FlexAnalytics: A Flexible Data Analytics Framework for Big Data Applications with I/O Performance Improvement

Hongbo Zou^{*}, Yongen Yu[†], Wei Tang[§], Hsuanwei Michelle Chen[†]

^{*}Queensland University of Technology, Brisbane, Qld 4001, AUS

[†]Illinois Institute of Technology, Chicago, IL 60616, USA

[§]Argonne National Laboratory, Argonne, IL 60439, USA

[†]San Jose State University, San Jose, CA 95192, USA

Abstract—Increasingly larger scale applications are generating an unprecedented amount of data. However, the increasing gap between computation and I/O capacity on High End Computing machines makes a severe bottleneck for data analysis. Instead of move data from its source to the output storage, in-situ analytics processes output data while simulations are running. However, in-situ data analysis incurs much more computing resource contentions with simulations. Such contentions severely damage the performance of simulation on HPE. Since different data processing strategies have different impact on performance and cost, there is a consequent need for flexibility in the location of data analytics. In this paper, we explore and analyze several potential data-analytics placement strategies along the I/O path. To find out the best strategy to reduce data movement in given situation, we propose a flexible data analytics (FlexAnalytics) framework in this paper. Based on this framework, a FlexAnalytics prototype system is developed for analytics placement. FlexAnalytics system enhances the scalability and flexibility of current I/O stack on HEC platforms and is useful for data pre-processing, runtime data analysis and visualization, as well as for large-scale data transfer. Two use cases – scientific data compression and remote visualization have been applied in the study to verify the performance of FlexAnalytics. Experimental results demonstrate that FlexAnalytics framework increases data transition bandwidth and improve the application End-to-End transfer performance.

Keywords—I/O Bottlenecks; In-situ Analytics; Data Preparation; Big Data; High-end Computing

I. INTRODUCTION

Big data applications running on High-End Computing (HEC) machines are generating large volumes of data in a single execution, and these data volumes are only expected to increase in future. Since this massive data is key to scientific discovery, the ability to rapidly store, move, analyze, and visualize data is critical for scientists' productivity. Yet the growth of data volume imposes numerous requirements on I/O performance, which results in I/O bottlenecks in current High-End Computing (HEC) machines. Furthermore, the enlarged gap between computational power and I/O performance further worsens I/O performance. These two folds combined together to result in undesirable situations, where I/O bottlenecks devastate the efficiency and scaling of scientific simulations and associated data analyses and visualizations. Scientists are forced to wait a substantial portion of the simulation runtime

writing data to the storage [16] or simply forgo writing out data in order to keep total I/O within reasonable bounds. These trends will be speeded up in the planning for exascale computing platforms in which the attainable I/O performance is further exacerbated by increased contention on shared resources [31][32][39] on those platforms.

Figure 1. Scientific Data Analytics System Architecture Overview

To improve I/O performance, in-situ data analytics has emerged as an effective way to substitute for the traditional offline data analytics and overcome the increasingly severe I/O bottleneck for scientific applications running at the petascale and beyond. Through processing data before placing it on disk, in-situ analytics can reduce I/O costs (both in time and in power), extract and deliver valuable insights from live simulation output in time, and gain improved end-to-end performance. The utility of in-situ processing is demonstrated by its wide use by leading scientific applications, like the Metagenomics [22], the Cosmology Simulation [27][28], and the Maya [20]. Furthermore and enabled by standard parallel I/O interfaces like ADIOS and MPI-IO [5], there are emerging infrastructures that support online data analytics, including FastBit [29], FlexQuery [36], and others [7][21][34].

Speeding up data processing to keep pace with simulation pushes the in-situ analytics to be separated from simulation and parallel executed on independent computational resources, such as shared compute cores (helper-core), separate nodes dedicated to analytics (termed ‘staging nodes’), or I/O nodes. The separated locations of placing simulation and data analytics, however, introduce the extra data movement along the I/O path. The challenge becomes how to balance the conflict between parallel using limited and costly computational resources versus the extra data movement cost. Figure 1 depicts a typical scientific data analytics architecture. It is comprised of the scientific simulation and its associated data pre-analytics functions for processing simulation results. The simulation is run on some high end HPC system, and data visualization is performed on the visualization cluster. The output of analysis or

visualization directly stores in storage system for science users access through Wide Area Networks (WAN) connection. Our study takes the further step of ‘augmenting’ the data analytics to any data processing techniques (such as data compression, query) for data reduction purpose.

Data compression has been proven to be a very effective data analytics method in improving end-to-end performance in the wide-area networking domain, in which the large size data are compressed and transferred with the worst networks bandwidth [15][37][40]. In HEC domain, the applicability of data compression has been studied and applied to many [23][35][37]. In addition, exploration and visualization of scientific data often are described as ‘visualization queries’ posed on simulation outputs. Typical queries include point selection, subsetting, cutting planes, global data summaries, and feature extraction. Such queries often return only subsets of the original data and can significantly reduce the amounts of data that have to be moved and/or further processed for display. Modern visualization systems exploit early data reduction within their visualization engine to optimize visualization performance [36]. Examples include ‘contracts’ in VisIt [25], ‘events’ in VTK [24], and ‘rendering information’ in ParaView [7]. So, data compression and visualization query can be considered as two kinds of effective in-situ data analytics for data reduction.

This paper presents a detailed assessment of using data compression and visualization query to reduce data movement cost along I/O path, especially with flexible placement strategies. Because the data compression and visualization query incur computational resource contention with simulation, such analytics servicing for data reduction should be carefully studied and devised. Understanding when, where, and how to use such data analytics in I/O path so as to improve the end-to-end application performance is the major goal of this paper. In this paper, we make the following contributions: (1) we present a latency based quantitative model to dynamically make a decision whether should reduce data for next step and what data analytics should be selected to apply on data with the comprehensive consideration on the cost of both I/O and computational resource. (2) we analyze the end-to-end latency of four potential analytics placement strategies with the above quantitative model. (3) we propose a flexible analytics framework and implement a prototype system for such framework. (4) we conduct experiments to show that flexibly choosing the best data movement mode in various resource changes can improve total execution time compared to the execution time without this service.

