task) into a hash table. In the second pass, the scheduler examines the input files of each task in the job script. When an input file is found in the hash table, the examined task is dependent on the output of the task in the hash table. If the input file is not found in the hash table, the scheduler checks if the input file already exists in the file system. If the file does not exist, the job script is considered to be invalid. In this way, the dependencies among the tasks can be resolved. The dependency information is used to build a DAG. In the following example, `getmax.reduce` cannot be launched until `getmax.1` and `getmax.2` produce their output objects. Therefore, the overall performance of AnalyzeThis depends on the efficient scheduling of the analysis kernels on the AFE array.

Listing 1: An example job script

```

name = "getmax";
workdir = "/scratch/getmax/";
tasks = (
{ name = "getmax.1"; kernel = "getmax.so";
  input = [ "1.dat" ]; output = [ "1.max" ]; },
{ name = "getmax.2"; kernel = "getmax.so";
  input = [ "2.dat" ]; output = [ "2.max" ]; },
{ name = "getmax.reduce"; kernel = "mean.so";
  input = [ "1.max", "2.max" ];
  output = [ "max.dat" ]; }
);

```

Scheduling Heuristics: The design of the workflow scheduler is driven by two objectives: (1) minimizing the overall execution time, and (2) reducing the data movement across the AFEs. We point out that minimizing data movement is critical as uncontrolled data movement may cause early wear-out of the SSDs and increase in the energy consumption [11]. We have designed and implemented several scheduling strategies that attempt to strike a balance between these two competing objectives.

Round-robin: A simple round-robin approach schedules tasks as soon as their dependency requirements are met, and ensures a homogeneous load-distribution across all AFEs in a best-effort manner since the tasks are scheduled without a priori information about their execution time. It picks the next available AFE in a round-robin fashion to balance the computational load. The round-robin strategy schedules the task on an available idle AFE controller, causing data movement, potentially in favor of a shorter execution time and load balance across the AFE controllers. Consequently, the technique may suffer from excessive data movement because

it does not account for the amount of data to be moved.

Input Locality: To minimize the data movement across the AFEs, this heuristic schedules tasks based on input locality. Tasks are scheduled on an AFE where maximum amount of input data is present. The scheduler maintains this information in memory during a job run, including the size and location of all involved files. Input-locality favors a reduction in data movement to performance (execution time). In our experiments with real workflows, we observed that this scheduling policy is effective in reducing the data movement. However, it can potentially increase the overall execution time considerably because it will execute the analysis on the AFE that stores larger input, even if other AFEs are idle.

Minimum Wait: To reconcile execution time and data movement, we propose to explicitly account for the data transfer time and queuing delays on the AFE controllers. The heuristic takes two inputs including a list of all available AFEs and the tasks to be scheduled next. The scheduler maintains information about the jobs currently queued on each AFE, their expected finish time and the size of the input file(s) for the task to be scheduled next. The scheduler iterates over each AFE to estimate the minimum wait time for the task to be scheduled. For each AFE, it calculates the queue wait time (due to other jobs) and data transfer time to that particular AFE. It chooses the AFE for which the sum of these two components is minimum. The minwait scheduler maintains and updates the “expected free time” of each AFE using the runtime history of jobs. When a task is ready to be executed, the scheduler calculates the expected wait time of the task for every AFE. The expected wait time at an AFE is calculated as: “expected free time” at the AFE plus the expected data transfer time (estimated using the input file size and AFE location). The scheduler assigns the task to an AFE that is expected to have the minimum wait time.

Hybrid: In the hybrid strategy, we exploit the storage server within the AnalyzeThis appliance (storage server in Figure 2(a)), in addition to the AFEs, to exploit additional opportunities. The storage server can run certain tasks in addition to anFS services. Running computation on the servers to which the disks are attached is a well-known practice adopted by several commercial vendors. However, hybrid processing offers more benefits by further exploiting the internal, aggregate bandwidth of the multi-channel flash device that exceeds the PCIe bandwidth between the storage server and the flash device by a factor of 2-4× [8]. AnalyzeThis does not blindly place all tasks on the AFEs or on the host storage server. The challenge is in carefully determining what to offload where (storage server vs. AFE) and when. Traditional solutions that simply perform server-side processing do not address this optimization. Reduce tasks in workflows involve high data movement cost because they gather multiple intermediate inputs on the AFEs, and can be moved to the storage server. This approach has the advantage of minimizing the overhead of data movement between the AFEs, beyond what Input Locality alone can achieve, without sacrificing the parallelism. Also, tasks that cause an uneven distribution on the AFEs cause stragglers, and can be moved to the storage server (unaligned tasks). Such tasks can be identified based on profiling of the workflows. The hybrid approach can work in conjunction with any of the aforementioned scheduling techniques.

