Hierarchical Collective I/O Scheduling for High-Performance Computing

Jialin Liu

Department of Computer Science
Texas Tech University
Lubbock, Texas, USA
Email: jaln.liu@ttu.edu

Yu Zhuang

Department of Computer Science
Texas Tech University
Lubbock, Texas, USA
Email: yu.zhuang@ttu.edu

Yong Chen

Department of Computer Science
Texas Tech University
Lubbock, Texas, USA
Email: yong.chen@ttu.edu

Abstract—The non-contiguous access pattern of many scientific applications results in a large number of I/O requests, which can seriously limit the data-access performance. Collective I/O has been widely used to address this issue. However, the performance of collective I/O could be dramatically degraded in today's high-performance computing systems due to the increasing shuffle cost caused by highly concurrent data accesses. This situation tends to be even worse as many applications become more and more data intensive. Previous research has primarily focused on optimizing I/O access cost in collective I/O but largely ignored the shuffle cost involved. Previous works assume that the lowest average response time leads to the best QoS and performance, while that is not always true for collective requests when considering the additional shuffle cost. In this study, we propose a new hierarchical I/O scheduling (HIO) algorithm to address the increasing shuffle cost in collective I/O. The fundamental idea is to schedule applications' I/O requests based on a shuffle cost analysis to achieve the optimal overall performance, instead of achieving optimal I/O accesses only. The algorithm is currently evaluated with the MPICH3 and PVFS2. Both theoretical analysis and experimental tests show that the proposed hierarchical I/O scheduling has a potential in addressing the degraded performance issue of collective I/O with highly concurrent accesses.

Keywords—collective I/O; scheduling; high-performance computing; big data; data intensive computing

I. INTRODUCTION

The volume of data collected from instruments and simulations for scientific discovery and innovations keeps increasing rapidly. For example, the Global Cloud Resolving Model (GCRM) project [1], part of DOE's Scientific Discovery through Advanced Computing (SciDAC) program, is built on a geodesic grid that consists of more than 100 million hexagonal columns with 128 levels per column. These 128 levels will cover a layer of 50 kilometers of atmosphere up from the surface of the earth. For each of these grid cells, scientists need to store, analyze, and predict parameters like the wind speed, temperature, pressure, etc. Most of these global atmospheric models process data in a 100-kilometer scale (the distance on the ground); however, scientists desire higher resolution and finer granularity, which can lead to significant larger sizes of datasets. Table I shows the data requirements of representative scientific applications run at Argonne Leadership Computing Facility

(ALCF) through the DOE's INCITE program [30]. The data volume processed online by many applications has exceeded TBs or even tens of TBs; the off-line data is near PBs of scale.

Table I: Data Requirements of Representative INCITE Applications at ALCF [30]

Project	On-Line Data	Off-Line Data
FLASH: Buoyancy-Driven Turbulent Nuclear Burning	75TB	300TB
Reactor Core Hydrodynamics	2TB	5TB
Computational Nuclear Structure	4TB	40TB
Computational Protein Structure	1TB	2TB
Performance Evaluation and Analysis	1TB	1TB
Climate Science	10TB	345TB
Parkinson's Disease	2.5TB	50TB
Plasma Microturbulence	2TB	10TB
Lattice QCD	1TB	44TB
Thermal Striping in Sodium Cooled Reactors	4TB	8TB

During the retrieval and analysis of the large volume of datasets on high-performance computing (HPC) systems, scientific applications generate huge amounts of non-contiguous requests [26, 34], e.g., accessing the 2-D planes in a 4-D climate dataset. Those non-contiguous requests can be considerably optimized by performing a two-phase collective I/O [12]. However, the performance of the collective I/O could be dramatically degraded when solving big data problems on a highly-concurrent HPC system [13, 28]. A critical reason is that the increasing shuffle cost of collective requests can dominate the performance. This increasing shuffle cost is due to the high concurrency caused by intensive data movement and concurrent applications in today's HPC system. The shuffle phase is the second phase of a two-phase collective I/O. A collective I/O will not finish until the shuffle phase is done. Previous research has primarily focused on the optimization of the other phase, the I/O phase, of a collective I/O for data-intensive applications. In this study, instead of only considering the service time during the I/O phase, we argue that a better scheduling algorithm in collective I/O should also consider the requests' shuffle costs on compute nodes. An aggregator who has the longest shuffle time can dominate an application's overall

performance, due to the reason that the slowest aggregator actually determines the overall performance of a collective I/O. In this research, we propose a new *hierarchical I/O (HIO)* scheduling to address this issue. The basic idea is, by saturating the aggregators' 'acceptable delay', the algorithm schedules each application's slowest aggregator earlier. The proposed algorithm is named as *hierarchical I/O* scheduling, because the predicted shuffle cost is considered at the MPI-IO layer on compute nodes and the server-side file system layer. Both layers leverage the shuffle cost analysis to perform an improved scheduling for collective I/O. The current analyses and experimental tests have confirmed the improvements over existing approaches. The proposed hierarchical I/O scheduling has a potential in addressing the degraded performance issue of collective I/O with highly concurrent accesses.

The contribution of this research is three-fold. First, we propose an idea of scheduling collective I/O requests with considering the shuffle cost. Second, we have derived functions to calculate and predict the shuffle cost. Third, we have carried out theoretical analyses and experimental tests to verify the efficiency of the proposed hierarchical I/O (HIO) scheduling. The results have confirmed that the HIO approach is promising in improving data accesses for high-performance computing. This work is an extension of our previous work[25]. The major difference is we generalize the HIO idea to a broader view, in which not only collective read but also write operation scheduling is designed, analyzed, and evaluated. The second difference is we add more evaluation results to demonstrate the potential of HIO. The third improvement is the Time Window concept. In the previous work, we only discussed how to use HIO to perform the scheduling on the queuing I/O requests, but we did not consider the starvation and interruption of aggregators within the same application, in other words, the aggregators from same application's different instance could also mess up with each other. We address the problem in this paper by applying a flexible time window concept. Besides, the shuffle cost prediction and HIO implementation are also extended a lot.