The reminder of the paper is organized as follows. Section 2 presents background information about data compression, visualization query, and in-situ analytics. Section 3 describes the analytics quantitative model and latency analysis of four placement strategies with model. Section 4 is the implementations of flexible analytics placement system with the integration with I/O middleware. Section 5 presents flexible analytics system built for the GKW data and the measurements of performance and cost. Parallel analytics is checked on this section for performance

optimization. Section 6 reviews related work, and Section 7 concludes the paper.

II. BACKGROUND

To develop a flexible data analytics framework, compression methods, data visualization query, and data analytics placement strategies are briefly introduced here. Data compression and visualization query are considered as two common data analytics to reduce data here. Data analytics placement strategies lists four potential placement strategies for the typical scientific data analytics architecture.

A. Data Compression

Data Compression aims to reduce redundant information in data to save the consumption of resources such as storage space or I/O bandwidth, and therefore this technique has very important application in the areas of data storage and transmission [6]. Apart from the space savings, data compression provides another benefit that compressed data could be stored and transmitted faster, thereby improving the performance of I/O intensive workloads. However, compression is a compute-intensive operation and imposes an extra processing cost and resource contention on compute node [15]. Therefore, the design and choosing of data compression algorithms involve trade-offs among various factors, including the degree of compression, the amount of distortion introduced (using in lossy data compression), and the computational resources required to compress and decompress the data [3]. To facilitate the following discussion in the paper, a brief introduction will be given about lossy and lossless compression algorithms.

Lossless data compression algorithms usually exploit statistical redundancy to represent data more concisely without losing information, and can reconstruct the original data exactly from the compressed data [18]. There are two key metrics – compression ratio and (de)compression speed that are used to evaluate the performance of lossless data compression algorithms. Different compression algorithm has different compression ratio. Even the same compression algorithm could have varying compression ratio with different original data format. In general, the compression ratio of text format data is higher than binary format, and the compression ratio of repetitive data is higher than random data. Processor power has direct relationship with compression speed. The compression speed is also affected by compression buffer, but we don’t want to extend this discussion at here.

Lossy data compression is contrasted with lossless compression and can reconstruct an approximation of the original data. In these algorithms, some loss of information is acceptable [10]. The lossy data compression is also could be evaluated by compression ratio and (de)compression speed. Compression accuracy is a new metric involved by lossy compression. Accuracy is used to judge the difference between original data and decompressed data. The inherited deficiency of information loss limits the application area of lossy compression on lossless requirement. This paper doesn’t explore compression algorithms in depth. We use a

lossy compression developed by [11] to evaluate our flexible analytics framework.

B. Data Visualization Query

Exploration and visualization of scientific data may be described as ‘data queries’ posed on simulation outputs. Typical queries include point selection, subsetting, cutting planes, global data summaries, and feature extraction. Such queries often return only subsets of the original data and can significantly reduce the amounts of data that have to be moved and/or further processed for display. Modern visualization systems exploit early data reduction within their visualization engine to optimize visualization performance [36]. Examples include ‘contracts’ in VisIt [25], ‘events’ in VTK [24], and ‘rendering information’ in ParaView [7].

TABLE I. DATA REDUCTION QUERIES

Query	Type	Description
Max, Min, Average, Sum, Stdev	Statistics	Max/Min/Average/Sum/Stdev (var_double_array)
Range	Similarity	Range(variable, dimensions, start_pos, end_pos)
Select, Lineout	Visualization	Select(variable, threshold1, threshold2) Lineout(variable, start_pos, end_pos)
Expression	Combination	Self-defined function, select (lineout (variables, start_pos, end_pos), threshold1, threshold2)

In this study, adaptive data reduction uses only four types of data reduction queries, as shown in Table I. Statistics queries calculate the features of a given variable. Similarity queries are used to filter out the values with close features. Visualization queries are similar to the former in terms of implementation, but the specified features are those relevant for visualization. Combination queries provide a simple pipeline method to help users combine basic queries to construct their won data requests. These four query categories have been widely adopted by the visualization community and provide the basic quantitative analysis for data visualization [36].

C. Data Analytics Placement Strategies

To explore the possible places on where data analytics can be placed in scientific data analytics architecture, analytics placement strategies is introduced here. In general, data analytics can be simply classified into four types according to execution time and location [33].

1) *Inline Processing*: analysis/visualization routines are synchronously performed by the simulation. ParaView [7], VisIt [25] and other in situ visualization works [36] fall into this category. Inline processing is easy to code and implement by scientists. The sequential execution of simulation and analytics introduces extra data waiting time for both of their execution.

2) *Helper-core Processing*: some cores on the compute nodes are dedicated to perform select analysis actions.

Examples include Functional Partitioning [12] and Software Accelerator. Helper cores parallel analyze and reduce data immediately after data produced. It involves some resource contention, such as memory, to simulation; In 1) and 2), data compression and analytics services share compute nodes with simulation. The compressed data is prepared for the further processing or storage saving.

3) *Dedicated-nodes Processing*: the analytics are executed on non-simulation nodes, such as staging area and active storage[17]. Remote Processing needs additional compute resource for analytics and incurs the extra data movement cost.

4) *Offline Processing*: data is read back for analysis [8] after written to storage. Offline processing can statically provide enough resource and tools to analyze data in details.

