Published in final edited form as:

Int J High Perform Comput Appl. 2017 January ; 31(1): 32–51. doi:10.1177/1094342015594519.

Application Performance Analysis and Efficient Execution on Systems with multi-core CPUs, GPUs and MICs: A Case Study with Microscopy Image Analysis

George Teodoro^{1,*}, Tahsin Kurc^{2,3}, Guilherme Andrade⁵, Jun Kong⁴, Renato Ferreira⁵, and Joel Saltz^{2,3}

¹Department of Computer Science, University of Brasília, Brasília, DF, Brazil

²Department of Biomedical Informatics, Stony Brook University, Stony Brook, NY, USA

³Scientific Data Group, Oak Ridge National Laboratory, Oak Ridge, TN, USA

⁴Department of Biomedical Informatics, Emory University, Atlanta, GA, USA

⁵Department of Computer Science, Federal University of Minas Gerais, Belo Horizonte, MG, Brazil

Abstract

We carry out a comparative performance study of multi-core CPUs, GPUs and Intel Xeon Phi (Many Integrated Core-MIC) with a microscopy image analysis application. We experimentally evaluate the performance of computing devices on core operations of the application. We correlate the observed performance with the characteristics of computing devices and data access patterns, computation complexities, and parallelization forms of the operations. The results show a significant variability in the performance of operations with respect to the device used. The performances of operations with regular data access are comparable or sometimes better on a MIC than that on a GPU. GPUs are more efficient than MICs for operations that access data irregularly, because of the lower bandwidth of the MIC for random data accesses. We propose new performance-aware scheduling strategies that consider variabilities in operation speedups. Our scheduling strategies significantly improve application performance compared to classic strategies in hybrid configurations.

Keywords

Hybrid Systems; GPGPU; Intel Xeon Phi; Cooperative Execution; Image Analysis

1. Introduction

Hardware accelerators (co-processors) are becoming increasingly important in scientific computing. Through massive parallelism and high-bandwidth memory sub-systems, a co-processor can deliver significant performance gains, sometimes at a lower energy footprint,

in many computationally difficult applications (1; 2; 3; 4; 5; 6). As a result, there is a rapid adoption in high-end systems of hybrid configurations equipped with multi-core CPUs and one or more co-processors. The architectures and execution models of co-processors differ from CPUs and among co-processor types, making it challenging for application developers to optimize their applications. Evaluating and understanding the performance characteristics of co-processors for common computation patterns and operations in an application class can help in designing more efficient operations and applications, in selecting the most suitable co-processors for an application, and in developing methods for efficient application execution on hybrid configurations with multiple types of co-processors.

Our work targets optimizations for and performance characterizations of operations and pipelines that are commonly employed in an important class of applications that analyze low-dimensional spatial datasets captured by high resolution sensors. This class includes applications that analyze and mine large datasets of tissue images captured by high resolution digital microscopes; that process data from satellites and ground-based sensors in weather and climate modeling; that make use of satellite data in large scale biomass monitoring and change analyses; that analyze seismic surveys in subsurface and reservoir characterization; and that process wide field survey telescope datasets in astronomy (7; 8; 9; 10; 11; 12). Input data and data products in these applications are typically contained in low-dimensional spaces (2D or 3D coordinate system with a temporal component). Common data processing steps include identification and segmentation of regions or objects of interest, computation of the texture, topology (or structure) and temporal features of segmented objects, classification of objects and object groups based on the features, and monitoring and quantifying changes over space and time. Table 1 presents the categories of common operations in this class of applications and presents examples in microscopy image analysis, weather prediction, and monitoring and change analysis. These operations can be composed into analysis pipelines to address different scientific questions. The data access and processing patterns of the operation categories range from local and regular to irregular and global access to data – see Table 2 (13). Local data access patterns correspond to accesses to a single data element or data elements within a small neighborhood in a spatial and temporal region (e.g., data cleaning and low-level transformations). Regular access patterns involve sweeps over data elements, while irregular accesses may involve accesses to data elements in a random manner (e.g., certain types of object classification algorithms, morphological reconstruction operations in object segmentation). Some data access patterns may involve generalized reductions and comparisons (e.g., aggregation) and indexed access (e.g., queries for data subsetting and change quantification).