The rest of this paper is organized as follows. Section II reviews collective I/O and motivates this study by analyzing a typical example of interrupted collective read. Section III introduces the HIO scheduling algorithm. Section IV presents the theoretical analysis of the HIO scheduling. Section V discusses the implementation. The experimental results are discussed in Section VI. Section VII discusses related work and compares them with this study. Section VIII summarizes this study and discusses future work.

II. BACKGROUND

A. Collective I/O

MPI is the dominant parallel programming model on all large-scale parallel machines, such as Cray XT5/XK6/XK7, IBM Blue Gene/P, IBM Blue Gene/Q supercomputers. We

briefly review its I/O interface, MPI-IO, in this subsection. We will also discuss the MPI-IO's common implementation, the most important optimization, collective I/O, and its nonblocking version.

MPI-IO is a subset of the MPI-2/MPI-3 specification [14]. It defines an I/O access interface for parallel I/O. The primary motivation for MPI-IO specification came from the observation that parallel I/O optimizations require two basic abstractions: the ability to define a set of processes, i.e., MPI communicators, and the ability to define complex data access patterns, i.e., MPI datatypes. By equipping the two abilities, the MPI-IO is designed as an interface that supports many parallel I/O operations and optimizations. The implementation of MPI-IO is usually a middleware connecting parallel applications and underlying various parallel file systems, providing the code-level portability across many different machine architectures and operating systems. ROMIO is a popular MPI-IO implementation [33]. It provides an abstract-device interface called ADIO for implementing the portable parallel I/O API. It performs various optimizations, including collective I/O and data sieving, for common access pattern of parallel applications.

Collective I/O is one of the most important I/O access optimizations. In collective I/O, multiple processes cooperate with each other to carry out large aggregated I/O requests, instead of performing many non-contiguous and small I/Os independently. The motivation of collective I/O is several-fold. First, collective I/O can filter overlapping and redundant requests from multiple processes. Second, for many parallel applications, even though each process may access several noncontiguous portions of a file, the requests of multiple processes are often interleaved and may instead result in the access of one large contiguous portion of a file. Third, the collective I/O can reduce the number of system calls by combining small and noncontiguous requests into large and contiguous ones.

A widely-used implementation of collective I/O is the two-phase I/O protocol [33]. This strategy serves the I/O requests using an I/O phase and a data exchange phase. As shown in the Figure 1, in the case of two phase collective read, the first phase consists of a certain number of processes that are assigned as aggregators to access large contiguous data. In the second phase, those aggregators shuffle the data among all processes to the desired destination.

Collective I/O is a technique to optimize one application's I/O, such optimization does not consider the interruption from other application, or other processes. while in current and future extreme scale HPC system, the highly concurrency is not neglectable. As the interruption increasing, the service order of the aggregator is random and one application's aggregators can have different waiting time (they are supposed to be served at the same time). To improve the average execution time and improve the performance, there should be an optimized scheduling methods.

Figure 1: Two Phase Collective I/O

Our scheduling idea for addressing the interruption issue originates from the nature of the two phase collective I/O itself. By scheduling the aggregators from multiple concurrent applications, we achieved lower average execution time.

B. Server I/O Scheduling

III. MOTIVATION

Collective I/O plays a critical role in cooperating processes to generate aggregated I/O requests, instead of performing non-contiguous small I/Os independently [10, 33]. As we discussed in the background section, a widely-used implementation of collective I/O is the two-phase I/O protocol[33]. For collective reads, in the first phase, a certain number of processes are assigned as aggregators to access large contiguous data; in the second phase, those aggregators shuffle the data among all processes to the desired destination. There is no synchronization during the shuffle phase, which means, as long as one aggregator gets its data, it will redistribute the data among processes immediately without waiting for other aggregators.

Figure 2: Interrupted Collective Read. (p_x is from other applications)

An observation is that, on today's HPC system with highly concurrent accesses, the service order of the aggregators on storage nodes can have an impact on the application's overall performance. The example in Figure 2 shows a two-phase collective read operation, which is interrupted by processes from other concurrent applications due to highly concurrent accesses. In Figure 2, five processes (p_0 - p_4) (on the same compute node for simplicity) from one MPI application are accessing the data striped across three storage nodes.

During the first phase, the I/O aggregators, p_0 , p_2 and p_4 , are assigned with evenly partitioned file domains. In this case, we can predict that only two aggregators (i.e., p_0 and p_2) will have to redistribute the data among other processes (i.e., p_1 and p_3) in the shuffle phase. The reason why p_4 does not need to participate in the shuffle phase is that p_4 's requests are only accessed by p_4 itself. From Figure 2, we can also find that the service order of each aggregator is different on the storage nodes, which means three aggregators of the same application are not serviced at the same time. For example, assuming other processes have the same service time, then a possible service order for these three I/O aggregators is p_4 , p_0 and p_2 . Such a service order can have variants and can have an impact. We compare two of them in Figure 3.

Figure 3: Two Different Service Orders.

In Figure 3, we analyze the cost for two different service orders. We assume that the service time and the shuffle cost is equal for the same amount of data movement and each process has service time $6t$, while the total cost is calculated as the sum of the read cost and the shuffle cost. In Figure 3(a), the aggregator p_4 is serviced first. After $6t$, p_0 receives the service, and then p_2 . During the shuffle phase, only p_0 and p_2 need to redistribute the data with other processes. Therefore, this application's total execution time is $3 \times 6t + 6t = 24t$. In Figure 3(b), p_0 is serviced first. We find that the execution time is reduced to $18t$. The performance gain comes from scheduling the 'slowest aggregator' first. The 'slowest aggregator' in this study refers the aggregator who takes the longest time to redistribute the data in the shuffle phase. Another observation from Figure 3(b) is that the service order of p_4 will not have impact on the total cost, which means even if p_4 comes first on node2 in some case, we can still service it later. In other words, we can delay p_4 at most $12t$, this delay time is acceptable (no performance degradation will be caused). This example only shows that scheduling aggregators properly can improve the performance for one application, whereas for multiple concurrent applications, how to achieve the average lowest execution time is a challenge. Besides, the shuffle cost of different aggregators varies. How to predict the shuffle cost and pass it to the server is a challenge too.