Any one of the above four placement strategies can be selected and executed independently. Or, some of placement strategies can be combined together and executed hybrid data processing strategy. The details will be extended in Section III according to our flexible data analytics framework design requirements.

III. METHODOLOGY

HEC machine is a complex system, which generally consists with three components – compute nodes, analytics nodes, and storage. To reduce I/O overheads on applications, some HEC machines allocate part of compute nodes to conduct staging functionality, called staging nodes. Data analytics can be placed and executed on staging nodes to reduce data before writing data into storage.

A. Flexible Analytics Placements

We introduce data analytics into peta-scale computing to reduce the data movement. To achieve this goal, we investigate where data analytics can be placed among I/O path. In general, a scientific application includes two parts – simulation and analytics, some application including visualization. Simulation and inline analytics is executed on compute node after an application submission. Among application execution, the output data is dumped to I/O nodes for storage saving. As shown in figure 2, data analytics can be placed at point 1 to reduce the output dumping to storage.

Figure 2. Potential locations of reduction analytics located in I/O path

If data analytics is placed on staging nodes, there are numerous data movements produced between compute nodes to staging nodes. The point 2 in figure 2 shows the potential position to place analytics to reduce the data size moving from compute nodes to staging nodes. Such case is the primary study case in the paper. After the data analysis, the reduced data needs to be written into storage for saving. In

figure 2, Point 3 shows that the data analytics is used to reduce movement cost between staging nodes and I/O nodes.

The listed three points in figure 2 show the potential places for data analytics. Adaptive placement algorithm will find out the best place to apply the analytics to the data and reduce the movement cost.

B. Analytics Algorithms Quantitative Model

To evaluate and compare the performance of different analytics method and placement strategies, we quantitate the analytics algorithms in this section. Because data queries are explicitly requested by client terminal with special use case, we only show the compression quantitative model for general data reduction request here. A set of performance and cost metrics are defined for quantitative model description. We then derive a simple performance model to compare the different compression algorithms. The model has been developed and successfully applied in practice to implement adaptive compression choosing.

TABLE II. NOTATIONS IN PERFORMANCE QUANTITATIVE MODEL

Notation	Description
$\rho^1 = \rho_c$ ρ_d	Available Processor Cycles Ratio (< 1 , sharing processor with other processes, > 1 using multiple processors to parallel execution). ρ_c is available cycles ratio on compression side. And ρ_d is available cycles ratio on decompression side.
BW	Available Data Transfer Bandwidth
B_{data}	The Size of Data Block
T_c, T_d	Compression/ Decompression Throughput
δ	Compression Ratio (compressed / original)
$\varepsilon^1 = \varepsilon_c$ ε_d	The Impact Factor of Cache. ε_c is impact factor in compression side. ε_d is impact factor in decompression side.
$t_{original}$	Total end-to-end latency to original (uncompressed) data transfer
$t_{compression}$	Time to compress the data (algorithm dependent)
$t_{decompression}$	Time to decompress the data (algorithm dependent)
$t_{transfer}$	Transfer time
t_{total}	Total end-to-end data transfer latency

Table II lists the notations used in compression performance quantitative model. Every notation is described and simply explained in the table. However, there are still two notations, ρ and ε , need to be explained here. ρ notates the available processor cycles ratio for compression computation. Because compression is a computation-intensive task, the processor idle cycles ratio indicates how many available processor cycles could be used for compression. $\rho = 1$ when a special core is exclusively used by compression process. Parallel compression is executed on multiple processors if $\rho > 1$. In addition, ε stands for the impact factor related cache. Compression process needs to swap data from memory for compression; therefore the available size of cache space directly impacts compression performance. $\varepsilon = 1$, when there is no other processes share multiple level caches with compression process. To simplify our discussion, the impact factor of cache is set as 1 ($\varepsilon = 1$) in the paper. The following equations calculate the cost (latency) of data transfer from sender to receiver when compression is running.

1. We give ρ and ε in equations to simplify the discussion. And, ρ_c , ρ_d , ε_c and ε_d are distinct when they need to be measured.

Equation (1) gives the end-to-end original data transfer latency on networking with available bandwidth BW . Equation (2) and (4) give compression and decompression time on sender and receiver respectively. Equation (3) gives the transfer time of compressed data from sender to receiver. Equation (5) shows the total cost.

$$t_{original} = B_{data} / BW; \quad (1)$$

$$t_{compression} = B_{data} / (T_c * \rho * \varepsilon); \quad (2)$$

$$t_{transfer} = (B_{data} * \delta) / BW; \quad (3)$$

$$t_{decompression} = B_{data} / (T_d * \rho * \varepsilon); \quad (4)$$

$$t_{total} = t_{compression} + t_{transfer} + t_{decompression}; \quad (5)$$

With the above quantitative analysis, the data transfer costs can be compared with/without compression under the selected compression algorithm. Therefore, the compression decision can be made by *If* statement as following:

```

If (  $t_{original} > t_{total}$  )
{
    Select the algorithm with Min( $t_{total}$ );
    Compress with the selected algorithm;
}
Else
    Don't Compress;

```

To implement *If* statement, the quantitative metrics are classified into two categories with different information collecting method:

Profiling Information:

Compression/ Decompression Speed (T_c, T_d),
Compression ratio (δ).

Monitoring Metrics:

Available Processor Cycle Ratio (ρ),
Available Data Transfer Bandwidth (BW),
The Impact Factor of Cache (ε).

C. Total End-to-End Latency Analysis

To explore the resource cost of the analytics, we analyze the total end-to-end latency in all cases. According to the different transfer destination, I/O path are roughly classified into two categories – memory-to-memory transfer and memory-to-storage one. Because the connection bandwidth of two categories is totally different, the analytics cost varies with the available resource. $t_{original} > t_{total}$ gives the condition that compression is selected. The following analysis is based on this condition.

