

ACURDION: An Adaptive Clustering-based Algorithm for Tracing Large-scale MPI Applications

Amir Bahmani
abahman@ncsu.edu

Frank Mueller
mueller@ncsu.edu

Abstract—Communication traces help developers of high-performance computing (HPC) applications understand and improve their codes. When run on large-scale HPC facilities, the scalability of tracing tools becomes a challenge. To address this problem, traces can be clustered into groups of processes that exhibit similar behavior. Instead of collecting traces information of each individual node, it then suffices to collect a trace of a small set of representative nodes, namely one per cluster. However, clustering algorithms themselves need to have low overhead, be scalable, and adapt to application characteristics. We devised an adaptive clustering algorithm for large-scale applications called ACURDION that traces the MPI communication of code with $O(\log P)$ time complexity where P is the number of processes. First, ACURDION identifies the parameters that differ across processes by using a logarithmic algorithm called Adaptive Signature Building. Second, it clusters the processes based on those parameters. Experiments show that collecting traces of just nine nodes/clusters suffices to capture the communication behavior of all nodes while retaining sufficient accuracy of trace events and parameters. In summary, ACURDION improves trace scalability and automation over prior approaches.

I. INTRODUCTION

The increasing size of HPC systems often requires applications to be carefully designed for scalability. One of the key challenges is to utilize communication efficiently. Programmers generally understand the semantics of MPI communication routines, but they may not know trade-offs and limitations in a concrete implementation of MPI, particularly when a problem surfaces only at a larger scale but not in constrained testing with small inputs. In such scenarios, communication traces often provide the insight to detect inefficiencies and help in problem tuning [21], [4]. Traces are also utilized to drive HPC simulations to determine the effect of interconnect changes for future procurements [17], [25], [14], [19], [26].

Today’s tracing tools either obtain lossless trace information at the price of limited scalability (e.g., Vampir [5], Tau [16], Intel’s toolsets [6], and Scalasca [8]) or preserve only aggregated statistical trace information to conserve the size of trace files (e.g., mpiP [20]). In general, the scalability of tracing tools becomes a challenge when an application is run on large-scale HPC facilities.

The work by Bahmani and Mueller [2] employed a hierarchical, **signature-based** clustering algorithm using two 64-bit signatures. The first level of clustering is call-path clustering where processes with different numbers or sequences of events are discovered. During the second phase, this algorithm applies parameter clustering using a 64-bit signature. This signature is composed of the parameters of an MPI event, such as its count, type, source, destination, etc.

Even though this previous work is suitable for large-scale applications, it has several limitations. The size of the signature is the first problem. The algorithm is very space efficient due to the 64-bit signature for all of the MPI parameters, but

compressing the parameters could result in loss of information. For instance, 16 bits for message sizes (*COUNT*) is not enough. On the other hand, expanding signatures blindly to a larger scale could increase space complexity significantly and cause scalability problems.

The second limitation is observed for benchmarks such as CG [3]. The unique communication behavior of the processes causes the number of parameter clusters to increase linearly with the number of processes, which is not scalable. To tackle the problem, user input plug-in functions were utilized to specify the communication pattern. The problem with the user plug-in functions is that finding user plug-ins for complex benchmarks is not easy.

The third problem is a scalability challenge due to hierarchical clustering, which already improves on ScalaTrace. ScalaTrace employs a two-stage trace compression technique, namely intra-node (loop level) and inter-node compression. The latter is consolidating traces in a reduction step over a radix tree. While intra-compression is fast and efficient, inter-compression is costly as it depends on the number of tasks. Hierarchical clustering lowered this overhead so that it depends on the number of clusters, but some application codes still require a number of clusters linear to the number of tasks in order to retain trace accuracy.

This paper contributes, ACURDION, a scalable clustering algorithm for large scale applications. ACURDION first finds parameter differences between processes through Adaptive Signature Building, a $O(\log P)$ algorithm. Then, ACURDION applies the K -farthest neighbor algorithm on the selected signatures to group processes into K clusters. This lowers the complexity of inter-node compression to just be dependent on K , which is a constant in practice ($K = 9$ for the programs studied). Consequently, user plug-ins are no longer needed for complex communication patterns, i.e., ACURDION advances trace automation as well.

Contributions:

- We develop, ACURDION, a $O(\log P)$ clustering algorithm with low time and space overheads.
- We develop novel signature finding algorithms that help prune unnecessary metrics and only consider parameters covering differences among the traces.
- We evaluate ACURDION for a set of HPC benchmarks showing its effectiveness at capturing representative application behavior for communication events. The resulting number of clusters is a constant for all benchmarks.
- We compare the accuracy of traces generated by ACURDION with prior work on signature-based trace clustering.

II. BACKGROUND

In this section, we briefly introduce several of the key ideas and techniques relevant to MPI tracing because our work builds on ScalaTrace as an MPI tracing toolset [23].

ScalaTrace uses the following two main data structures: Regular Section Descriptors (RSD) that capture MPI events in the innermost loop, and Power-RSDs (PRSDs) that capture RSDs from higher-level loop nests. Consider the example in the following code segment:

```

for  $i = 0 \rightarrow 1000$  do
  for  $k = 0 \rightarrow 100$  do
     $MPI\_Irecv(\dots);$ 
     $MPI\_Send(\dots);$ 
     $MPI\_Wait(\dots);$ 
  end for
   $MPI\_Barrier(\dots)$ 
end for

```


RSDs and PRSD of the above code segment are RSD1:<100, MPI_Irecv1, MPI_Send1, MPI_Wait>, and PRSD1:<1000, RSD1, MPI_Barrier1>, respectively.