In the following sections, we introduce a hierarchical I/O scheduling (HIO) algorithm to address these issues. To the best of our knowledge, the hierarchical I/O scheduling is the first method that considers the increasing shuffle cost in collective I/O due to highly concurrency accesses

IV. HIERARCHICAL I/O SCHEDULING

From the previous analysis, we can see that by scheduling slower aggregators earlier, the execution time (access cost and shuffle cost) can be reduced. In the case of with concurrent applications, however, the goal of the optimal scheduling should be achieving the lowest ‘average’ execution time. From the observation in Figure 3(b), we know that application’s aggregators may have ‘acceptable delay’ time. If such time is well utilized to service other applications, we can potentially achieve a win-win situation. In the following subsections, we first discuss how to predict the shuffle cost, and then formally introduce the concept of ‘acceptable delay’. We also apply the time window concept [32] to divide the long I/O queue on each node into sub-sequences. Finally, we present the proposed hierarchical I/O scheduling algorithm.

A. Analysis and Prediction of Shuffle Cost

In order to analyze the shuffle cost, we need to know how aggregators are assigned and how will they communicate with other desired processes. Most previous works assumed that only one aggregator is assigned in one node [5], which is also a default configuration in MPICH. The communication cost thus includes inter-node and intra-node cost, as shown in Figure 4. The amount of data redistributed impacts the cost too. Therefore, the shuffle costs are mainly determined by exchanging data and the number and position of desired processes.

Figure 4: Example of Process Assignment in Three Node Multi-core System. The arrows show the inter and intra communication in the shuffle phase.

In Figure 4, we illustrate an example of different communication patterns in the shuffle phase. In this example, there are totally three compute nodes, where each node is assigned with different number of processes, with only one process acting as the aggregator. During the shuffle phase, the aggregator either sends the data to processes on other nodes, which results in inter-node communication, or redistributes

the data within the same node, which results in intra-node communication. As a consequence, each aggregator will have different shuffle cost.

Figure 5: Timing aggregator and non-aggregator's cost in the shuffle phase

To predict the shuffle cost of each aggregator, we then conducted an initial experiment by timing one aggregator and its coresponding processes, as shown in Figure 5. This initial test is performed in two nodes, each node has 12 cores. We did a collective I/O using 16 processes, in which 12 processes including one aggregator are from the same node, while the other four processes are from the second node (This does not need to manually specify, just by simply launching 16 processes). The data amount that the aggregator will send to each non-aggregator is same. We modify the MPI-IO to timing the sending/receiving and other steps. The first bar in the Figure 5 shows the aggregator's sending time, which is the time this aggregator takes to send the data to the collective buffer. The lower dashed portion in other bars (2-16) shows the waiting time of non-aggregators and the higher portion refers to their receiving time. We can see that the data exchange time between aggregator and non-aggregators is various, due to which, we can argue that the shuffle cost is determined by the maximum data exchange time (in this case, i.e., 2.3ms). Based on the above theoretical analysis and the initial experiment, we can derive the following equation to predict the shuffle cost:

$$T = \max(\max(\frac{m_{ai}}{B_a}), \max(\frac{m_{ej}}{B_e})) + \gamma$$

$$= \max(\frac{m_{ai}}{B_a}, \frac{m_{ej}}{B_e}) + \gamma \quad (1)$$

where T is the total shuffle cost of one aggregator; m_{ai} is the i^{th} intra message size (MByte) and $0 < i < A$, where A is the total number of intra communication; m_{ej} is the j^{th} inter communication message size (MByte), $0 < j < E$, where E is the total number of inter-node communication; B_a is the saturated throughput of intra-node communication (MB/sec); B_e is the saturated throughput of inter-node communication

of a given cluster system (MB/sec); and γ is the latency. the difference with ours is that their approach calculates the shuffle cost of each node, while ours calculate the cost in terms of each aggregator.

In order to distinguish the intra-communication and inter-communication at the runtime, we need to know the Hydra's process-core binding strategy. Hydra is a process management system compiled into the MPICH2 as a default process manager. Without any user-defined mapping, we assume the basic default allocation strategy is used, i.e., round-robin mechanism, using the OS specified processor IDs. Whether the communication will be intra or inter can be determined with the following equation:

$$Comm\ is\ \begin{cases} intra & \text{if } a_id \% n_c = p_id \% n_c \\ inter & \text{else} \end{cases} \quad (2)$$

where, Comm is short for communication, a_id is the rank of the aggregator, p_id is the rank of non-aggregator processes, and n_c is number of cores per node.

B. Acceptable Delay

As we have analyzed in the previous section, aggregators with lower shuffle costs can be scheduled later, whereas slow aggregators who have higher shuffle costs are better to be serviced first. We introduce the *Acceptable Delay* (AD) in this study to support the hierarchical I/O scheduling. An aggregator's AD refers to the maximum acceptable time it can be delayed. The AD is defined as follows:

$$AD_i = \max\{T_0, T_1, T_2, \dots, T_n\} - T_i \quad (3)$$

where AD_i is i^{th} aggregator's acceptable delay, and T_i is i^{th} aggregator's shuffle cost. Usually, I/O requests from the same application are better to be serviced at the same time in order to achieve lower average response time. However, due to aggregators' various ADs, it is not necessary to do that, which means we can utilize every aggregator's AD to better schedule I/O requests, by saturating one aggregator's AD and servicing other applications first. We also define a Relative Acceptable Delay (RAD) as the rank in the ascending order of AD.

C. Time Window

In previous work, a time window is used to avoid starvation and to maintain fairness [32]. With the time window concept, all I/O requests waiting in a queue are divided equally by a predefined time interval. Each divided window consists of a sequence of I/O requests. In the same window, I/O requests are ordered by the value of 'Application ID'; whereas in different windows, requests in an earlier 'Time Window' will be serviced prior to those in a later one, to avoid starvation.

As shown in Figure 6, we utilize the time window to organize the I/O queue too but applied for the new hierarchical I/O scheduling algorithm. The motivation here is

Figure 6: Divided Requests Queues with Time Window.

different compared to the existing work. In our design, we argue that the earlier request from one application should always be serviced earlier than the later one from the same application. The earlier I/O is from an aggregator of the earlier collective access. After this earlier collective access finishes, the application needs a synchronization before issuing another collective read/write. Scheduling the later aggregator earlier will cause delay. Therefore, instead of a fixed-width time window, we set a flexible time window on each server. On each storage node, the original queue will be divided into several windows. In each window, there are no more than two aggregators from the same application. In other words, the aggregators within a window are all from different applications.