3.3 anFS File System Interface

The functionalities of AnalyzeThis are exposed to the clients via a specialized file system interface, “anFS.” Since the analysis workflows operate on but do not modify the original data from scientific experiments and simulations, anFS is designed as a write-once-read-many file system. As discussed earlier (Section 3), anFS is exported to clients via NFS. Thus, operations on shared files follow the NFS consistency semantics. anFS provides standard APIs such as `open()`, `read()`, and `write()`, as well as support special virtual files (SVFs), serving as an interaction point, e.g., to submit and track jobs, between users and AnalyzeThis.

Figure 2(b) shows the overall architecture of anFS. It is implemented using the FUSE user-space file system, and can be mounted on the standard file system, e.g., `/mnt/anFS/`. FUSE provides an elegant way to develop user-level file systems. anFS provides several custom features, such as workflow execution and provenance management, which are more appropriate to be implemented in the user-space than the kernel-level. Also, a FUSE-based, user-level implementation offers better portability than a kernel-based solution. anFS is composed of the following components. The *Namespace Manager* consolidates the array of available AFEs, and provides a uniform namespace across all the elements. The *Workflow Manager* implements the workflow engine of AnalyzeThis (Section 3.2). The *Device Manager* provides the *control path* to the AFEs, implementing AFE-OSD (Section 3.1.1) to allow interactions with the AFEs. Finally, the *exoFS (Extended Object File System) layer* [12] provides the *data path* to the AFEs.

Namespace: anFS exposes a consolidated standard hierarchical UNIX namespace view of the files and SVFs on the AFEs. To this end, the storage server metadata table (Section 3.1.1) includes additional information associated with every stored object and information to track the AFEs on which the objects are stored (Figure 3). For example, there is an AFE identifier and an object identifier associated with every inode of a file stored by anFS. All file system operations are first sent to the Namespace Manager that consults the metadata to route the operation to an appropriate AFE. To manage the large amount of metadata that increases with increasing number of files, provide easy and fast access, and support persistence across failures we employ the SQLite RDBMS [48] to store the metadata. We note that anFS implements features such as directories, special files, and symbolic links, entirely in the metadata database; the AFEs merely store and operate on the stored data objects. Instead of striping, anFS stores an entire file on a single AFE to facilitate on-element analysis and reduce data movement. The placement of a file on an AFE is either specified by the workflow manager, or a default AFE (i.e., inode modular number-of-AFEs) is used.

Data and Control Path: To provide a data path to the AFEs, anFS uses exoFS, an ext2-based file system for object stores. anFS stores regular data files via the exoFS mount points, which are created one for each AFE. For reads and writes to a file, anFS first uses the most significant bit of the 64-bit inode number to distinguish between a regular file (MSB is 0) and a SVF (MSB is 1). For regular files, the Namespace Manager locates the associated AFE and uses exoFS to route the operation to the AFEs as shown in Figure 2(b). Upon completion, the return value is returned to the user similarly as in the standard file system. To provide

a control path for active operations, anFS intercepts the files and routes it to the Workflow Manager, which uses the Device Manager to route the operations to the AFEs using the AFE-OSD library for further analysis and actions.

Active File Operations — Job Submission: anFS supports SVFs to allow interaction between users and AnalyzeThis operations, e.g., running an analysis job, checking the status of the job, etc. Specifically, we create a special mount point (`.analyzethis`) under the root directory for anFS (e.g., `/mnt/anFS/`), which offers similar functionality as that of `/proc` but for workflow submission and management (Figure 2(b)). To submit a job, e.g., JobA, the user first creates a submission script (`/home/alice/jobA-submission`) that contains information about how the job should be executed and the data that it requires and produces. Next, the job is submitted by writing the full path of the submission script to the submission SVF, e.g., by using `echo/home/alice/jobA-submission>/mnt/anFS/.analyzethis/alice/submit`. This SVF write is handed to the Workflow Manager for processing, which parses the script, assigns a unique opaque 64-bit job handle to the script, and takes appropriate actions such as creating a task schedule, and using the appropriate Device Manager thread to send the tasks to the AFEs. The Workflow Manager also updates a per-user active job list, e.g., `SVF /mnt/anFS/.analyzethis/alice/joblist` for user Alice, to include the job handle for the newly submitted job. Each line in the joblist file contains the full path of the submission script and the job handle. Moreover, the Workflow Manager also monitors the task progress. This information can be retrieved by the user by reading from the job handle SVF `/mnt/anFS/.analyzethis/alice/jobA-status`. When the user accesses the job handle, the request is directed to the Device Manager thread for the AFE, via the Workflow Manager. The Device Manager thread sends the `get_attribute` command via the AFE-OSD protocol to the *Task Tracker* in the Mini Workflow Engine on the AFE to retrieve the status of the jobs.