The main three properties of ScalaTrace are as follows: (1) Location-independent encodings: ScalaTrace leverages relative encodings of communication end-points, i.e., an end-point is denoted as $\pm c$ for a constant c relative to the current MPI task ID [15]. Fig. 1(a) depicts the relative encoding of node 7 in terms of communication end-points, namely -4 , -1 , $+1$ and $+4$, i.e., the node has identical *relative* communication end-points.

(2) Call sequence identification: To distinguish between MPI calls from different locations, ScalaTrace captures the calling context by recording the calling sequence that leads to the MPI event. The calling context is obtained from the stack backtrace of an MPI event where each location is a unique signature of the stack trace called the stack signature [15].

(3) Communication group encoding: To store communication groups in a scalable way in traces, ScalaTrace leverages a special data structure called ranklist. A ranklist is represented as $\langle \text{dimension, start_rank, iteration_length, stride, iteration_length, stride} \rangle$, where dimension is the dimension of the group, start rank is the rank of the starting node, and the iteration length stride pair is the iteration and stride of the corresponding dimension [22].

In Fig. 1(b), the shaded nodes are presented as ranklist $\langle 2 \ 0 \ 4 \ 4 \ 2 \ 1 \rangle$. The ranklist reads as a 2D ranklist starting at task 0 with four entries in the first dimension and a stride of 4 (implying tasks 0, 4, 8 and 12) with two entries in the second dimension with a stride of 1 (implying tasks 1, 5, 9 and 13). The ACURDION algorithm introduced in the next section is working on top of the intra-node compression step.



(a) Communication Encoding (b) Group Encoding

Fig. 1: Communication and Group Encodings

III. A NOVEL CLUSTERING ALGORITHM

ACURDION has two main phases. First, the adaptive signature building phase discovers signatures suitable for the clustering. Second, the clustering phase involves a single-step (where all signatures are considered in the K-farthest algorithm), or a double-step where during the first phase it applies K-farthest clustering on the *Call-Path* signature and

then in the second phase it applies the K-farthest algorithm on all other signatures within the *Call-Path* clusters. This section details design and implementation of the single- and double-step algorithms of ACURDION and the reference clustering algorithm.

A. Signature Building

There are eleven signatures by default involved in clustering. The first one, the *Call-Path* signature, helps to cluster processes with different sequences of MPI calls.

As noted previously, to capture calling context, ScalaTrace uses the stack signature consisting of a number of backtrace addresses of the program counter (return addresses), one for each stack frame. The *Call-Path* signature, a 64-bit signature, is the *XOR* of all 64-bit stack signatures. In order to create the *Call-Path* signature, capturing the calling context is sufficient for distinguishing MPI events from each other in most benchmarks. Moreover, to order events, we multiply the modulo 10 plus 1 of the sequence number of each event by the 64-bit stack signature and then use this value in the *Call-Path* signature. This ensures that signatures cannot cancel out each other due to permutations on call sequences and recursion, which could otherwise happen in rare cases (e.g., NAS MG code).

All other signatures are averaged parameter signatures composing parameters of the MPI call event (*COUNT, SRC, DEST, KEY, COLOR, Tag, Computation time, Communication time, LOOP iteration* and *Data+Operation+Communicator type*). For the first eight above-mentioned signatures, aggregating their values and then taking the average could result in an overflow. To address the overflow problem, we utilized an estimation function.

For the *Loop* signature, considering the importance of nested loops, we know that multiplying the bounds of nested loops could cause overflow as well. To avoid this, we divide 64 bits into four sections. The least significant 16 bits are assigned to the average of the inner most loop sizes (or the least important loop), then the second 16 bits are assigned to loop above the first section, etc. Anything above three levels is considered part of the fourth section (most significant bits). As a *Data+Operation+Communicator Type* signature, we assign a bit such as 0:MPI CHAR, 32:MPI MAX, and 55:MPI COMM SELF, etc., per MPI data type, MPI operation type, and MPI communicator types. After creating signatures at the node level, the first step of ACURDION is to enter an adaptive signature building phase.

Set your signature format to 0;

if a left/right child exists **then**

Receive its signatures;

if your signature \neq child signature **then**

update the signature format;

end

Receive the child's signature format;

signature format = signature format OR child's signature format;

end

if a parent exists **then**

Send your signatures to your parent;

Send your format signature to your parent;

end

Broadcast signature format by rank root;

Algorithm 1: Adaptive Signature Building

The time complexity of the algorithm is $\log(P)$. Fig. 2

shows an example of signature building. This procedure uses a set in which there is a bit (per signature) indicating whether or not the signature should be stored. At the beginning, all bits are zero. When two different clusters are partially overlapping, bits corresponding to the different signatures are set to 1. As shown, leaf nodes (such as 5 and 6) send their signatures to the parent node 2. Then, node 2 compares its own signatures with the received ones. If it finds any difference, it flips the corresponding bit in the signature format to 1. Here, nodes 6 and 2 differ (captured at node 2).




Fig. 2: A Sample of Signature Building

This process continues up to the root of the tree. The root then broadcasts the bitmap to all nodes. At the end of this stage, all nodes know which signatures are subjected to clustering.