D. HIO Algorithm

The main idea of the hierarchical I/O scheduling algorithm is to utilize the aggregator's 'acceptable delay' to minimize the shuffle cost. Because aggregators from different applications have various ADs, we can not directly compare the AD from different applications. The algorithm first generates an initial order, and then tunes the order by comparing the aggregator's AD and read cost. If one aggregator's AD is larger than its successor's read cost, then the order of the two requests can be exchanged. The algorithm is described in Algorithm 1-HIO scheduling algorithm. Before the HIO scheduling algorithm performed, the I/O queue is divided into fixed windows (e.g., 10) on each node. If there are more than one aggregator from the same application, the window width is reduced as discussed in the previous subsection. The outer loop (i to $n - 1$) in the algorithm is carried out in parallel because the scheduling is performed on each node separately. The actual scheduling starts from the second loop (j to $m - 2$). Each request's AD is compared with its successor's read cost. If $agg_j.ad > agg_{j+1}.read$, then exchange the order of agg_j and agg_{j+1} . At the same time, the AD is updated as: $agg_j.ad = agg_j.ad - agg_{j+1}.read$, $agg_{j+1}.ad = agg_{j+1}.ad + agg_j.read$.

input :

n: number of storage nodes;
m: number of applications;
threshold: 0.2;
agg[i][j].ad: the acceptable delay of j_{th}
aggregator on i_{th} node;
agg[i][j].read: the read cost of j_{th}
aggregator on i_{th} node
agg[i][j].rad: the relative ad of j_{th}
aggregator on i_{th} node

output: Optimal service order on each node

```

for  $i \leftarrow 0$  to  $n - 1$  do
  ratio = sum(agg[i].shuffle)/sum(agg[i].read);
  if ratio > threshold then
    | qsort agg[i] by rad;
  end
  else
    | qsort agg[i] by app_id;
  end
  for  $j \leftarrow 0$  to  $m - 2$  do
    for  $k \leftarrow j + 1$  to  $m - 1$  do
      if agg[i][j].ad > agg[i][k].read then
        temp = agg[i][j];
        agg[i][j] = agg[i][k];
        agg[i][k] = temp;
        agg[i][j].ad += agg[i][k].read;
        agg[i][k].ad -= agg[i][j].read;
      end
      else
        j++;
        break;
      end
    end
  end
end

```

Algorithm 1: HIO Scheduling Algorithm

The initial order is generated by sorting the RAD (if $shuffle/read > threshold$), through which each application's slowest aggregator is scheduled earlier. The reason why a threshold is set that if the shuffle cost is too small compare to the read cost, the algorithm just sorts I/Os by application ID. The detailed reason is discussed in Section IV. The algorithm then tunes the initial order by saturating each aggregator's AD. These two steps make sure that the slower aggregator moves ahead, and the faster aggregator moves back. The tuning counts the read cost in order to balance the service order among all applications.

V. THEORETICAL ANALYSIS

The HIO scheduling can reduce the service time and shuffle cost. In this section, we analyze the cost reduction through an analytical model for collective I/O and compare

against one latest Server I/O scheduling [32].

Assuming the number of concurrent applications is m , and the number of storage nodes is n . On each node, assuming every application has a request, then there will be m aggregators on each node. Suppose each request needs time t to finish the service on the server side and s_{ij} to finish the shuffle phase (s_{ij} is the j_{th} aggregator's shuffle cost of the i_{th} application). The longest finish time of requests on all nodes determines the application's completion time on the server side, which could be $\{t, 2t, 3t, \dots, mt\}$. The application's total cost is the sum of the completion time on the server side and the maximum shuffle cost on the client side. Without any scheduling optimization, applications' server-side completion time has the same distribution, i.e., the density is $g(x)$, while the probability distribution function is $G(x) = (x/m)^n$. Therefore, the completion time of each application can be derived as shown in Equation 4:

$$\begin{aligned}
T_i &= \text{Service time} + \text{Shuffle cost} \\
&= E(\max(Tc_i)) + \max(s_{ij}) \\
&= \left(\sum_{x=1}^m xg(x) \right) t + \max(s_{ij}) \\
&= \left(\sum_{x=1}^m x(G(x) - G(x-1)) \right) t + \max(s_{ij}) \quad (4) \\
&= \left(\sum_{x=1}^m x \left(\left(\frac{x}{m} \right)^n - \left(\frac{x-1}{m} \right)^n \right) \right) t + \max(s_{ij}) \\
&= mt + \max(s_{ij}) - \frac{t}{m^n} \sum_{x=1}^{m-1} x^n
\end{aligned}$$

in which, Tc_i is the i_{th} application's completion time on one node.

With the Server I/O scheduling, in which the same applications' requests are serviced at the same time on all nodes, the service time for those applications are fixed: $t, 2t, 3t, \dots, mt$. The average execution time is:

$$\begin{aligned}
T_i &= \frac{1}{m} \left(\sum_{x=1}^m xt + m(\max(s_{ij})) \right) \\
&= \frac{m+1}{2} t + \max(s_{ij}) \quad (5)
\end{aligned}$$

For the potential of the HIO scheduling, we analyze the best case and the worst case separately. The best case requires two conditions: first, each application's slowest aggregator comes to different node; second, the slowest aggregator dominates the application's execution time, which can be described as $\max(s_{ij}) - \min(s_{ij}) > (m-1)t$. With the HIO scheduling, the slowest aggregator is serviced first on each node and determines each application's execution

time. Therefore, we have the average execution time:

$$T_i = \frac{1}{m}(mt + m(\max(s_{ij}))) \quad (6)$$

$$= t + \max(s_{ij})$$

For the worst case, either the first condition or the second condition is not satisfied. If the first condition is not met, it indicates that applications' slowest aggregators arrive at the same node. Thus, the average execution time is:

$$T_i = \frac{1}{m}((t + 2t + 3t + \dots + mt) + m(\max(s_{ij}))) \quad (7)$$

$$= \frac{m+1}{2}t + \max(s_{ij})$$

If the second condition is not met, the slowest aggregator's shuffle cost is close to zero. With the HIO scheduling, the initial order will be sorted by application id, which means that the same application will be serviced at the same time. Then we have the average execution time same with Equation 5. In another word, the worst case of the HIO scheduling at least has the same performance as that of the Server I/O scheduling.