Supporting Internal Data Movement: Ideally, AnalyzeThis will schedule a task to an AFE that also stores the (*most*) data needed by the task. While we attempt to minimize data movement through smart heuristics, there is still the need to move data between AFEs as a perfect assignment is not feasible. To this end, anFS may need to replicate (or move) data from one AFE to another by involving the storage server. However, this incurs added overhead on the storage server. In the future, direct PCI to PCI communication can help expedite these transfers.

Data and Workflow Security: anFS ensures data security for multiple users via the OSD2 standards. To protect the stored data, OSD maintains the ownership of objects as object attributes. When a data item is stored, it also provides the kernel-level user-id of the data owner, which is then stored in the OSD-level object ownership metadata automatically by the device. When the data is accessed, the user-id information is provided along with the request, and the OSD2 protocol ensures that only the allowed user(s) are given access to the data. Similarly, when a task is scheduled on the AFE, it is associated with the user-id information, and must present these credentials to access data. The access control for the SVFs are set such that the submit SVF (`.anFS/submit`) is world writable, but the resulting joblist and status files are user specific. The sub-directories

are named on a per-user basis, e.g., Alice’s jobs are under `.anFS/alice/`, and the associated POSIX ACLs protect user files.

3.4 Provenance

AnalyzeThis tracks the lineage of the data produced as a result of a workflow execution at very minimal cost. This allows the user to utilize the intermediate data for future analysis. We have implemented provenance support on top of the distributed database infrastructure (Figure 3) between the storage server (workflow table, *anFS_wf*) and AFEs (Dataset task collection membership table, *tcid_oid*). Recall that *anFS_wf* stores information about the task and the AFE on which the task is executed; *tcid_oid* stores the task collection to data object mapping and will also need to be maintained on the storage server. Upon receiving a provenance query regarding a dataset, AnalyzeThis searches the *anFS_dirent* table to get the *anFS_inode* of the file, which is used to get the object id. The object id is then used to retrieve the task collection id from the *tcid_oid* table. The task collection id is used to obtain the AFE id from the *anFS_wf* table. Alternatively, if *tcid_oid* is not maintained on the storage server as well, we can broadcast to the AFEs to determine the task collection id for a data object id. Further analysis of the lineage is performed on that AFE. Using the task collection id and the attribute page we get the task collection attribute page number. Using the predefined attribute offset all the information regarding the task is fetched. The task provenance from multiple AFEs is merged with similar job-level information that is maintained at the storage server in the *anFS_wf* table.

4. EVALUATION

4.1 Experimental Setup

Testbed: Our emulation testbed (Figure 4) is composed of the following: (1) client desktop computer that submits analysis workflows, (2) storage server within the AnalyzeThis appliance, and (3) the networked machines that emulate the AFEs (four AFEs are connected to the storage server). For the desktop computer and the storage server, we used a 1.8 GHz Intel Xeon E5-2603 processor. We emulated the AFEs using Atom machines with RAM disks, to mimic the flash controller and the internal flash chips with high I/O bandwidth to the controller. The Atom-based AFEs use a single 1.8 GHz Atom processor as the controller, a 3 GB RAM disk as the flash chip, and a 1 Gbps Ethernet connection to the storage server within the AnalyzeThis appliance. All servers run the Linux kernel 2.6.32-279. anFS offers a read and write bandwidth of 120 MB/s and 80 MB/s, respectively.

Software: AnalyzeThis has been implemented using 10 K lines of C code. We extended the OSD iSCSI target emulator from the open-osd project [33], for the AFE target. The task executions in an AFE are serialized by spawning a dedicated thread, which mimics dedicating a device controller for active processing. For the AFE-OSD driver in the storage server, we extended the OSD initiator driver in the Linux kernel. We also extended exoFS [12] to synchronize the OSD object id space with the userspace anFS. anFS has been implemented using FUSE [13], and it keeps track of metadata using SQLite [48].

Scientific Workflows: We used several real-world complex workflows. We used Montage [27], Brain Atlas [28],

Figure 4: AnalyzeThis testbed.

Figure 5: The DAGs representing the workflows.

Sipros [54], and Grep [14] workflows. The DAG representations and the details of the workflows are shown in Figure 5 and Table 1.