B. Single-Step and Double-Step Clustering

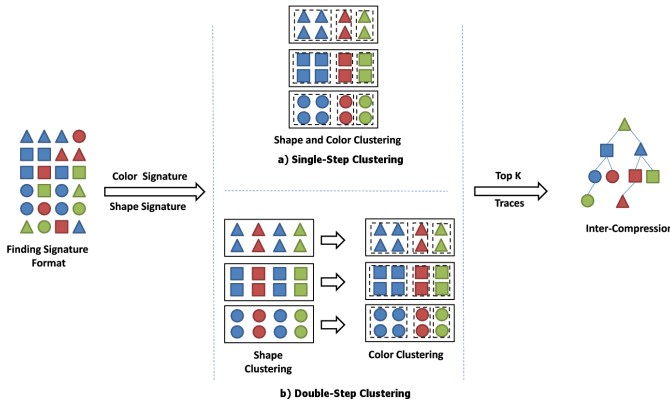


Fig. 3: Overview of Proposed Clustering Algorithm

Fig. 3 provides a simple illustration of ACURDION, where it first finds the signature format, i.e., shape and color in this example. It then either follows the single-step K-farthest algorithm (all signatures are considered in clustering) or the double-step K-farthest algorithm (first, only *Call-Path* or shape in this figure, and second, all other signatures are considered for the clusters created in the first step). In this figure, *color* is the second dimension. After finding the top K clusters and selecting the top K representatives (nodes) to create the global trace file, inter-compression on the selected traces is applied.

In our design and implementation, we considered both Euclidean and Manhattan distances. Before calculating the distances, ACURDION dynamically normalizes them, i.e., it groups signatures based on their importance. The importance ordering is as follows: Group1={COUNT, COMM Time and COMP Time}, Group2={SRC, DEST, KEY, COLOR, TAG} and Group3={Call-Path, Data+Operation+Communicator Type}. Compromising on group 1 does not perturb the application time significantly, as experiments show. *Loops* are covered by group 2 and 3. Group 3 is the most important group

because we do not want to lose any events by compromising on *Call-Path*. We further observed that *Call-Path* also covers the *Data + Operation + Communicator* signature in practice.

Since *COUNT* is a 32-bit integer, its average is also 32-bits. The largest value for *COMP/COMM* time is 136 years, so the largest value of group1 needs 32 bits. In our implementation, we assumed there is a boundary on the maximum value of group 2 which is also 32. With these assumptions, we can shift group 2 based on group 1, and then group 3 based on the smaller groups in such a way that the value of the larger group becomes larger than the smaller ones.

We tested ACURDION on all benchmarks with both the Manhattan and Euclidean distance functions. We then calculated the distance between the output of clustering to the non-clustering version. According to our experiments, both metrics are giving close distances, so we chose the Manhattan distance for our experiments.

if a left/right child exists then

```

Receive list of left_K / right_K clusters;
Receive signature of head of top left_K / right_K clusters;
Merge left_K / right_K clusters + yourself into Top K list;
if left_K + right_K + 1 > K then
    Calculate the distance matrix for Top K list;
    while Size of selected list < K do
        Find the farthest cluster to the selected ones;
    end
    while Other cluster have not been selected do
        Find the closest cluster;
        Assign the cluster to the closest one;
    end
end

```

end

if a parent exists then

```

Send your list of K clusters to your parent;
Send signature of head of top K clusters to your parent;

```

end

Broadcast Top K by root;

Algorithm 2: ACURDION K-farthest algorithm clustering algorithm

ACURDION algorithm operates over a radix tree. At each node, if it has a child, it first receives the K selected clusters and their head's signatures. Then, the received node compares the number of clusters plus one (i.e., its own cluster) to determine whether or not it is less than K . If the result is positive, then it adds its own cluster to the list of selected clusters and sends it to its parent. If not (the number of clusters is larger than K) then it needs to find the top K clusters.

First, it calculates the distance matrix between all potential clusters based on the signature format. Second, it selects the top K clusters farthest from all other selected clusters. Third, it distributes the rest of the clusters (which have not be selected) to their nearest cluster. Fourth, if the current node has a parent, it will send the top K clusters and their signatures to its parent.

In single-step ACURDION, clustering happens over the entire signature format. For example, if *COLOR* and *Call-Path* are parts of a signature format then the algorithm calculates the normalized distances from the signatures. The double-step

version, in contrast, first clusters over *Call-Path* and then (at the second level) within the clusters created at the first level over other dimensions such as *COLOR*.

Note that the computational cost of our clustering algorithm is $O(\log P)$, where P is the number of processes. K can be any constant value. In our experiments, $K = 9$ was shown to preserve sufficient accuracy (discussed in Section V).

By the end of this stage, the algorithm has clustered all processes with disjoint behavior. Then, the algorithm creates the complete trace based on the cluster information.

C. Aggregating the Traces

The time complexity of inter-compression reduction of ScalaTrace is $O(n^2)$, where n is the size of the PRSD-compressed intra-node event trace. Since for ScalaTrace without clustering, all processes are participating in this operation over a radix tree, the time complexity is $O(n^2 \log P)$. On the other hand, for ScalaTrace with the clustering algorithm, only a set of representative K nodes with different signatures are participating in this operation. During the last phase of Fig. 3, K , different nodes are merged.

The cost of inter-compression with clustering is $O(n^2 \log K)$, where $K = 9$ in our experiments and the cost of the clustering algorithm is $O(\log P)$.

To create the full trace out of the clustered trace, before merging, K selected representatives update their trace file considering all members of the cluster. This operation is linear, i.e., representatives linearly traverse their trace and replace a cluster ranklist (representing all members' IDs) with an event ranklist (compatible with the original ScalaTraceV2 format).

D. Reference Signature

We use a reference clustering approach to evaluate the accuracy and scalability of our algorithm. The reference signature utilized by the reference clustering is a sequence of events. It concatenates the *Call-Path* signatures by adding a sequence number to each MPI event, and features parameter clustering by keeping the parameters of each MPI event uncompressed.