Comparing Equation 4 and Equation 5, the Server I/O scheduling can achieve an average execution time reduction as the following:

$$T_{reduction} = mt + \max(s_{ij}) - \frac{t}{m^n} \sum_{x=1}^{m-1} x^n$$

$$- \frac{m+1}{2}t - \max(s_{ij}) \quad (8)$$

$$= \frac{m-1}{2}t - \frac{t}{m^n} \sum_{x=1}^{m-1} x^n$$

$$= \frac{m-1}{2}t \quad (n \rightarrow \infty)$$

The best case of the HIO scheduling can further reduce the execution time of Equation 5 by:

$$T_{reduction}^b = \frac{m+1}{2}t + \max(s_{ij})$$

$$- t - \max(s_{ij}) \quad (9)$$

$$= \frac{m-1}{2}t$$

The theoretical analysis and this comparison show that the HIO scheduling achieves better scheduling performance, especially when the shuffle cost keeps increasing due to highly concurrent accesses from large-scale HPC systems and/or big data retrieval and analysis problems.

VI. DISCUSSION OF HIO FOR WRITE OPERATION

The HIO was originally designed for the collective read operation. Since the read operation's two phase procedure is different with the write operation. Our scheduling only works for the read operation, which is I/O phase first and shuffle phase second. The reason is that the I/O phase of

read on server side is not the last step for the application, and the later shuffle phase on client side, if well utilized by the HIO, the overall performance can be improved. However, for collective write operation, the data are first shuffled on the client side, then after each aggregator getting its file view and all the data within that view, the aggregator will be sent to the storage nodes to do the I/O. This procedure is totally opposite of the read operation. The HIO idea seems to fail in this case. Fortunately, by rethinking the HIO idea, we found that it is not difficult to apply the HIO for write operations without modification. Same with collective read, the aggregator's shuffle phase for collective write also has various cost. Such various cost will lead to different arriving order on the storage server.

Suppose we have a 'third' phase for collective write, which means when the aggregators are returned to the compute nodes, they will also do the "shuffle" similar with the collective read's second phase. But assuming the cost of the third phase equals to zero, then it is not difficult to find that the collective write is just one 'worst' case of collective read for HIO. In section V, we have discussed how our HIO addresses the worst cases, in which the shuffle cost close to zero and the we have Equation 5.

VII. IMPLEMENTATION

The aggregators' shuffle cost and AD are calculated at the MPI-IO layer. Our evaluation was carried out on the ROMIO that is included in MPICH2-1.4. It provides a high-performance and portable implementation of MPI-IO including collective I/O. The MPICH2-1.4 and ROMIO provide a PVFS2 ADIO device. We modified this driver to integrate the shuffle cost analysis and pass it to the PVFS server side scheduler as a hint. When an application calls the collective read function ADIOI_Read_and_exch in ad_read_coll.c under the src/mpi/romio/adio/common, the shuffle cost is calculated after the aggregators are allocated, i.e., ADIOI_Calc_file_domains. The message size m is calculated with ADIOI_Calc_my_req and ADIOI_Calc_others_req. The calculated shuffle cost is stored into a variable of PVFS-hint type. The hint is passed to file servers along with I/O requests.

On the PVFS server side, in the request scheduling function PINT_req_sched_post(), we implemented the HIO algorithm. The original function only enqueues the coming requests into the tail of the queue, while the HIO algorithm first divides the waiting queue into several sequences, and performs the scheduling within each sub-queue following the scheduling algorithm discussed in Section III.

VIII. EXPERIMENTS AND ANALYSES

A. Experimental Setup

We have conducted tests on a 16-node Linux testbed. This cluster is composed of one PowerEdge R515 rack server node and 15 PowerEdge R415 nodes, with a total

Figure 7: Implementation

of 32 processors and 128 cores. Nodes are fully connected via a PowerConnect 2848 network switch. The PowerEdge R515 server node has dual quad-core 2.6GHz AMD Opteron 4130 processors, 8GB memory, and a RAID-5 disk array with 3TB storage capacity composed of 7200 RPM Near-Line SAS drives. Each PowerEdge R415 node has dual quad-core 2.6GHz AMD Opteron 4130 processors, 4GB memory and a 500GB 7200RPM Near-Line SAS hard drive. We conducted experiments with the mpi-io-test parallel I/O benchmark [2]. The proposed hierarchical I/O scheduling algorithm was compared with other scheduling strategies through tests. We have also evaluated the HIO scheduling algorithm with a real climate science application. The HIO scheduling algorithm is evaluated and compared with two other scheduling algorithms, Server I/O scheduling (denoted as SIO) [32] and the normal collective I/O (denoted as NIO).

B. Results and Analyses

In the first test, we run multiple instances of MPI-IO-TEST simultaneously. We conducted the experiments by specifying the number of aggregator as 6 and the number of processes as 50 for each application. The I/O request size was set to a fixed value, 16 MBs. We run with 6, 12, 24, and 48 processes simultaneously. Six storage nodes were deployed. The results are plotted in Figure 8. From the figure, we can observe that the HIO scheduling outperformed other scheduling. The total execution time was decreased by up to 34.1% compared with NIO and by up to 15.2% compared with SIO. Furthermore, when the number of concurrent applications increased, the performance gain was even better.

We have conducted experiments with varying the request size too. As reported in Figure 9, the I/O request size was set as 64KB, 1MB, 5MB, and 10MB respectively. The number of concurrent applications was set as six, and the number of aggregators was configured as six too. During this test, we compared the ratio of shuffle cost against the total cost. It was found that the ratio increased from 0.7% to 5.6%, as the request size increased. This fact matches with our

Figure 8: Average execution time with concurrent applications

observation that the shuffle cost considerably increases when applications become more and more data intensive.

When the requests size increased, the performance gain of using HIO scheduling was increased too, from 6.8% to 18.3% in terms of the execution time reduction rate. This result also matches with our theoretical analysis discussed in Section IV.