1) Simulation to Storage (Inline compression): in this case, the data is transferred from computational nodes to the storage. Such case does not include data decompression. Therefore,

$$\begin{aligned}
 t_{original} > t_{total} &\Rightarrow t_{original} > t_{compression} + t_{transfer}; \text{ And,} \\
 B_{data} / BW > B_{data} / (T_c * \rho * \varepsilon) + (B_{data} * \delta) / BW; \\
 &\Rightarrow 1 - \delta > BW / (T_c * \rho * \varepsilon);
 \end{aligned}$$

$$\text{So, } T_c * (1-\delta) * \rho * \varepsilon > BW; \quad (6)$$

In equation (6), δ and T_c (compression ratio) are constants if the compression algorithm and input data format don't change. Therefore, keeping ρ and BW in a proper proportion decides the improvement of end-to-end latency. ρ can be greatly increased with the parallel. For data transmission from simulation to storage, although BW is limited by storage, the simulation can provide available processor-cycles for inline compression to offset such limitation. For example, if $BW = 400 \text{ Mb/s}$, $T_c = 100 \text{ Mb/s}$, $\delta = 50\%$, and $\varepsilon = 1$, ρ should be equal or larger to 4. $BW = 400 \text{ Mb/s}$ is very common configuration in Supercomputer.

2) Simulations to Staging (Inline compression): in this case, the data is transferred from computational nodes to staging nodes. This case includes data decompression in staging nodes. Therefore equation (6) becomes

$$(T_c * T_d) / (T_c + T_d) * (1-\delta) * \rho * \varepsilon > BW; \quad (7)$$

Actually, $T_c < T_d$ (this happens often in compression algorithms), and compression and decompression are executed in parallel with multiple blocks transfer, therefore ρ_c can be calculated with equation (6) and ρ_d in decompression side can be derived by

$$\rho_d = T_c * \rho_c / T_d; \quad (8)$$

Equation (7) gives us the same conclusion with equation (6). If the BW is very large between computational nodes and staging nodes with high-speed networking connection, compression becomes an expensive operation in this case. For example, if $BW = 16 \text{ Gb/s}$ (16384 Mb/s), $T_c = 100 \text{ Mb/s}$, $T_d = 200 \text{ Mb/s}$, $\delta = 50\%$, and $\varepsilon = 1$, ρ should be equal or larger to almost 491.520. $\rho \geq 491.520$ stands for that the computational area needs more than 492 processors for data compression. With equation (8), there are at least 246 processors to decompression. The values given in example are very common configuration in HEC.

3) Helper cores to Staging: this case is same with case 2). Case 2) and 3) stand for high-speed memory-to-memory data transfer. Such available bandwidth needs fast data compression to catch the small end-to-end latency.

4) Staging to Storage: this case is same with case 1). Case 1) and 4) stand for memory-to-storage data transfer. Compression becomes very effective method to release the I/O bottleneck in this case.

IV. DESIGN AND IMPLEMENTATION

A. Overview

The quantitative model clears that different use case requests different data analytics to reduce data. Many factors, such as transfer bandwidth, available memory, and data format etc., impact the end-to-end performance. With this conclusion, a flexible data analytics framework is developed. Figure 3 shows the flexible data analytics system framework (FlexAnalytics). Such framework can dynamically deploy data analytics on the I/O path to overcome the bandwidth limitations between data source (simulation) and sink

(storage/user terminal). FlexAnalytics includes five components: Policy Engine, Resource Monitoring, Analytics Profiling, Analytics Placement Plug-in, and Adapter. Figure 4 shows the FlexAnalytics software stack. FlexAnalytics clearly separates the simulation, analytics, and storage/user terminal into two parts: the data source, which is the outputs of the simulation performed on the high end HPC system, and the data sinks, which consume the output data with data storing and visualization.

Figure 3. Flexible Data Analytics System Framework

FlexAnalytics leverages the ADIOS [9] parallel I/O library which provides meta-data rich read/write interfaces to simulation and analysis codes. ADIOS has a set of built-in I/O methods under its higher level APIs to support various file I/O methods (such as MPI-IO, HDF5, and NetCDF) as well as data staging methods. ADIOS has been used by several leadership scientific codes and is integrated with popular analysis and visualization tools that include ParaView, VisIt.

Figure 4. Flexible Data Analytics Software Stack

B. Data Compression Profiling

We profile end-to-end latency with 222MB float data transferring between two nodes with three data compression algorithms – lossy [11], bzip2 [1], and gzip [4]. The profiling results can help us to evaluate the selected compression algorithm. Three compression algorithms emphasize the different requirements on the compression speed and compression ratio. Bzip2 and gzip are lossless data compression. Bzip2 is an implementation of the Burrows-Wheeler algorithm, which is more effective on data compression ratio. Gzip is based on the DEFLATE algorithm, which provides better compression speed and worse compression ratio than bzip2. We check the performance of lossy algorithm described in [11]. The lossy compression is better than lossless on speed and ratio. However, information loss constraints its wide use.

TABLE III. COMPRESSION ALGORITHM FEATURES PROFILING FOR BINARY DATA ON TWO NODES MEASUREMENT

	lossy	bzip2	gzip
T_c	51.002 MB/s	5.417 MB/s	10.991 MB/s
T_d	242.468 MB/s	8.703 MB/s	23.632 MB/s
δ	23.46%	46.18%	71.75%
ε	Approximate Linear (we generally set $\varepsilon = 1$ as impact factor to to eliminate the cache impact with cache exclusively using in this benchmark testing.)		

The testing nodes are equipped with 2.13GHz Intel Xeon Processor with 4MB L3 cache and 12GB memory. The nodes connected with 16GB infiniband networks. Table III lists the profiling information of three compression algorithms: lossy, bzip2, and gzip. Lossy compression will be introduced in experimental section.