In Section V, we provide the results of the experiments conducted with different benchmarks to compare the space complexity of the ACURDION and the reference signature algorithms.

IV. EXPERIMENTAL SETUP

TACC's Stampede [18], a state-of-the-art HPC cluster, is utilized to conduct experiments. It consist of a total of 6400 nodes, each with two Intel Xeon E5 processors and one Intel Xeon Phi coprocessor. The compute nodes are interconnected with Mellanox FDR InfiniBand technology (56 Gb/s) in a two-level fat-tree topology.

Each experiment was run five times, and their averages and standard deviation are reported. The aggregate wall-clock times across all nodes for these benchmarks are reported. We conducted experiments with a variety of codes: (1) the NPB suite (version 3.3 for MPI) with class D input size [3]; (2) Sweep3D [11], a solver for the 3-D, time-independent, particle transport equation on an orthogonal mesh, which uses a multidimensional wavefront algorithm for "discrete ordinates" deterministic particle transport simulation with a problem size of $100 \times 100 \times 1000$; (3) Lulesh, which approximates the hydrodynamics equations discretely by partitioning the spatial problem domain into a collection of volumetric elements [1]. Results of ACURDION are compared to reference clustering and related work on signature-based clustering [2] of which

we obtained a copy.

V. RESULTS AND ANALYSIS

To assess the accuracy of the proposed clustering algorithm, we conducted two types of experiments. First, we tested the accuracy of point-to-point communication through heatmaps. Second, to verify the accuracy of collective operations, we replayed the traces and compared the wall-clock time of the clustered and non-clustered versions.

Our first experiment assesses the effect of ACURDION on point-to-point communication. Fig. 4 depicts heatmaps of Lulesh for 64 processes each. The (a) original version and the (b) ACURDION version of are shown. The x- and z-axes denote mutual communication end-points, and the *number of sends* is depicted on the y-axis. The average time in seconds is depicted as heatmaps (dark=low to white=high). Since the ACURDION heatmaps are a perfect match to the non-clustering ones for BT, Sweep3D, LU and SP, we only show Lulesh's heatmap. The heatmaps for Lulesh differ slightly before (Fig. 4(a)) and after (Fig. 4(b)) clustering (same for CG and MG, figures omitted due to space).

Table I indicates varying parameters selected during the signature building phase for these benchmarks. On average, the size of the signature was reduced by 43%.

TABLE I: Signature Format

Code	Callpath	Count	Src	Dest	Comp	Comm	Tag	Loop
BT		✓	✓	✓	✓	✓	✓	
CG			✓	✓	✓	✓		
LU	✓	✓	✓	✓	✓	✓		
MG	✓	✓	✓	✓	✓	✓		✓
SP		✓	✓	✓	✓	✓		
Sweep3D	✓	✓	✓	✓	✓	✓	✓	✓
Lulesh	✓	✓	✓	✓	✓	✓		

We next define an accuracy metric of trace replay as

$$ACC = 1 - \frac{|t - t'|}{t}$$

where t and t' are the replay times without and with clustering, respectively.

Table II covers the percentage of matching clustered parameters relative to non-clustered ones. The similarity of point-to-point events was already depicted in Fig. 4. Some codes may experience different endpoints in sends/receives after clustering, but the overall patterns are preserved. In other words, if concrete endpoints differ then only by a slight shift so that the overall behavior remains close to the original program, which is also confirmed in terms of wallclock time later.

TABLE II: Matching Percentage

Benchmark	COUNT	LOOP	# EVENTS	SRC	DEST	TAG
BT	99.9%	100%	100%	100%	100%	100%
CG	100%	100%	100%	80.90%	80.90%	100%
LU	97.12%	98.16%	100%	100%	100%	100%
MG	99.7%	100%	100%	96.80%	96.80%	99.03%
SP	100%	100%	100%	100%	100%	100%
Sweep3D	96.86%	87.37%	100%	100%	100%	100%
Lulesh	82.35%	75%	100%	70.37%	70.37%	100%

We chose a maximum number of nine clusters ($K = 9$) for ACURDION, which we experimentally determined based on captured communication patterns of related work [2]. This suffices to represent average communication time, send count, and source and destination ranks for point-to-point communication. Table III also shows that for CG increasing the number of clusters does not improve the accuracy of trace. In fact, we observed that the key element with respect to trace accuracy is the number of *Call-Path* clusters. Covering all distinct events over all traces results in acceptable accuracy.

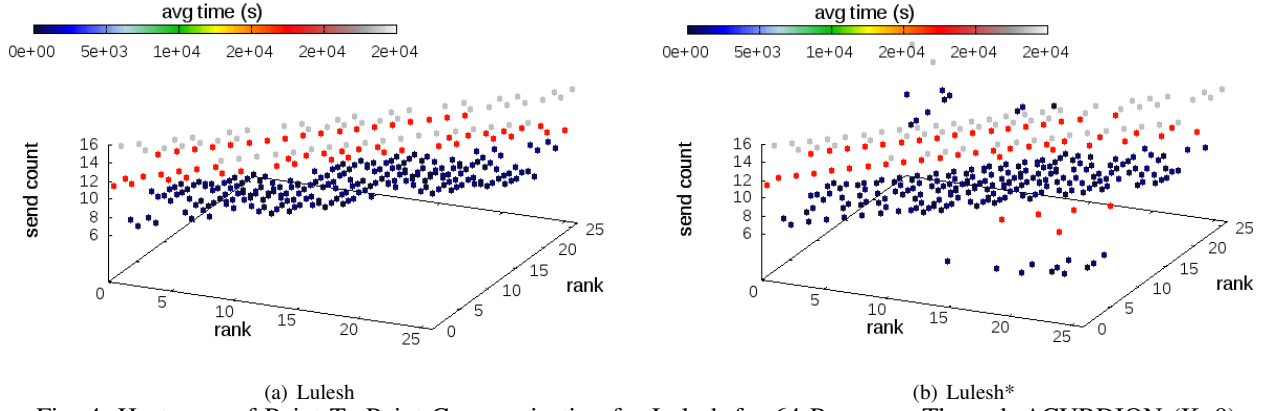


Fig. 4: Heatmaps of Point-To-Point Communication for Lulesh for 64 Processes Through ACURDION (K=9)