The Montage workflow [27] creates a mosaic with 10 astronomy images. It uses 8 analysis kernels, and is composed of 36 tasks, several of which can be parallelized to run on the AFEs. The Brain workflow [28] creates population-based brain atlases from the fMRI Data Center’s archive of high resolution anatomical data, and is part of the first provenance challenge [28] used in our provenance evaluation. The Sipros workflow runs DNA search algorithms with database files to identify and quantify proteins and their variants from various community proteomics studies. It consists of 12 analysis tasks, and uses three analysis kernels. The Grep workflow counts the occurrences of ANSI C keywords in the Linux source files.

	Input	Intermediate (MB)	Output	Total	Object (#)
Montage	51	222	153	426	113
Brain	70	155	20	245	35
Sipros	84	87	1	172	45
Grep	463	363	1	827	13

Table 1: Workflow input, output and intermediate data size.

Figure 6: Comparison of AnalyzeThis round-robin, hybrid, and offline-anFS. Multiple runs for each case, without much variance.

4.2 AnalyzeThis Performance

We compare offline-anFS and AnalyzeThis. In offline-anFS, data analyses are performed on desktops or clusters by pulling data from the anFS, whereas in *AnalyzeThis*, they are performed on the AFE cores. In Figure 6, we show the total runtime in terms of computation and I/O time. We further break down the I/O time into anFS I/O and anFS-internal I/O times (i.e., data movement between the AFEs). Therefore, a break-down of the workflow run time comprises of the following: (i) time to read the data from anFS (only offline-anFS incurs this cost), (ii) compute time of the workflow on either the desktop or the AFEs, (iii) I/O time to write the intermediate output to anFS during analysis (only for offline-anFS), (iv) data shuffling time among the AFEs (only for AnalyzeThis), and (v) time to write the final analysis output to anFS. We specifically compared the following scenarios: (i) offline analysis using one client node and anFS (offline-anFS), (ii) AnalyzeThis using four Atom-based AFEs and round-robin scheduling across the AFEs, and (iii) AnalyzeThis-hybrid using the storage server, four AFEs and round-robin across the AFEs.

In the Montage, Brain and Grep experiments for offline-anFS, the time to write the analysis outputs to anFS noticeably increases the run time (more than 20% of the run time), while, for AnalyzeThis, the I/O time, anFS-internal I/O, is much smaller compared to the overall run time. The run time for offline-anFS for Montage and Brain is slightly lower than AnalyzeThis due to relatively less computing power on the AFEs. However, as AFEs begin to have multicores in the future, this small difference is likely to be overcome. In contrast, for Sipros and Grep, AnalyzeThis performs better than offline-anFS. This is because the tasks are memory-bound. The results indicate that offline’s performance is heavily affected by the data movement costs, whereas AnalyzeThis is less impacted. Further, AnalyzeThis can free up compute resources of desktops or clusters, enabling “true” out-of-core data analysis.

Next, we evaluate (AnalyzeThis-Hybrid). For Montage, AnalyzeThis-Hybrid significantly reduced the total run time over AnalyzeThis and offline-anFS. Unaligned *mProjectPP* tasks (Figure 5(a)) are executed on the storage server, which removed task stragglers. Also, more than 50% of data copies between AFEs are reduced by executing reduce tasks on the storage server. Similarly, for Brain, executing a single reduce task (*softmax* in Figure 5(b)) on the storage server eliminated more than 75% of data copies, which results in a 37% runtime reduction compared to AnalyzeThis. Similarly, for Sipros, AnalyzeThis-hybrid is better than both AnalyzeThis and offline-anFS as it ran unaligned tasks on

the storage server.

4.3 Scheduling Performance

Here, we discuss the performance of scheduling techniques.

Impact of Scheduling Heuristics: Figure 6 compares the performance of round robin (RR), input locality (IL), minimum wait (MW), and hybrid (HY) based on AFE utilization and data movement. Figure 7(a) compares the sum (first bar) of the computation time of the workflow and the data shuffling time among the AFEs against the AFE utilization time (other two bars). AFE utilization is denoted by the slowest (second bar) and the fastest (third bar) AFEs, and the disparity between them indicates a load imbalance across the AFEs. The smaller the difference, the better the utilization. Figure 7(b) shows the amount of data shuffled between the AFEs. An optimal technique strikes a balance between runtime, data movement, and AFE utilization.

Figure 7: Performance of scheduling heuristics.