Figure 5. End-to-End Latency Comprison with Different Compression

To evaluate the behavior of compression with real-world use cases, we measure end-to-end latency on varying transfer bandwidth. The compression and decompression have been exclusively executed on one single processor. Figure 5 shows the measurement results. The transfer latency of raw data without compression is shown in figure as baseline.

The results shows the (de)compression time of lossless data compression accounts for over 98% of end-to-end latency when the transfer bandwidth is larger than 1.034Gb/s. For lossy compression, the (de)compression accounts for over 95% of end-to-end latency. If we compress the data in parallel, there are at least 4 processors for lossy compression, 63 processors for gzip, and 74 processors for bzip2 to catch up the data transfer with 1.034Gb/s bandwidth. When the transfer bandwidth reduces to 0.258 Gb/s, the proportion of (de)compression in latency starts to decrease. The latency of data compression and transfer with gzip is better than non-compression after the available bandwidth less than 0.016Gb/s. The measured results verify our quantitative model. In our case, all three compressions are not effective method to optimize data transfer when the bandwidth is larger than 264.19Mb/s. Therefore, the new challenge is how to reduce data (de)compression time to match the high-speed transfer in HEC.

C. Visualization Query Profiling

It is similar with compression analytics. We profile end-to-end latency with 222MB float data transferring between

two nodes with three different visualization queries – statistics, similarity, and combination. Three queries have different workload for computing resource requirement. Statistic query(ang_vel) is a lightweight query, which requests the minimal value of angle velocity in the data set. Because the return results are very small data set, the statistic query can significantly increase the transmission bandwidth online. We inject controlled additional network traffic between two nodes variability, resulting in available bandwidth varying from 20MB/s to 120MB/s. Figure 6 (a) shows the latency comparison of 222 MB data transfer with and without statistics query. The lightweight query has great data reduction ratio in our profiling. Therefore, the latency is significantly reduced over 150 times when the available bandwidth is less than 20 MB/s.

Figure 6. Data Query Features Profiling with Different Query Types

Figure 6 (b) and (c) show the latency comparison of similarity and combination query on temperature and velocity mesh. These two queries stand for heavy load and joint query, which request much more computing resource for data filtering. So, such queries are good for data reduction with high available transfer bandwidth. 20MB/s available bandwidth is a special case in the profiling. Although similarity and combination queries are heavy load operations, the low available bandwidth still needs data analytics to improve I/O performance. Because the data

query is the operation which is explicitly specified by user terminal, the query cannot be applied on the output data in general cases.

D. Flexible Analytics Placement Algorithm

The flexible analytics placement algorithm places the data analytics to the best place, to reduce data movement cost, subject to resource availability and performance penalties due to the potential need to share the source's compute resources with other simulation.

TABLE IV. PROFILING AND MONITORING METRICS

<i>Profiling Metrics (Profiled on idle nodes)</i>	
Notation	Parameter
t_m	Data Transfer Time per Iteration
t_c	Query Execution Time per Iteration
c_p	Processor Usage (percentage) for the Analytics
m_p	Memory Usage (percentage) for the Analytics
<i>Monitoring Metrics (Periodically collected on the candidate place)</i>	
Notation	Parameter
T_c	Data Transfer Time per Iteration
C_p	Query Execution Time per Iteration
M_p	Processor Usage (percentage) for the Analytics

Placement takes as inputs the potential locations, resource status, and analytics performance profiles. Table IV shows the profiling and monitoring metrics used as input. Profiling metrics are collected by placing analytics onto some idle node, which provides the policy engine with baseline data about the analytics' running (this should be changed with changed analytics method, especially for data query). For example, if a analytics requires 35% of compute resource for its computation, placement must select a place that is presently utilized less than 65%. Profiling and monitoring are collected on the candidate places for decision-making.

The algorithm 1 operates as follows. When the data is generated by the simulation, an analytics is selected to reduce the data on the idle node, thereby profiling it. Next, the algorithm checks the potential places (as shown in the pseudo-code). Monitoring information $(C_p, M_p)_i$ and the profiling metrics (t_m, t_c, c_p, m_p) are collected by algorithm on all potential places. If the analytics is a compute-intensive task, the algorithm needs to allocate appropriate computational resources to execute the query in parallel (as shown on the algorithm, we currently consider an analytics task to be compute intensive if its processor utilization is 90% or above when it is profiled). A placement decision diverging from the query's initial placement onto an idle node involves analytics relocating. If the analytics is relocated, analytics execution time T_c is again checked on the new place. Placement decisions are re-evaluated at regular time intervals, an example being an analytics relocation from the current (busy) place to the idle place, i.e., if $T_c > t_c + t_m$. The pseudo-code of algorithm 1 shows the main idle of flexible analytics placement.

Data: Input data per iteration
Input: An analytics
Result: Place the analytics to the best location

Select the potential locations ($N_1, N_2, N_3, \dots, N_i$)
Place the analytics on an idle node to profile its performance
for every iterations do
 collect $(C_p, M_p)_i$ on every potential place N_i ;
 $\text{Max}(\text{CM}) = (C_p, M_p)_i$; // N_i is the place with maximum resources
 if analytics on the idle node then
 profile (t_m, t_c, c_p, m_p) on the idle node;
 if $c_p \geq 90\% \parallel m_p \geq 90\%$ then
 execute analytics in parallel;
 else
 if $\text{Max}(\text{CM}) > (c_p, m_p)$ then
 relocate the analytics to the place with $\text{Max}(\text{CM})$;
 end
 else
 collect (T_c) on the location where analytics is running;
 if $T_c > t_c$ then
 relocate the analytics to the place with $\text{Max}(\text{CM})$;
 collect (T_c) on the new place;
 end
 if $T_c > t_c + t_m$ then
 relocate the analytics back to the previous place;
 end
 end
 end

Algorithm 1. Flexible Analytics Placement

E. Placement Policy Engine

The placement policy engine decides where to execute each analytics requested by the terminal adaptor. The placement decision is enforced by deploying the analytics code to the best place among the data path. The flexible analytics placement algorithm is performed in the placement policy engine.