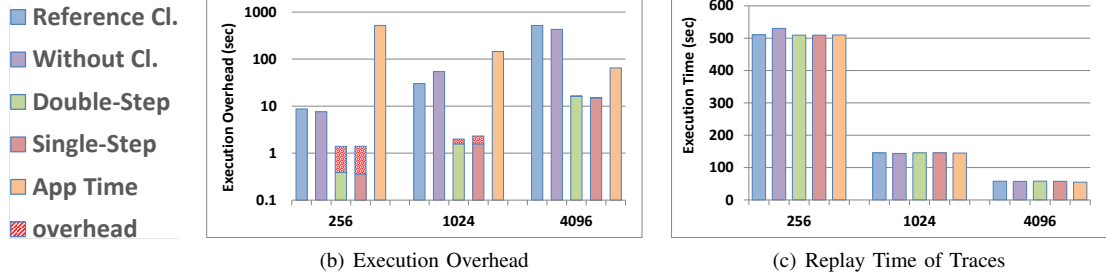


Fig. 5: Execution Overhead and Replay Time for BT benchmark: Strong Scaling, Nodes/Tasks=1/16

Due to the iterative nature of parallel programs, the number of different *Call-Path* patterns is limited to a small number. For our tested benchmarks, this value was nine, which is sufficient to cover stencil codes. Nonetheless, we can change the value of K dynamically should the number of *Call-Path* patterns increase during adaptive signature building.

TABLE III: Accuracy and Number of Clusters: CG Class D and P=256

# Clusters	9	27	81	243
Overhead (s)	1.09	3.22	9.54	27.91
Accuracy	98.92%	99.05%	96.88	97.75%

All benchmarks result in the same number of main clusters and sub-clusters under ACURDION, which often reflects the shape of communication patterns in these codes.

But we also observed minor differences. For example, *BT* and *SP* have the same communication pattern but slightly different message payloads (*COUNT* parameter) after clustering, a difference of around 0.1% reported in Table II. Communication patterns are always retained after clustering, e.g., the stencil pattern of *Sweep3D* and *LU*. Here, *S3D** retains pattern and send volumes while *LU** reflects the communication pattern but diverges slightly in send volume (*COUNT*) under ACURDION.

The Lulesh (Livermore Unstructured Lagrange Explicit Shock Hydrodynamics) proxy application (proxy app) [1] is a shock hydrodynamics code developed at Lawrence Livermore National Laboratory (LLNL). Lulesh is large enough to be more complex than traditional benchmarks, yet compact enough to support a large number of implementations [10]. We observe a slight diffusion in the communication pattern in Fig. 4(b). However, clustering covered all events and retained parameters such as *COUNT* and *LOOP* as Table II showed.

MG has different clusters as data partitioning (sub-grid creation) depends on input size, number of processes, and two different communication patterns (halo/boundary exchange and cross-grid interpolation). Due to changes in grid resolution

per iteration, boundary geometry also changes. As the algorithm moves from coarser to finer, more boundaries are created. Nonetheless, heatmaps (omitted due to space) show that ACURDION clustering captures this pattern relatively accurately as it closely resembles that without clustering (albeit with some shifts of individual points). Table II also indicates that clustering accurately covers the parameters.

CG shows similar shifts that diffuse the regularity of the pattern without clustering (figure omitted due to space), but the overall number of communication exchanges and the send volume are retained. Table II indicates that ACURDION clustering captures the parameters accurately. In related work [2], *CG*'s pattern could only be captured via a user-supplied plugin function, which captured unique parameters that otherwise would significantly increase the total number of clusters. It would be preferable to avoid such user plugins where possible as it is difficult for users to provide such functions for complex benchmarks. ACURDION provides the means to retain a concise trace representation without user plugins, but the price is a more diffuse communication matrix. We will later see that this has little effects on replay accuracy, which shows the benefits of K -farthest clustering.

A. Strong Scaling

The second set of experiments focuses on the accuracy of wall-clock time comparing replayed traces with ACURDION and without clustering, first under strong and then under weak scaling. Some of the benchmarks only support strong scaling (most NAS codes) while others support weak scaling (*Sweep3D*, *Lulesh*) so that different sets of benchmarks are reported in these experiments.

Strong scaling features a set of experiments where the number of tasks is changed while the program input remains the same. This effectively reduces the amount of work per task as the input problem is partitioned into smaller pieces while potentially inflicting more (but smaller) messages as the

number of tasks increases. In these experiments, 16 MPI tasks were mapped onto one node (with 16 cores).

Fig. 5(b) depicts the wall-clock time on a logarithmic scale (y-axis) for different number of MPI tasks (x-axis) of the respective benchmarks. Per task size, the average execution time over five runs is reported for (a) reference clustering, (b) no clustering (vanilla ScalaTrace V2 with intra- and inter-node reduction), (c) double-step and (d) single-step ACURDION clustering, and (e) base application time without instrumentation. For ACURDION results, bars are stacked to distinguish the base instrumentation overhead (blue/bottom) from the clustering overhead (red/top). This distinction is omitted for reference clustering. The last bar (e) is shown as a reference to get an idea how much time would be spent on tracing compared to base application runtime.