HY and RR show a balanced load distribution across the AFEs with the least variability in utilization. However, RR incurs the most data movement. IL can improve runtime by significantly reducing data movement, however it may degrade the overall performance due to inefficient load distribution. IL shows higher runtimes than RR in all workflows. In fact for IL, we observed in Montage that the slowest AFE was assigned 21 tasks among 36 tasks; in Brain, only two AFEs out of four executed all of the tasks; and in Sipros,

Figure 8: AFE scaling experiments for AnalyzeThis.

only one AFE was used during analysis. HY and MW perform best in reconciling AFE utilization and data movement cost. For Montage, MW shows a 10% lower runtime than IL by incurring a 6% increase in data movement. For Brain, RR and MW show very close runtimes, but MW further reduces the data movement cost of RR by 16%, with less core utilization, suggesting that it is likely to be more energy-efficient. For Sipros, MW shows a 2.4% lower runtime than RR while reducing the data movement cost by 3%. By executing reduce tasks on the storage server, HY significantly reduces the runtimes over other scheduling algorithms for all workflows. In Montage and Brain, this also reduces data movement cost by 52% and 76% over RR, respectively.

Scaling Experiments: We performed scalability experiments for AnalyzeThis by increasing the number of AFEs. Figure 8 shows the results with MW. Interestingly, we observe that the overall performance scales only up to a certain number of AFEs, since the maximum task parallelism in the workflow can limit the performance gain. In Montage, for instance, `mProjectPP` is the most time-consuming kernel, used by ten analysis tasks. After five AFEs, at least one AFE will run one more `mProjectPP` than others. Thus, there is little improvement in performance after five AFEs. Likewise, Brain and Sipros scale only up to four or five AFEs, respectively. Also, each workflow shows a different speed-up curve with the increase in AFEs, depending on its task dependency, number of kernels, inputs and outputs. For instance, Sipros shows a linear speedup up to five AFEs. This is because the less complex Sipros DAG (Figure 5(c)) allowed for more parallelism in task execution with the increase in AFEs than others. In some cases, however, the increased data movement cost with more AFEs can degrade the performance. For instance, Brain shows a slightly higher run time from four to eight AFEs, due to the increased data movement.

Provenance Performance: We used the Brain workflow (Figure 5(b)) and the provenance queries from the first provenance challenge [28] to evaluate AnalyzeThis. The provenance test consists of five queries, where Query 1 finds the provenance data up to the start of a job for a file; Query 2 finds the provenance data up to a task name for a file; Query 3 finds the provenance details for the levels within a DAG for a file; Query 4 finds all the invocations of a task that ran with the certain parameters and on a specific date; and Query 6 finds all output images produced by a task when another task was executed with certain arguments. In addition to the decentralized approach (Section 3.4), we implemented a centralized technique, where all of the lineage information is maintained on the storage server. For both approaches, one million entries were added and retrieved from the storage server side tables. For the decentralized approach, four

Query	Centralized (sec)	Decentralized (sec)
1	12.930	8.180
2	1.750	2.830
3	13.746	10.720
4	0.410	3.400
6	7.410	7.640

Table 2: Response time for five provenance queries.

AFEs were used and each had 0.25 million entries in all the database tables. For three out of five cases (Table 2), centralized performed better in terms of the query response time. Even though the query execution is distributed to multiple AFEs, there is no improvement in the processing time for some queries as the storage server side provenance information still needs to be parsed. Reduction in query time for decentralized is primarily due to the parallelization opportunity available in the workflow. On the other hand, the centralized approach offers no fault tolerance in the event of a storage server crash. Thus, decentralization is desirable even if the query time is higher in certain cases.

5. RELATED WORK

Migrating tasks to disks has been explored before [37, 18]. There is a renewed interest in active processing given the recent advances in SSD technology [39]. Recent efforts, such as iSSD [8], SmartSSD [17], and Active Flash [51] have demonstrated the feasibility and the potential of processing on the SSD. These early studies lay the foundation for AnalyzeThis.

The active storage community has leveraged the object storage device (OSD) protocol to enable computation within a storage device. The OSD T10 standard [55, 56, 38] defines a communication protocol between the host and the OSD. Recent efforts leverage the protocol for different purposes, including executing remote kernels [38], security, and QoS [36, 56]. In contrast, we extend the OSD implementation to support entire workflows, and to integrally tie together the data with both the analysis sequence and its lineage.