V. EXPERIMENTS AND ANALYSIS

Experiments with the Gyro-kinetic Simulation (GKW) [30] application are run on an 80-node Linux cluster. A SGI Altix UV 1000 systems is used as ancillary visualization machine for output data analysis. Every node of simulation cluster is equipped with four quad-core 2.0GHz AMD Opteron processors, 32 GB of memory, a gigabit Ethernet networking connection. The compute node runs CentOS 5.7 OS. GKW is a three-dimensional Particle-In-Cell (PIC) code used to study the microburbulence. SGI Altix system is a shared memory system, which has 1024 2.0GHz cores (Intel Nehalem EX processors), 4 TB of global shared memory, and 8 GPUs in a single system image. The data movement throughput between two machines via shared Lustre file system can reach at most 200MB/s. GKW simulation outputs particle data containing two particle arrays for zions and electrons respectively. The particle arrays consist of millions of unique double precision values. The values are high degree of randomness and are the most "hard-to-compress" scientific dataset [11].

TABLE V. COMPRESSION ALGORITHM FEATURES PROFILING FOR GWK ON 80-NODE CLUSTER

	lossy	bzip2	gzip
T_c	49.672 MB/s	2.491 MB/s	16.242 MB/s
T_d	819.767 MB/s	6.523 MB/s	108.698 MB/s
δ	24.33%	97.44%	94.08%

In our configuration, GWK submits four processes in every node. Every process generates output data for further data processing. Four analytics processes are created to service data transfer in every node. To find out the metrics for analytics selection, compression algorithms are profiled. Table V shown the profiled features of compression algorithms in the cluster. All results presented in the paper are averaged with at least three runs.

A. Bandwidth Analysis from Simulation to Storage

In the experiment, the compression is deployed on computational side to reduce data size. The available end-to-end bandwidth reaches to 62.74MB/s with 64 computational nodes. We compare the expected, estimated, and real-tested bandwidth among end-to-end connection. Expected bandwidth is the best one, which considers compress time is very small and negligible. Therefore, $BW_{expected} = BW_{IO} / \delta$. The estimated bandwidth is calculated by equation (6) in the paper. Estimated bandwidth shows the maximal optimization of data transfer with parallel compression. Any data transfer bandwidth is smaller than the $\min\{BW_{expected}, BW_{estimated}\}$. The experiment is mapped to two scenarios: simulation to storage and staging to storage. Both of them can be summarized as memory to storage data transfer model. Figure 7 shows the tested bandwidths with/without compression.

Figure 7. Data Transfer Bandwidth Comparison Without Staging

B. Bandwidth Analysis from Simulation to Staging nodes

In the experiment, the output of GWK is written to staging nodes. The compression service is running on computing nodes to reduce the data size. The available end-to-end bandwidth reaches to 132.94MB/s with 64 computational nodes. We configure 4 staging nodes in the experiment. The experiment simulates two cases: simulation to staging and helper core to staging. Both of them can be summarized as memory-to-memory transfer model. Figure 8 shows the tested bandwidth with or without compression.

Figure 8. Data Transfer Bandwidth Comparison With Staging

In the experiment, multiple data blocks transfer from simulation nodes to staging nodes. (De)compression and compression are executed on computational nodes and staging nodes separately. Because compression processes are executed on same compute node and share multi-level cache hierarchy (including L3 cache), the cache impact factor should be ($\epsilon = 1/2$).

C. Latency Analysis with Data Query

For data visualization, end users are interested in i) query latency - the elapsed time between when data is available to when query output appears, and ii) simulation I/O throughput, which is impeded if data cannot be moved off the high end machine before the application's next output step - due to application blocking on output.

Figure 9. Data Request Response Time Test with/without Query

Figure 9 shows the performance impact of flexible analytics placement on data transfer latency, given variable connection bandwidth between the visualization machine and the simulation. We inject controlled additional network traffic between simulation cluster and ancillary visualization machine, resulting in available bandwidth varying from 20MB/s to 120MB/s. The GWK simulation generates 1GB output data per output step.

We compare two cases: one is to move raw simulation output data to ancillary visualization machine and run query locally on it; the other is to place the query on the simulation cluster to reduce data movement between the two machines. Figure 9 shows that the FlexAnalytics system can significantly improve query latencies by reducing data movement volumes for all three tested queries. Even when the available bandwidth between two machines is at its maximum (120MB/s), there is 23.7% improvement. When

available bandwidth decreases to 20MB/sec, the improvement brought by FlexAnalytics can be up to 80%.

D. Throughput Analysis with Data Query

Simulation execution throughput (the amount of simulation output over a given interval of time) is impaired when output data cannot be moved off the high end machine between subsequent simulation time steps (i.e., due to blocking on I/O). Figure 10 shows, under 120MB/s communication bandwidth, the throughput of the simulation running at different scales. When the simulation scales beyond 432 cores, the data movement cost become significant enough so that placing queries near data source becomes beneficial. At the scale of 1000+ cores, the advantage of FlexAnalytics is as large as 100%.

Figure 10. The Improvement of Simulation Throughput

Results show the FlexAnalytics approach to be capable of placing queries so as to alleviate communication bottlenecks between data sources and analytics processing sites.

VI. RELATED WORK

In-situ data analytics has been widely used in many data-intensive applications to reduce data size. Data compression and query is useful because it helps minimize storage utilization, reduce networks bandwidth and energy consumption in hardware. However, most of compressions and queries are not transparently operated at the I/O layer, but rather explicitly worked at the application level [26] [36][40].