Figure 9: Speedup of HIO and SIO with different request size

We have also evaluated the impact of the number of storage nodes and report the results in Figure 10. In this test, there are 6 applications running simultaneously, and the number of aggregators in each application was set the same as the number of storage nodes, in order to have each application access all storage nodes. The request size was set as 15MBs, and the aggregator's request size is equal. The number of storage nodes was varied as 2, 4, 6, 8, and 16.

We observe that, from the Figure 10, the normal collective I/O did not scale well with the increasing size of the system. While both HIO and SIO achieved better scalability, we also find that the HIO performed and scaled better than SIO. The advantage of the HIO is due to the reduced

Figure 10: Average execution time with different number of storage nodes

shuffle cost. As the number of storage nodes increased, the inter-communication between aggregators and processes on different nodes also increased, which has been confirmed in a prior study too [5]. It can be projected that, as the system scale keeps increasing in the big data computing era, the shuffle cost in the two-phase collective I/O will become a critical issue. The proposed HIO scheduling in this study is essentially for addressing this issue and is likely to be promising at the exascale/extremescale of HPC system.

Figure 11: Average execution time with different number of aggregators

We also evaluated the HIO with various number of aggregators. In Figure 11, we test the HIO, SIO and NIO with 5, 10 and 15 aggregators separately. As the number of aggregator increasing, the interruption and the shuffle cost will also increase. We can find that the NIO shows a linear increasing trend. Both the HIO and SIO reduce the average execution time. The performance gain of SIO comes from the reduction of interruption of aggregators on storage nodes. While the HIO achieves more by reducing the shuffle cost.

As we have discussed in section VI, the HIO was originally designed for collective read. For collective write, since the shuffle phase is already done before any scheduling, so

Figure 12: Collective Write with HIO

the HIO can not utilize the various acceptable delay any more, therefore, we can see from Figure 12, the HIO and SIO do not distinguish a lot. Both of them achieve an average speedup about 12%.

Figure 13: FASM with HIO

We have evaluated the HIO scheduling with a real climate science application and datasets from the Bjerknes Center for Climate Research as well [24]. This set of tests was specifically for understanding the benefits of the HIO scheduling for a special access pattern, accessing 2D planes in scientific datasets. In scientific computing, scientists are interested in understanding the phenomenon behind the data by performing subsets queries [24]. Those subsets queries usually happen in the 2D planes, e.g., parameters along time dimension and level dimension in a climate science data. The datasets can range between GBs and TBs. Previous studies have shown the poor scalability of collective I/O due to high concurrency and I/O interruption. The proposed HIO algorithm addresses this issue by better scheduling concurrent I/Os. The total dataset size evaluated in this series of tests is more than 12 GBs. We run multiple 2D subsets queries concurrently using the FASM system [24] and performed the HIO scheduling. We run 20 queries for each dataset. A sample query statement is like "select

temperature from dataset where $10 < \text{temperature} < 31$ ". These queries were generated randomly and followed a global normal distribution. The performance gain with the HIO scheduling compared to the conventional collective I/O is shown in Figure 13. It can be observed that the HIO scheduling improved the average query response time clearly and by up to 59.8%.

All these tests have well confirmed that the proposed hierarchical I/O scheduling in this study can improve the performance of collective I/O given highly concurrent and interrupted accesses. It holds a promise for big data problems and scientific applications on large-scale HPC systems.

IX. RELATED WORK

A. I/O Scheduling

Parallel I/O scheduling has been widely studied by many researchers at a hope of obtaining the peak sustained I/O performance. Few of them, however, meets the current demand of data-intensive applications and big data analysis yet. Disk-directed I/O[22] and server-directed I/O[31] have been proposed to improve the bandwidth of disks and network servers respectively. There are also numerous scheduling algorithms targeting the quality of service (QoS) [15, 16, 18, 29, 36]. The proposed hierarchical I/O scheduling in this study takes one step further to optimize the scheduling of collective I/O while considering the highly concurrent I/O requests from data-intensive applications.

In [32], a server-side I/O coordination method is proposed for parallel file systems. Their idea is to coordinate file servers to serve one application at a time in order to reduce the average completion time, and in the meantime maintain the server utilization and fairness. By re-arranging I/O requests on the file servers, the requests are serviced in the same order in terms of applications on all involved nodes. However, without considering the shuffle cost in the collective I/O, it is unlikely to achieve the optimal performance. In [38], the authors proposed a scheme namely IOrchestrator to improve the spatial locality and program reuse distance by calculating the access distances and grouping the requests with small distance together. These two works seem similar but differ in the motivation. The first one is based on the observation that the requests with synchronization needs will be optimized if they are scheduled at the same time, whereas the latter one is motivated by exploring the program's spatial locality.

In [17], the authors proposed three scheduling algorithms, with considering the number of processes per file stripe and the number of accesses per process, to minimize the average response time in collective I/O. In servicing one aggregator, instead of scheduling one stripe at a time in the increasing file offset order, they propose to prioritize the file stripes based on their access degree, the number of accessing processes. Their work optimized the scheduling of stripes within an aggregator, whereas our work focuses

on the scheduling of aggregators. Besides, their work only considers the average I/O response time. The reduced I/O response time, however, does not always lead to the reduced total cost that includes the I/O response time and the shuffle cost.

B. Two Phase Collective I/O Optimization

Collective I/O has been proposed about 15 years [33], optimization of collective I/O has never been stopped. We have studied most of the work related to two phase collective I/O, from our classification, the research efforts have been focused on the following points:

- 1) Implementation[9, 33, 35], in which the collective I/O are designed implemented and advanced feature are supported;
- 2) File view partition [3, 20, 37], which are related to the two phase collective I/O's global file view partition to optimize the aggregator's I/O access;
- 3) Aggregator selection[4, 6]. Interesting ideas are about how to selection the processes as aggregators and how to define the number of aggregators;
- 4) Cache and buffering [21, 27]. Researchers find the traditional cache and buffering idea can be utilized to optimize the collective I/O, which is done on the client side;
- 5) Data compression[11] is another example that other well-studied ideas are used in collective I/O.

All these works have successfully improved the performance of two phase collective I/O, and proved that the collective I/O is promising in the parallel computing. But we can also find that the new challenges in today's large scale, big data, and power constrained era, drive the further development of collective I/O. and our HIO is proposed under this circumstance. Among the previous work in collective I/O, there are two which are most related.