Some extensions to parallel file systems, e.g., PVFS [47] and Lustre [35], provide support for analysis on the I/O node’s computing core. However, they are not workflow-aware, a key trait for efficient analysis execution, and neither is the analysis conducted on the storage device. The ADIOS [23] I/O middleware uses a subset of staging nodes alongside a running simulation on a cluster to reduce the simulation output on-the-fly; while workflow-aware, it also only uses the computing elements of the staging nodes. Instead, AnalyzeThis uses the AFEs on the storage themselves, obviating the need for a separate set of staging nodes for analysis. Enterprise solutions such as IBM Netezza [44] enable provenance tracking and in-situ analysis, but lack an easy-to-use file system interface and workflow-awareness. Workflow- and provenance-aware systems, such as PASS [29], LinFS [40], BadFS [4], WOSS [1], and Kepler [2], are not meant for in-situ analysis. Compared to dedicated provenance systems like PASS and LinFS, lineage tracking in AnalyzeThis is a natural byproduct of executing workflows in the storage. Distributed execution engines, such as Dryad [57], Nephelē [3], Hydracks [6], and MapReduce [9], can execute data-intensive DAG-based workflows on distributed computing resources. AnalyzeThis fundamentally differs from these systems as it exploits the SSDs as the primary computing resources.

6. CONCLUSION

The need to facilitate efficient data analysis is crucial to derive insights from mountains of data. However, extant techniques incur excessive data movement on the storage system. We have shown how analysis-awareness can be built into each and every layer of a storage system. The concepts of building an analysis object abstraction atop an Active Flash array, integrating a workflow scheduler with the storage, and exposing them via a /proc-like file system bring a fresh perspective to purpose-built storage systems. We have developed the AnalyzeThis storage system on top of an emulation platform of the Active Flash array. Our evaluation of AnalyzeThis shows that is viable, and can be used to capture complex workflows.

7. ACKNOWLEDGMENTS

This research was supported in part by the U.S. DOE's Office of Advanced Scientific Computing Research (ASCR) under the Scientific data management program, and by NSF through grants CNS-1405697 and CNS-1422788. The work was also supported by, and used the resources of, the Oak Ridge Leadership Computing Facility, located in the National Center for Computational Sciences at ORNL, which is managed by UT Battelle, LLC for the U.S. DOE, under the contract No. DE-AC05-00OR22725. It was also supported in part by LBNL under Contract Number DE-AC02-05CH11231.

8. REFERENCES

- [1] S. Al-Kiswani, E. Vairavanathan, L. B. Costa, H. Yang, and M. Ripeanu. The Case for Cross-Layer Optimizations in Storage: A Workflow-Optimized Storage System. *arXiv preprint arXiv:1301.6195*, 2013.
- [2] I. Altintas, C. Berkley, E. Jaeger, M. Jones, B. Ludascher, and S. Mock. Kepler: An Extensible System for Design and Execution of Scientific Workflows. In *Proceedings of the 16th International Conference on Scientific and Statistical Database Management, SSDBM '04*, pages 423–, Washington, DC, USA, 2004. IEEE Computer Society.