For instance, BT for 256 tasks has about an order of magnitude lower trace overhead with ACURDION clustering than without (or with reference clustering), which is nearly two orders of magnitude smaller than application runtime. Within ACURDION, half the time is spent in clustering. For 4096 tasks, ACURDION incurs an order of magnitude lower overhead than reference/no clustering but results in application overhead of about 20%. The clustering time, however, within ACURDION is negligible. Similar observations were made for CG and SP. LU has low instrumentation overhead under clustering (even reference clustering) while the overhead without clustering is significant and outstrips application runtime at 4096 tasks. ACURDION cuts down overheads to about half or even a quarter of that for reference clustering, which is nearly two orders of magnitude smaller than application runtime regardless of the number of tasks (up to 4096 tasks). Most of the tracing overhead is due to clustering under ACURDION. MG’s overhead for ACURDION clustering changes from being two orders of magnitude smaller than application runtime at 256 tasks to match application runtime at 4096 tasks, yet remains about an order of magnitude smaller than reference and no clustering. Overall, single- and double-step clustering perform equally well, and ACURDION outpaces the other techniques due to the lower number of processes involved in inter-node compression after clustering.

The next set of experiments assess the accuracy of the trace information obtained in the techniques featured so far. To this end, a trace replay tool, ScalaReplay [23], is utilized to issue MPI events in the same order and over the same number of nodes that they were originally recorded during application execution. Yet, instead of computing, the recorded time spent for computation is “replayed” as sleep time to resemble the same distance between communication calls. The communication calls themselves are issued with the same parameters as recorded, except for slight differences in send volume and end-points due to clustering (see previous discussion about communication patterns). The message payload is a buffer of the indicated size (but with some random content as content is not recorded during tracing, nor is it required for correct replay as computation has been replaced by sleep). Nodes interpret the same trace file during replay but transpose MPI communication endpoints relative to their task ID (due to the relative encoding of end-points in ScalaTrace). For clustering, a different event is generated per cluster, which results in up to $K = 9$ different events for subsets of tasks (compared to a single event without clustering). All other parameters are replayed directly from the trace.

Fig. 5(c) depicts the replay time in seconds on a linear y-axis for traces resulting from clustering as opposed to *not* using reference clustering, no clustering, single-/double-step ACURDION clustering and also the corresponding application time without instrumentation for comparison. A match to the latter means that traces retain application behavior. We observe that the replay times over all methods, task sizes and applications matches the original application very closely. We observe an accuracy level of more than 95% across the set of benchmarks and experimental parameters. This illustrates that ACURDION is competitive with any other scheme, even though it retains only a subset of the trace information of other methods and requires lower overhead.

B. Weak Scaling

The next experiments cover weak scaling, which features a sequence of experiments where the input size and the number of tasks are increased at about the same rate. The objective is to ensure that the input size per task (after input partitioning) remains constant so that execution times (in the ideal case) also remain constant as we scale up. Of course, changes in communication overhead may influence this behavior. Input constraints on several benchmarks limit the set of experiments that we could conduct for weak scaling to Sweep3D and Lulesh. The input size for Sweep3D is chosen to be $100 \times 100 \times 1000$ per node. For Lulesh, weak scalability tests were run at a problem size of 32^3 per node.

Fig. 6(b) depicts the overheads in seconds on a logarithmic scale (y-axis) for different number of processors (x-axis) for Lulesh and Sweep3D. Lulesh results in about one order of magnitude lower overhead for any clustering approach than without clustering. Single/double-step ACURDION takes about the same time as reference clustering, even though it has to perform the K-farthest algorithm. And clustering techniques incur instrumentation overhead 1-2 orders of magnitude smaller than application runtime. Sweep3D results in even lower overheads of 1-2 orders of magnitude for clustering over no clustering. Its instrumentation cost is 3-5 orders of magnitude smaller than the corresponding application runtime. Here, single-step outperforms double-step slightly whereas in all previous experiments no clear winner could be declared between the two. Reference clustering outperforms ACURDION here for the first time. However, we will later show that ACURDION outperforms reference clustering in terms of space complexity, which can have a significant impact for large-scale tracing.

Fig. 6(c) features the overhead per replay method and in comparison to original application runtime in seconds on a linear scale (y-axis) for a different number of tasks (x-axis). We observe that replay times are uniformly resembling the original application runtime irrespective of which tracing method was used. The overall accuracy of ACURDION is 95%-97%.

C. Signature-based vs. ACURDION

To compare the accuracy of trace files generated by ACURDION to signature-based clustering [2], we conducted two experiments covering both strong and weak scaling on the same platform, where all machines were 2-way SMPs with AMD Opteron 6128 processors with 8 cores per socket. Each node is connected by QDR InfiniBand. In the first experiment, we compared ACURDION and signature-based clustering for CG class C under strong scaling. Figure 7 shows that the accuracy of traces for signature-based and ACURDION are




Fig. 6: Execution Overhead and Replay Time for Lulesh benchmark: Nodes/Tasks=1/16




Fig. 7: Replay Time of Traces: Nodes/Tasks=1/16

98.73% and 98.93% compared to non-clustering. The second experiment covers weak scaling for Sweep3D in Fig. 7. The average accuracy of traces is 97.96% and 97.91% for ACURDION and signature-based, respectively. While ACURDION has a high level of accuracy, it does not have the limitations of the signature-based approach, i.e., no user plug-in is required and it has lower space complexity.