The first one is in [5, 6], the authors discussed the increasing shuffle cost in today's HEC system too. Their discussions are for motivating the importance of node re-ordering for reducing the collective I/O's shuffle cost. Their work provides a method to evaluate the shuffle cost and designed algorithms to automatically assign the aggregators at the node level, whereas our work focuses on the scheduling of aggregators considering highly concurrent accesses to achieve the optimal collective I/O performance.

The second one is the LACIO idea [7], in which the author discuss the gap between logical access and physical storage. When the aggregators are assigned with different file view, the view itself is just logical space, therefore, the aggregator may not know the physical distribution of data on the parallel storage nodes. In other words, the aggregator from same application will interrupt with each other on the storage nodes and the LACIO idea well addressed the issue. Our HIO idea targets the interruption amount concurrent applications, which in another side, to improve

the performance (we did our evaluation by first removing the interruption from same application, such that we can tell the HIO really reduces the interruption among different applications). Previous work, like LACIO, focus on one application's collective I/O, while our work goes one more step to optimize the performance of concurrent case. Both of the direction, when combined, is indeed a trend for extreme scale systems in the future.

C. Data Organization and File Systems

Our work focuses on I/O scheduling, while the data is not moveable. In fact, there are bunch of work targeting the data organization and file systems. We concluded some related work in this last subsection. For example, there exist other works that address the scientific data retrieval issues by optimizing the data organization [19]. These works provide efficient mechanisms from the data level and fit the access pattern of scientific applications. Our work also improves the application's accesses, most of which are non-contiguous, but through a hierarchical scheduling. There are also works utilizing existing database techniques and compression algorithms to boost the big data analysis. For example, Fastbit implemented bitmap index in the large datasets [8]. ISABELA improved the big data query by compressing the datasets [23]. Our work focuses on collective I/O scheduling, which is beneficial and critical to big data retrieval and analysis too.

X. CONCLUSION

Collective I/O has been proven a critical technique in optimizing the non-contiguous access pattern in many scientific applications run on high-performance computing systems. It can be critical for big data retrieval and analysis too as non-contiguous access pattern also commonly exists in big data problems. The performance of collective I/O, however, could be dramatically degraded due to the increasing shuffle cost caused by highly concurrent accesses and interruptions. This problem tends to be more and more critical as many applications become highly data intensive. In this study, we propose a new hierarchical I/O scheduling for collective I/O to address these issues. This approach is the first considering the increasing shuffle cost involved in collective I/O. Through theoretical analyses and experiments, it has been confirmed that the hierarchical I/O scheduling can improve the performance of collective I/O. In the future, we will apply the similar approach for write operations. We will analyze the feasibility of implementing hierarchical I/O scheduling only at the MPI-IO layer as well. More experiments will be conducted to analyze how the shuffle cost can affect the big data analysis and further refine our algorithm. We will also try to apply similar approaches for write operations and develop different scheduling methods for different parallel file systems.

ACKNOWLEDGMENT

This research is sponsored in part by the National Science Foundation under grant CNS-1162488 and the Texas Tech University startup grant. The authors are thankful to Yanlong Yin of Illinois Institute of Technology and Wei-Keng Liao of Northwestern University for their constructive and thoughtful suggestions toward this study. We also acknowledge the High Performance Computing Center (HPCC) at Texas Tech University for providing resources that have contributed to the research results reported within this paper.

REFERENCES

- [1] The global cloud resolving model (gcrm) project. <http://kiwi.atmos.colostate.edu/gcrm/>.
- [2] Mpi-io test. <http://public.lanl.gov/jnunez/benchmarks/mpiioetest.htm>.
- [3] F. J. G. Blas, F. Isaila, D. E. Singh, and J. Carretero. View-based collective I/O for MPI-IO. In *CCGRID*, pages 409–416. IEEE Computer Society, 2008.
- [4] K. Cha and S. Maeng. An efficient I/O aggregator assignment scheme for collective I/O considering processor affinity. In *International Conference on Parallel Processing Workshops, ICPPW, Taipei, Taiwan, Sept. 13-16, 2011*, pages 380–388. IEEE, 2011.
- [5] K. Cha and S. Maeng. Reducing communication costs in collective I/O in multi-core cluster systems with non-exclusive scheduling. *The Journal of Supercomputing*, 61(3):966–996, 2012.
- [6] M. Chaarawi and E. Gabriel. Automatically selecting the number of aggregators for collective I/O operations. In *CLUSTER*, pages 428–437. IEEE, 2011.
- [7] Y. Chen, X.-H. Sun, R. Thakur, P. C. Roth, and W. D. Gropp. LACIO: A new collective I/O strategy for parallel I/O systems. In *IPDPS*, pages 794–804. IEEE, 2011.
- [8] J. Chou, M. Howison, B. Austin, K. Wu, J. Qiang, E. W. Bethel, A. Shoshani, O. Rübel, Prabhat, and R. D. Ryne. Parallel index and query for large scale data analysis. In *Conference on High Performance Computing Networking, Storage and Analysis, SC 2011, Seattle, WA, USA, November 12-18, 2011*. ACM, 2011.
- [9] P. Dickens and R. Thakur. Improving collective I/O performance using threads. In *Proceedings of the Joint International Parallel Processing Symposium and IEEE Symposium on Parallel and Distributed Processing*, pages 38–45, Apr. 1999.
- [10] P. M. Dickens and R. Thakur. Evaluation of collective I/O implementations on parallel architectures. *Journal of Parallel and Distributed Computing*, 61(8):1052–1076, Aug. 2001.
- [11] R. Filgueira, D. E. Singh, J. C. Pichel, and J. Carretero. Exploiting data compression in collective I/O

techniques. In *Proceedings of the IEEE International Conference on Cluster Computing (10th CLUSTER'08)*, pages 479–485, Tsukuba, Japan, Sept.-Oct. 2008. IEEE.