- [3] D. Battré, S. Ewen, F. Hueske, O. Kao, V. Markl, and D. Warneke. Nephel/PACTs: A Programming Model and Execution Framework for Web-scale Analytical Processing. In *Proceedings of the 1st ACM Symposium on Cloud Computing, SoCC '10*, pages 119–130, New York, NY, USA, 2010. ACM.
- [4] J. Bent, D. Thain, A. C. Arpaci-Dusseau, R. H. Arpaci-Dusseau, and M. Livny. Explicit Control a Batch-aware Distributed File System. In *Proceedings of the 1st Conference on Symposium on Networked Systems Design and Implementation - Volume 1, NSDI'04*, pages 27–27, Berkeley, CA, USA, 2004. USENIX Association.
- [5] S. Boboila, Y. Kim, S. Vazhkudai, P. Desnoyers, and G. Shipman. Active Flash: Out-of-core Data Analytics on Flash Storage. In *Mass Storage Systems and Technologies (MSST), 2012 IEEE 28th Symposium on*, 2012.
- [6] V. Borkar, M. Carey, R. Grover, N. Onose, and R. Vernica. Hyracks: A Flexible and Extensible Foundation for Data-intensive Computing. In *Proceedings of the 2011 IEEE 27th International Conference on Data Engineering, ICDE '11*, pages 1151–1162, Washington, DC, USA, 2011. IEEE Computer Society.
- [7] P. Carns, K. Harms, W. Allcock, C. Bacon, S. Lang, R. Latham, and R. Ross. Understanding and Improving Computational Science Storage Access Through Continuous Characterization. *Trans. Storage*, 7(3):8:1–8:26, Oct. 2011.
- [8] S. Cho, C. Park, H. Oh, S. Kim, Y. Yi, and G. R. Ganger. Active Disk Meets Flash: A Case for Intelligent SSDs. In *Proceedings of the 27th International ACM Conference on Supercomputing, ICS '13*, pages 91–102, New York, NY, USA, 2013. ACM.
- [9] J. Dean and S. Ghemawat. MapReduce: Simplified Data Processing on Large Clusters. In *Proceedings of the 6th Conference on Symposium on Operating Systems Design & Implementation - Volume 6, OSDI'04*, pages 10–10, Berkeley, CA, USA, 2004. USENIX Association.
- [10] HDF5 - A New Generation of HDF. <http://hdf.ncsa.uiuc.edu/HDF5/doc/>.
- [11] The Opportunities and Challenges of Exascale Computing. http://science.energy.gov/~media/ascr/ascac/pdf/reports/exascale_subcommittee_report.pdf.
- [12] exofs [LWN.net]. <http://lwn.net/Articles/318564/>.
- [13] Filesystem in Userspace. <http://fuse.sourceforge.net/>.
- [14] Grep - Hadoop Wiki. <http://wiki.apache.org/hadoop/Grep>.
- [15] R. L. Henderson. Job scheduling under the portable batch system. In *Job scheduling strategies for parallel processing*, pages 279–294. Springer, 1995.
- [16] DOE Exascale Initiative Technical RoadMap. <http://extremecomputing.labworks.org/hardware/collaboration/EI-RoadMapV21-SanDiego.pdf>, 2009.
- [17] Y. Kang, Y.-s. Kee, E. L. Miller, and C. Park. Enabling cost-effective data processing with smart SSD. In *Mass Storage Systems and Technologies (MSST), 2013 IEEE 29th Symposium on*, 2013.
- [18] K. Keeton, D. A. Patterson, and J. M. Hellerstein. A Case for Intelligent Disks (IDISKs). *ACM SIGMOD Record*, 27(3):42–52, 1998.
- [19] Y. Kim, R. Gunasekaran, G. Shipman, D. Dillow, Z. Zhang, and B. Settlemeyer. Workload Characterization of a Leadership Class Storage Cluster. In *Petascale Data Storage Workshop (PDSW), 2010 5th*, pages 1–5, Nov 2010.
- [20] Computational Science Requirements for Leadership Computing. https://www.olcf.ornl.gov/wp-content/uploads/2010/03/ORNL_TM-2007_44.pdf, 2007.
- [21] J. Li, W. Liao, A. Choudhary, R. Ross, R. Thakur, W. Gropp, R. Latham, A. Siegel, B. Gallagher, and M. Zingale. Parallel netCDF: A High-Performance Scientific I/O Interface. In *Proceedings of SC2003: High Performance Networking and Computing*, 2003.
- [22] libconfig.