D. Space Complexity

Let us consider the space complexity of tracing with and without clustering. The memory space allocated during trace compression with optional clustering differs from method to method. Table IV depicts the number of clusters and the required memory space in KB per method and per benchmark for 256 tasks. We observe that ACURDION (single/double-step do not differ) requires about an order of magnitude (and up to two orders of magnitude) less space than reference clustering or no clustering. The number of clusters for ACURDION is always nine due to the K-farthest method, which is often significantly smaller than any other technique. (Notice that without clustering, each task is considered a cluster of its own).

TABLE IV: Average Space Complexity Per Process for P=256 (P=216 for Lulesh)

Code	ACURDION		Reference Clustering		Without Clustering	
	# clusters	avg space	# clusters	avg space	# clusters	avg space
BT	9	2.73KB	41	108.49KB	256	71.71KB
CG	9	1.70KB	256	376.32KB	256	43.82KB
LU	9	2.73KB	16	25.05KB	256	71.71KB
MG	9	11.8KB	72	733.83KB	256	215.15KB
SP	9	2.50KB	53	133.06KB	256	67.73KB
Sweep3D	9	1.50KB	9	4.86KB	256	27.89KB
Lulesh	9	1.51KB	26	17.56KB	216	27.87KB

The required space of ACURDION is on average one order of magnitude larger compared to signature-based clustering, i.e., it ranges from slightly smaller (MG) to two orders of magnitude larger (LU). And the number of clusters of ACURDION is sometimes smaller and sometimes larger compared to signature-based clustering. But more significantly, the case (MG) where the number of clusters grows linearly with the number of tasks for signature-based clustering presents a non-scalable behavior. In contrast, ACURDION always requires only $K = 9$ clusters and still retains similar overheads at a constant trace size. This is a significant advance in terms of scaling behavior.

This scalability result is corroborated by a complexity analysis of the algorithms. Without clustering, all processes contribute to the inter-compression step, so the space complexity is linear to number of processes. But for clustering, only a constant number of representative nodes are involved in this operation (e.g., one node per cluster). Furthermore, the size of signatures is a key player for clustering. Thus, we considered the size of signatures and related algorithms for space complexity analysis (e.g., adaptive signature building for ACURDION).

Prior work established the complexity without clustering (1), reference clustering (2) and signature-based clustering (3), which resulted in the lowest overhead. In contrast, the complexity of K-farthest clustering (4) is even lower than (3) since it depends on the constant K for ACURDION, and we showed that $K = 9$ suffices in experiments. Specifically, the average trace and signature sizes per node are multiplied by the constant K (instead of the sum of main clusters and sub-clusters in prior work, which is not constant for some programs, such as MG).

VI. RELATED WORK

Bahmani and Mueller [2] proposed a signature-based clustering algorithm for ScalaTraceV2. ACURDION enhances this work in two directions. First, the parameter signature of [2] is a concatenation of several truncated parameters. At a large scale, 16 bits for representing average COUNT may not be enough to cover all differences in COUNT. The eleven 64-bit signatures that are created based on characteristics of the application by the Adaptive Signature Building phase avoid such deficiencies in ACURDION. Second, the strength of ACURDION lies in its independence of user input plug-ins for benchmarks that have unique communication behavior (e.g., CG).

CYPRESS [27] combines static program analysis with dynamic runtime trace compression. at least an One of the main problems of CYPRESS is the overhead of static and dynamic analysis while the combination of ScalaTraceV2+Clustering [2] does not have any overhead at the compile time and extends to binaries/libraries as well.

A density-based clustering analysis [13], [9]uses an arbitrary number of performance metrics to characterize the application K -means clustering is used in *CHARM++*

to select representative data for migration of objects [12]. These clustering algorithms are expensive in terms of time complexity, especially for large-scale sizes while our work is a low overhead clustering algorithm with $O(\log P)$ complexity.

Phantom [26], a performance prediction framework, uses deterministic replay techniques to execute any process of a parallel application on a single node of the target system. To reduce the measurement time, Phantom leverages a hierarchical clustering algorithm to cluster processes based on the degree of computational similarity. First, the computational complexity for most hierarchical clustering algorithms is at least quadratic in time, and this high cost limits their application in large-scale data sets [24]. Second, since Phantom focuses on performance prediction, it emphasizes computational similarity and does not sufficiently cover communication behavior.

Sampling methods such as Scalasca [28] and CAPEK [7] cannot produce accurate data but rather represent statistical and lossy methods. If the sampling frequency is too low, results may not be representative. Conversely, if it is too high, measurement overhead can significantly perturb the application. Finding an appropriate rate of sampling is complicated, and the cost of having a dense CCT is high. In contrast, ACURDION provides a full trace file without sampling and does so at very low cost by leveraging 64-bit stack signatures.

VII. CONCLUSION AND FUTURE WORK

This work contributes ACURDION, a novel signature-based, K-farthest clustering algorithm with a low time complexity of $O(\log P)$ and low space overheads. A signature finding algorithm prunes unnecessary metrics and only considers parameters representing differences among the traces of nodes.

We evaluated ACURDION in comparison to other clustering algorithms for a set of HPC benchmarks showing its effectiveness at capturing representative application behavior for communication events. ACURDION is superior to past work because it is more scalable in terms of space and time complexities at sustained accuracy. And in contrast to other work, it does not rely on user plugins, which may be hard to construct. Experiments showed that without loss of accuracy, only nine clusters suffice to represent the behavior of a wide set of HPC benchmarks codes.

ACKNOWLEDGMENT

This research was supported in part by the National Science Foundation, awards 1058779 and 1217748.