- [12] K. Gao, W. keng Liao, A. N. Choudhary, R. B. Ross, and R. Latham. Combining I/O operations for multiple array variables in parallel netCDF. In *CLUSTER*, pages 1–10. IEEE, 2009.
- [13] J. Gray, D. T. Liu, M. A. Nieto-Santisteban, A. S. Szalay, D. J. DeWitt, and G. Heber. Scientific data management in the coming decade. *SIGMOD Record*, 34(4):34–41, 2005.
- [14] W. Gropp, E. Lusk, and R. Thakur. *Using MPI-2: Advanced Features of the Message Passing Interface*. Scientific and engineering computation. MIT Press, pub-MIT:adr, 2000.
- [15] A. Gulati, I. Ahmad, and C. A. Waldspurger. PARDA: Proportional allocation of resources for distributed storage access. In *FAST*, pages 85–98. USENIX, 2009.
- [16] L. Huang, G. Peng, and T. cker Chiueh. Multi-dimensional storage virtualization. *ACM SIGMETRICS Performance Evaluation Review*, 32(1):14–24, June 2004.
- [17] C. Jin, S. Sehrish, W. keng Liao, A. N. Choudhary, and K. Schuchardt. Improving the average response time in collective I/O. In *EuroMPI*, volume 6960 of *Lecture Notes in Computer Science*, pages 71–80. Springer, 2011.
- [18] M. Karlsson, C. Karamanolis, and X. Zhu. Triage: Performance isolation and differentiation for storage systems. Technical Report HPL-2004-40, Hewlett Packard Laboratories, Sept. 2 2004.
- [19] W. Kendall, M. Glatter, J. Huang, T. Peterka, R. Latham, and R. B. Ross. Terascale data organization for discovering multivariate climatic trends. In *SC*. ACM, 2009.
- [20] W. keng Liao and A. Choudhary. Dynamically adapting file domain partitioning methods for collective I/O based on underlying parallel file system locking protocols. In *SC'08*. ACM/IEEE, Austin, TX, Nov. 2008.
- [21] W. keng Liao, K. Coloma, A. N. Choudhary, L. Ward, E. Russell, and S. Tideman. Collective caching: application-aware client-side file caching. In *HPDC*, pages 81–90. IEEE, 2005.
- [22] D. Kotz. Disk-directed I/O for MIMD multiprocessors. Technical Report PCS-TR94-226, Dartmouth College, July 1994.
- [23] S. Lakshminarasimhan, J. Jenkins, I. Arkatkar, Z. Gong, H. Kolla, S.-H. Ku, S. Ethier, J. Chen, C.-S. Chang, S. Klasky, R. Latham, R. B. Ross, and N. F. Samatova. ISABELA-QA: query-driven analytics with ISABELA-compressed extreme-scale scientific data. In *Conference on High Performance Computing Networking, Storage and Analysis, SC 2011, Seattle, WA, USA,*

November 12-18, 2011.

- [24] J. Liu and Y. Chen. Improving data analysis performance for high-performance computing with integrating statistical metadata in scientific datasets. In *HPCDB, SC'12*, 2012.
- [25] J. Liu, Y. Chen, and Y. Zhuang. Hierarchical I/O scheduling for collective I/O. In *Proc. of the 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid'13)*, 2013.
- [26] J. F. Lofstead, M. Polte, G. A. Gibson, S. Klasky, K. Schwan, R. Oldfield, M. Wolf, and Q. Liu. Six degrees of scientific data: reading patterns for extreme scale science IO. In *HPDC*, pages 49–60. ACM, 2011.
- [27] X. Ma, M. Winslett, J. Lee, and S. Yu. Improving MPI IO output performance with active buffering plus threads. In *Proceedings of the International Parallel and Distributed Processing Symposium*. IEEE Computer Society Press, Apr. 2003.
- [28] J. No, R. Thakur, and A. Choudhary. Integrating parallel file I/O and database support for high-performance scientific data management. In *Proceedings of SC2000: High Performance Networking and Computing*, Dallas, TX, Nov. 2000. IEEE Computer Society Press.
- [29] A. Povzner, D. Sawyer, and S. A. Brandt. Horizon: efficient deadline-driven disk I/O management for distributed storage systems. In *HPDC*, pages 1–12. ACM, 2010.
- [30] R. Ross, R. Latham, M. Unangst, and B. Welch. Parallel I/O in practice, tutorial notes. In *SC'08*. ACM/IEEE, Austin, TX, Nov. 2008.
- [31] K. E. Seamons, Y. Chen, P. Jones, J. Jozwiak, and M. Winslett. Server-directed collective I/O in Panda. In *Proceedings of SC'95*, San Diego, CA, Dec. 1995. IEEE Computer Society Press.
- [32] H. Song, Y. Yin, X.-H. Sun, R. Thakur, and S. Lang. Server-side I/O coordination for parallel file systems. In *SC*. ACM, 2011.
- [33] R. Thakur, W. Gropp, and E. Lusk. Data sieving and collective I/O in ROMIO. In *Proc. of the 7th Symposium on the Frontiers of Massively Parallel Computation*, pages 182–189. IEEE, Feb. 1999.
- [34] Y. Tian, S. Klasky, H. Abbasi, J. F. Lofstead, R. W. Grout, N. Podhorszki, Q. Liu, Y. Wang, and W. Yu. EDO: Improving read performance for scientific applications through elastic data organization. In *CLUSTER*, pages 93–102. IEEE, 2011.
- [35] V. Venkatesan, M. Chaarawi, E. Gabriel, and T. Hoefler. Design and evaluation of nonblocking collective I/O operations. In *18th European MPI Users' Group Meeting, EuroMPI 2011, Santorini, Greece, September 18-21*, volume 6960, pages 90–98. Springer, 2011.
- [36] Y. Wang and A. Merchant. Proportional service allocation in distributed storage systems. Technical Report HPL-2006-184, Hewlett Packard Laboratories, Feb. 18

2007.

- [37] W. Yu and J. S. Vetter. Parcoll: Partitioned collective I/O on the cray XT. In *ICPP*, pages 562–569. IEEE Computer Society, 2008.
- [38] X. Zhang, K. Davis, and S. Jiang. IOrchestrator: Improving the performance of multi-node I/O systems via inter-server coordination. In *SC'10*. ACM/IEEE, New Orleans, LA, USA, Nov. 2010.

*Graphical Abstract (for review)

(a) service order: p_4 , p_0 , p_2

(b) service order: p_0 , p_2 , p_4