- <http://www.hyperrealm.com/libconfig/>, 2013.
- [23] J. F. Lofstead, S. Klasky, K. Schwan, N. Podhorszki, and C. Jin. Flexible IO and Integration for Scientific Codes Through the Adaptable IO System (ADIOS). In *Proceedings of the 6th international workshop on Challenges of large applications in distributed environments*, 2008.
 - [24] The New Sky | LSST. <http://www.lsst.org/lsst/>.
 - [25] N. Mi, A. Riska, E. Smirni, and E. Riedel. Enhancing Data Availability in Disk Drives through Background Activities. In *Dependable Systems and Networks With FTCS and DCC, 2008. DSN 2008. IEEE International Conference on*, pages 492–501, June 2008.
 - [26] N. Mi, A. Riska, Q. Zhang, E. Smirni, and E. Riedel. Efficient Management of Idleness in Storage Systems. *Trans. Storage*, 5(2):4:1–4:25, June 2009.
 - [27] Montage - An Astronomical Image Mosaic Engine. <http://montage.ipac.caltech.edu/docs/m101tutorial.html>.
 - [28] L. Moreau, B. Ludäscher, I. Altintas, R. S. Barga, S. Bowers, S. Callahan, G. Chin, B. Clifford, S. Cohen, S. Cohen-Boulakia, et al. Special Issue: The First Provenance Challenge. *Concurrency and computation: practice and experience*, 20(5):409–418, 2008.
 - [29] K.-K. Muniswamy-Reddy, D. A. Holland, U. Braun, and M. Seltzer. Provenance-Aware Storage Systems. In *Proceedings of the annual conference on USENIX '06 Annual Technical Conference*, 2006.
 - [30] NetCDF Documentation. <http://www.unidata.ucar.edu/packages/netcdf/docs.html>.
 - [31] Nexus. <http://trac.nexusformat.org/code/wiki>.
 - [32] OCZ RevoDrive 3 X2 (EOL) PCI Express (PCIe) SSD. <http://ocz.com/consumer/revodrive-3-x2-pcie-ssd>.
 - [33] Open-OSD project. <http://www.open-osd.org>, 2013.
 - [34] HDF5 Tutorial: Parallel HDF5 Topics. <http://hdf.ncsa.uiuc.edu/HDF5/doc/Tutor/parallel.html>.
 - [35] J. Piernas, J. Nieplocha, and E. J. Felix. Evaluation of Active Storage Strategies for the Lustre Parallel File System. In *Proceedings of the 2007 ACM/IEEE conference on Supercomputing*, 2007.
 - [36] L. Qin and D. Feng. Active Storage Framework for Object-based Storage Device. In *Proceedings of the 20th International Conference on Advanced Information Networking and Applications*, 2006.
 - [37] E. Riedel, G. Gibson, and C. Faloutsos. Active Storage for Large Scale Data Mining and Multimedia Applications. In *Proceedings of 24th Conference on Very Large Databases*, 1998.
 - [38] M. T. Runde, W. G. Stevens, P. A. Wortman, and J. A. Chandy. An Active Storage Framework for Object Storage Devices. In *Mass Storage Systems and Technologies (MSST), 2012 IEEE 28th Symposium on*, 2012.
 - [39] Samsung SSD. <http://www.samsung.com/uk/consumer/memory-cards-hdd-odd/ssd/830>.
 - [40] C. Sar and P. Cao. Lineage File System. <http://crypto.stanford.edu/cao/lineage.html>, 2005.
 - [41] F. B. Schmuck and R. L. Haskin. GPFS: A Shared-Disk File System for Large Computing Clusters. In *FAST*, 2002.
 - [42] P. Schwan. Lustre: Building a File System for 1000-Node Clusters. In *Proceedings of the 2003 Linux Symposium*, 2003.
 - [43] SDSS-III DR12. <http://www.sdss.org>.
 - [44] M. Singh and B. Leonhardi. Introduction to the IBM Netezza warehouse appliance. In *Proceedings of the 2011 Conference of the Center for Advanced Studies on Collaborative Research*, 2011.
 - [45] HDF 4.1r3 User's Guide. http://hdf.ncsa.uiuc.edu/UG41r3_html/.
 - [46] Spallation Neutron Source | ORNL Neutron Sciences. <http://neutrons.ornl.gov/facilities/SNS/>.
 - [47] S. W. Son, S. Lang, P. Carns, R. Ross, R. Thakur, B. Ozisikyilmaz, P. Kumar, W.-K. Liao, and A. Choudhary. Enabling Active Storage on Parallel I/O Software Stacks. In *Mass Storage Systems and Technologies (MSST), 2010 IEEE 26th Symposium on*, 2010.
 - [48] SQLite. <https://sqlite.org/>.
 - [49] DAGMan: A Directed Acyclic Graph Manager. <http://research.cs.wisc.edu/htcondor/dagman/dagman.html>.
 - [50] Introducing Titan. <https://www.olcf.ornl.gov/titan/>.
 - [51] D. Tiwari, S. Boboila, S. S. Vazhkudai, Y. Kim, X. Ma, P. J. Desnoyers, and Y. Solihin. Active Flash: Towards Energy-Efficient, In-Situ Data Analytics on Extreme-Scale Machines. In *Proceedings of the 11th USENIX Conference on File and Storage Technologies (FAST'13)*, 2013.
 - [52] Top 500 Supercomputer Sites. <http://www.top500.org/>.
 - [53] G. Velampampil. Data Management Techniques to Handle Large Data Arrays in HDF. Master's thesis, Department of Computer Science, University of Illinois, Jan. 1997.
 - [54] Y. Wang, T.-H. Ahn, Z. Li, and C. Pan. Sipros/ProRata: A Versatile Informatics System for Quantitative Community Proteomics. *Bioinformatics*, 29(16), 2013.
 - [55] R. O. Weber. Information Technology - SCSI Object-Based Storage Device Commands (OSD). *Technical Council Proposal Document*, 10:201–225, 2004.
 - [56] Y. Xie, K.-K. Muniswamy-Reddy, D. Feng, D. D. E. Long, Y. Kang, Z. Niu, and Z. Tan. Design and evaluation of Oasis: An Active Storage Framework Based on T10 OSD Standard. In *Proceedings of the 2011 IEEE 27th Symposium on Mass Storage Systems and Technologies*, 2011.
 - [57] Y. Yu, M. Isard, D. Fetterly, M. Budiu, Ú. Erlingsson, P. K. Gunda, and J. Currey. DryadLINQ: A System for General-Purpose Distributed Data-Parallel Computing Using a High-Level Language. In *OSDI*, volume 8, pages 1–14, 2008.