REFERENCES

- [1] Hydrodynamics Challenge Problem, Lawrence Livermore National Laboratory. Technical Report LLNL-TR-490254.
- [2] A. Bahmani and F. Mueller. Scalable performance analysis of exascale mpi programs through signature-based clustering algorithms. In *International Conference on Supercomputing*, pages 155–164. ACM, 2014.
- [3] D. H. Bailey, E. Barszcz, L. Dagum, and H. D. Simon. Nas parallel benchmark results. *Parallel & Distributed Technology: Systems & Applications, IEEE*, 1(1):43–51, 1993.
- [4] D. Becker, F. Wolf, W. Frings, M. Geimer, B. J. N. Wylie, and B. Mohr. Automatic trace-based performance analysis of metacomputing applications. In *International Parallel and Distributed Processing Symposium*, pages 1–10, 2007.
- [5] H. Brunst, M. Winkler, W. E. Nagel, and H.-C. Hoppe. Performance optimization for large scale computing: The scalable vampir approach. In *Computational Science-ICCS 2001*, pages 751–760. Springer, 2001.
- [6] I. Corp. Intel trace analyzer and collector. <https://software.intel.com/en-us/intel-trace-analyzer>.
- [7] T. Gamblin, B. R. De Supinski, M. Schulz, R. Fowler, and D. A. Reed. Clustering performance data efficiently at massive scales. In *International Conference on Supercomputing*, pages 243–252. ACM, 2010.
- [8] M. Geimer, F. Wolf, B. J. N. Wylie, E. Abraham, D. Becker, and B. Mohr. The scalasca performance toolset architecture. In *International Workshop on Scalable Tools for High-End Computing*, June 2008.
- [9] J. Gonzalez, J. Gimenez, and J. Labarta. Automatic detection of parallel applications computation phases. In *International Parallel and Distributed Processing Symposium*, pages 1–11. IEEE, 2009.
- [10] I. Karlin, A. Bhatlele, J. Keasler, B. Chamberlain, J. Cohen, Z. DeVito, R. Haque, D. Laney, E. Luke, F. Wang, D. Richards, M. Schulz, and C. Still. In *International Parallel and Distributed Processing Symposium*.
- [11] K. R. Koch, R. S. Baker, and R. E. Alcouffe. Solution of the first-order form of the 3-d discrete ordinates equation on a massively parallel processor. *Transactions of the American Nuclear Society*, 65(108):198–199, 1992.
- [12] C. W. Lee, C. Mendes, and L. V. Kalé. Towards scalable performance analysis and visualization through data reduction. In *International Parallel and Distributed Processing Symposium*, pages 1–8. IEEE, 2008.
- [13] G. Llort, J. Gonzalez, H. Servat, J. Gimenez, and J. Labarta. On-line detection of large-scale parallel application’s structure. In *International Parallel and Distributed Processing Symposium*, pages 1–10. IEEE, 2010.
- [14] G. Marin and J. Mellor-Crummey. Cross architecture performance predictions for scientific applications using parameterized models. In *SIGMETRICS Conference on Measurement and Modeling of Computer Systems*, 2004.
- [15] M. Noeth, P. Ratn, F. Mueller, M. Schulz, and B. R. de Supinski. Scalatrace: Scalable compression and replay of communication traces for high-performance computing. *Journal of Parallel and Distributed Computing*, 69(8):696–710, 2009.
- [16] S. S. Shende and A. D. Malony. The tau parallel performance system. *Int. J. High Perform. Comput. Appl.*, 20(2):287–311, May 2006.
- [17] A. Snively, L. Carrington, N. Wolter, J. Labarta, R. Badia, and A. Purkayastha. A framework for performance modeling and prediction. In *Supercomputing*, Nov. 2002.
- [18] T. A. C. C. Stampede. <http://www.tacc.utexas.edu/resources/hpc/stampede>.
- [19] R. Susukita et al. Performance prediction of large-scale parallel system and application using macro-level simulation. In *Supercomputing*, 2008.
- [20] J. Vetter and M. McCracken. Statistical scalability analysis of communication operations in distributed applications. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, 2001.
- [21] J. Vetter and F. Mueller. Communication characteristics of large-scale scientific applications for contemporary cluster architectures. *Journal of Parallel Distributed Computing*, 63(9):853–865, Sept. 2003.
- [22] X. Wu and F. Mueller. Scalaextrap: Trace-based communication extrapolation for spmd programs. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 113–122. ACM, 2011.
- [23] X. Wu and F. Mueller. Elastic and scalable tracing and accurate replay of non-deterministic events. In *Conference on International Conference on Supercomputing, ICS '13*, pages 59–68. ACM, 2013.
- [24] R. Xu, D. Wunsch, et al. Survey of clustering algorithms. *IEEE Transactions on Neural Networks*, 16(3):645–678, 2005.
- [25] T. Yang, X. Ma, and F. Mueller. Cross-platform performance prediction of parallel applications using partial execution. In *Supercomputing*, Nov. 2005.
- [26] J. Zhai, W. Chen, and W. Zheng. Phantom: Predicting performance of parallel applications on large-scale parallel machines using a single node. In *ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 305–314, 2010.
- [27] J. Zhai, J. Hu, X. Tang, X. Ma, and W. Chen. Cypress: Combining static and dynamic analysis for top-down communication trace compression. In *Supercomputing*. To appear, 2014.
- [28] I. Zhukov and B. J. N. Wylie. Assessing measurement and analysis performance and scalability of scalasca 2.0. In *Proc. of the Euro-Par 2013: Parallel Processing Workshops*, volume 8374 of *Lecture Notes in Computer Science*, pages 627–636. Springer, 2014.