In many recent studies[11][26], compression is applied on big data research to reduce transfer latency and improve response time. All of them concentrate on the study on compression algorithm and data saving on parallel file system. Our work differs from these compression studies, as it does not focus on the study of compression algorithms. The exploration of compression selection is conducted in the paper. We investigate the effects of data compression on the whole I/O path rather the file system to find the best location to place the compression. We show that the compression improves network end-to-end transmission time, network bandwidth, and consequently the I/O performance with the well-planned placing.

The area of data compression in cloud computing has received many attentions in recent literature. A recent study [3] investigated the effects of compression in MapReduce

clusters. This study focused on increasing I/O performance in order to reduce cluster power consumption. Our work focuses on improving I/O performance to achieve better End-to-End application performance. Another data compression studies focus on applying the data compression into grid computing. The highly parallel-distributed data management in grid has the different requirements on data transmission with HEC. Such research is designed for use in Grid computing environments. So, the existing data compression methods cannot be applied on HEC directly.

There are a number of related efforts that provide solutions to data visualization and analysis in the HPC community [1][13][36]. They all use the data sever/client architecture supporting visualization by permitting remote users to acquire images via their visual client from visual servers located on large-scale machines placed ‘next’ to simulation. The problem of how to move data from the simulation to the visualization machine remains to be unsolved. Our work addresses it. In addition, related work on query-driven visualization [19] considers static data queries rather than ad hoc, dynamic queries enabled by FlexAnalytics. There are also contributions on data indexing, which we have not yet leveraged. Instead, we flexibly place analytics among I/O path to improve both simulation and data access performance. Additional input to our work can be derived from our previous work on adaptive compression and FlexQuery [36][38].

VII. CONCLUSIONS AND FUTURE WORK

In this paper, we studied questions about how to introduce data analytics into HEC to reduce data movement and release severe I/O performance bottleneck. To explore the possible solutions, a quantitative model is built to evaluate the potential analytics algorithms and placement strategies. Based on this model, we propose an flexible analytics framework for I/O performance optimization. This flexible placement strategy combines data compression and visualization query - two algorithms, which can be dynamically selected and switched to achieve the best I/O performance with features profiling and real-time system resource status monitoring.

The experiments with the real application of GKW are conducted on an 80-node 1280 cores cluster machine and a SGI visualization machine. The analysis of results shows us that the data reduction ratio and available processors are two most important factors to impact the analytics selection. The experiments investigate these two factors in details. Thus, our future work should focus on the improvement of on these two factors.

VIII. ACKNOWLEDGEMENTS

The work at Argonne is supported in part by the U.S. Department of Energy (DOE), Office of Science, under Contract DE-AC02-06CH11357.

REFERENCES

- [1] J. P. Ahrens etc. “Interactive remote large-scale data visualization via prioritized multi-resolution streaming,” Proc. The 2009 workshop on Ultrascale Visualization, UltraVis’09, 2009.

- [2] High-quality Data Compressor, bzip2, <http://www.bzip.org>.
- [3] Y. Chen, A. Ganapathi etc., "To Compress or Not to Compress Compute vs. I/O Tradeoffs for MapReduce Energy Efficiency," Proc. The SIGCOMM Workshop on Green Networking, 2010, pp. 23-28.
- [4] Compression Utility, Gzip, <http://www.gzip.org>.
- [5] MPI Forum. MPI-2: Extension to The Message-Passing Standard.
- [6] D. A. Lelewer, and D. S. Hirschberg, "Data Compression," Proc. ACM Computing Surveys (CSUR), 1987, pp. 261-296.
- [7] N. Fabian, K. Moreland etc., "The ParaView Coprocessing Library: A Scalable, General Purpose In Situ Visualization Library," Proc. IEEE Symp. on Large-Scale Data Analysis and Visualization (LDAV2011), 2011, pp. 89-96.
- [8] A. Gerndt, B. Hentschel, M. Wolter, T. Kuhlen, and C. Bischof, "VIRACOCOA: An Efficient Parallelization Framework for Large-Scale CFD Post-Processing in Virtual Environments," Proc. ACM/IEEE Conference on Supercomputing (SC04), 2004, pp. 50-61.
- [9] S. Hodson, S. Klasky, Q. Liu, J. Lofstead, N. Podhorszki, F. Zheng etc., "ADIOS 1.3.1 User's Manual," Oak Ridge National Laboratory, 2011, pp. 1-95.
- [10] I. Kontoyiannis, "Pattern Matching and Lossy Data Compression on Random Fields," Proc. IEEE Transactions on Information Theory, 2003, pp. 1047-1051.
- [11] S. Laksh, N. Shah, S. Ethier, S. Klasky, R. Latham, R. Ross, and N. Samatova, "Compressing the Incompressible with ISABELA: In-situ Reduction of Spatio-temporal Data," Proc. International European Conference on Parallel and Distributed Computing (Euro-Par 2011), 2011, pp. 366-379.
- [12] M. Lin, S. S. Vazhkudai, A. R. Butt, F. Meng, X. Ma, Y. Kim, C. Engelmann, and G. Shipman, "Functional Partitioning to Optimize End-to-End Performance on Many-core Architectures," Proc. ACM/IEEE Conference on Supercomputing (SC10), 2010, pp. 1-12.
- [13] K.-L. Ma, "A New Approach to Remote Visualization of Large Volume Data," SIGGRAPH Comput. Graph., 44(3):5:1-5:2, 2010.
- [14] S. N. Laboratory, Cth shock physics, <http://www.sandia.gov/CTH/>.
- [15] B. Nicolae, "High Throughput Data-Compression for Cloud Storage," Proc. The 3rd International Conference on Data Management in Grid and Peer-to-Peer Systems (Globe'10), 2010, pp. 1-12.
- [16] R. Oldfield, S. Arunagiri, P. J. Teller, S. R. Seelam, M. R. Varela, R. Riesen, and P. C. Roth, "Modeling the Impact of Checkpoints on Next-generation Systems," Proc. IEEE Conference on Mass Storage System and Technologies (MSST 2007), 2007, pp. 30-46.
- [17] J. Piernas, J. Nieplocha, E. J. Felix, "Evaluation of Active Storage Strategies for the Lustre Parallel File System," Proc. ACM/IEEE Conference on Supercomputing (SC07), 2007, pp. 1-10.
- [18] K. Sayood, "Introduction to Data Compression," 3rd edition, Morgan Kaufmann, 2005.
- [19] K. Stockinger etc. "Query-driven visualization of large data sets," 2005.
- [20] M. Slawinska, T. Bode, H. Zou etc. "A Maya Use Case: Adaptable Scientific Workflows with ADIOS for General Relativistic Astrophysics," Proc. The Conference on Extreme Science and Engineering Discovery Environment, 2013.
- [21] XH Sun, C. Du, H. Zou, "V-mcs: A Configuration System for Virtual Machines," Proc. IEEE International Conference on Cluster Computing (Cluster'09), 2009.
- [22] W. Tang, J. Wilkening, N. Desai, W. Gerlach, A. Wilke, and F. Meyer, "A Scalable Data Analysis Platform for Metagenomics," in Proc. of IEEE Conference on Big Data, 2013.
- [23] W. Tang, Z. Lan, N. Desai, and Y. Yu etc., "Reducing Fragmentation on Torus-connected Supercomputers", Proc. IEEE International Parallel and Distributed Processing Symp (IPDPS'11), pp 828-839.
- [24] VTK 6.0.0 Documentation. 2013.
- [25] Paraview VisIt, OpenSource Scientific Visualization and Graphical Analysis Tool.
- [26] B. Welton, D. Kimpe, J. Cope, C. Patrick etc., "Improving I/O Forwarding Throughput with Data Compression," Proc. IEEE International Conference on Cluster Computing (Cluster'11), 2011.
- [27] J. Wu, Z. Lan, and Y. Yu etc., "Performance Emulation of Cell-based AMR cosmology Simulations," Proc. ACM/IEEE Conference on Supercomputing (SC'12), 2012.
- [28] J. Wu, Z. Lan, and X. Xiong etc., "Hierarchical Task Mapping of Cell-based AMR Cosmology Simulations," Proc. The International Conference on High Performance Computing, Networking (Cluster'11), 2011.
- [29] K. Wu, S. Ahern etc., "FastBit: Interactively Searching Massive Data," Proc. SciDAC, Journal of Physics: Conference Series, 2009.
- [30] W. X. Wang, Z. Lin, W. M. Tang, W. W. Lee, S. Ethier, J. L. V. Lewandowski, G. Rewoldt, T. S. Hahm, and J. Manickam, "Gyrokinetic Simulation of Global Trubulent Transport Properties in Tokamak Experiments," Proc. Physics of Plasmas, 2006, pp 59-64.
- [31] Y. Yu, D. Rudd, Z. Lan, J. Wu etc., "Improving Parallel IO Performance of Cell-based AMR Cosmology Applications," Proc. IEEE International Parallel and Distributed Processing Symp (IPDPS'12), pp. 933-944, 2012.
- [32] Y. Yu, J. Wu, and Z. Lan etc., "A Transparent Collective I/O Implementation," Proc. IEEE International Parallel and Distributed Processing Symp (IPDPS'13), 2013, pp 297-307, 2013.
- [33] F. Zheng, H. Zou, J. Cao, J. Dayal, T. Nugye, G. Eisenhauer, S. Klasky, "FlexIO: Location-flexible Execution of In-Situ Data Analytics for Large Scale Scientific Applications," Proc. IEEE International Parallel and Distributed Processing Symp (IPDPS'13), 2013, pp 320-331, 2013.
- [34] H. Zou, H. Jin, Z. Han etc. "A Virtual-Service-Domain based Bidding Algorithm for Resource Discovery in Computational Grid," Proc. IEEE/WIC/ACM International Conference on Web Intelligence, pp5-53, 2005.
- [35] H. Zou, H. Jin, Z. Han etc. "HRTC: Hybrid Resource Information Service Architecture based on GMA", Proc. IEEE International Conference on e-Business pp541-544, 2005.
- [36] H. Zou, K. Schwan, M. Slawinska, G. Eisenhauer, F. Zheng etc. "FlexQuery: An Online Query System for Interactive Remote Visual Data Exploration at Large Scale," Proc. IEEE International Conference on Cluster Computing (Cluster'13), 2013.
- [37] H. Zou, W. Wu, XH Sun etc. "An Evaluation of Parallel Optimization for OpenSolaris Network Stack," Proc. IEEE 35th Conference on Local Computer Networks (LCN), pp296-299, 2010.
- [38] H. Zou, Y. Yu, W. Tang, M. Chen, "Improving I/O Performance with Adaptive Data Compression for Big Data Applications", Proc. IEEE International Parallel and Distributed Processing Symp (IPDPS'14), 2014.
- [39] H. Zou, XH Sun, etc. "A Source-aware Interrupt Scheduling for Modern Parallel I/O Systems," Proc. IEEE International Parallel and Distributed Processing Symp (IPDPS'12), pp. 156-166, 2012.
- [40] H. Zou, F. Zheng, K. Schwan etc. "Quality-Aware Data Management for Large Scale Scientific Applications," Proc. High Performance Computing, Networking, Storage and Analysis 2012 (SC'12), pp 816-820, 2012.