Motivated by the requirements of this class of applications, in this paper we study the performance impact of GPUs and Intel Xeon Phi (Many Integrated Core – MIC) co-processors on a set of core operations in the object segmentation and feature computation categories from Table 1. We also develop and experimentally evaluate several task scheduling strategies for execution of pipelines of these operations on hybrid systems with multiple co-processors. We have chosen a microscopy image analysis application, which has been developed to support large-scale brain tumor studies (14; 8; 15), as a real application scenario for performance evaluation. Our choice is driven by the fact that our group has extensive experience with this application domain; we have implemented methods for object

segmentation, feature computation, and classifications and analysis pipelines and developed high performance implementations of the methods and pipelines. In digital pathology high resolution images obtained from tissue specimens using advanced digital microscopy instruments are processed to detect and segment nuclei and cells and compute shape and texture features on the segmented objects. These features are then mined to investigate the morphology at sub-cellular level of disease and how it correlates with disease onset and progression and with response to treatment. In general, execution of a microscopy image analysis application requires a lot of computing power and memory, because of the computational complexity of analysis pipelines and the sizes of images – a three-channel color image captured by a state-of-the-art scanner can reach 120K×120K pixels in resolution. In addition, modern scanners are able to collect images quickly and studies with a few thousand images are becoming common. Adding to the computational cost is the fact that analysis pipelines may need to be executed multiple times to tune its parameters and methods for best results or to determine how sensitive an analysis pipeline is to its input parameters. Microscopy image analysis application has many similarities, in terms of data access and processing patterns, to other applications in our target application class. For example, a change monitoring and analysis application will use measurements and images obtained from remote sensing instruments (e.g., satellites) (16; 17; 18; 19; 20). It will apply similar object detection, feature computation, and classification pipelines (17; 18) to detect and classify different objects and regions on earth (e.g., trees, forests, residential areas) – rather than nuclei and tumor regions in a microscopy image analysis application. It may track changes in these areas, like a microscopy image analysis application may track changes in distribution of certain types of nuclei and tumor regions.

Our main contributions can be summarized as follows. First, we characterize operations in the object segmentation and feature computation categories for the microscopy image analysis application according to their data access and computation patterns. Second, we evaluate the performance of the operations on modern CPUs, GPUs and MIC co-processors. We empirically show that the performance gains of the operations vary significantly according to their data access pattern and parallelization strategies. For instance, the performance on a MIC for operations that perform regular data access is comparable or sometimes better than that on a GPU, but GPUs are more efficient for algorithms that irregularly access data or rely on atomic operations. Third, we develop new performance-aware scheduling strategies: PADAS and PAMS, for efficient cooperative execution on hybrid systems, equipped with CPU and one or multiple accelerators. PAMS attains better performance than all other policies, including Heterogeneous Earliest Finish Time (HEFT), for all experiments with multiple accelerators with gains of at least 1.16× on top of the competitors. Finally, we experimentally show that our implementation of the application scales well on distributed memory systems with different hybrid node configurations: (i) nodes with CPU and MIC and (ii) nodes with CPU + MIC + GPU.

This paper extends our previous work (21), which carried out a performance comparison of image processing operations on multiple devices, by proposing and evaluating task scheduling strategies for execution on hybrid configurations with multiple types of co-processors. The rest of the paper is organized as follows: We describe the motivating application and the data access and computation characteristics of its core operations in

Section 2. Section 3 presents the the programming models and parallelization strategies used to implement the core operations on GPUs, MICs, and multi-core CPUs. The deployment of the application for execution on distributed memory systems and new performance-aware scheduling policies are presented in Section 4. We perform an experimental evaluation in Section 5, present related work in Section 6, and conclude the paper in Section 7.

2. Example Motivating Application and Core Operations

We provide a brief overview of the microscopy image analysis application. Our work is presently focused on the development of operations in the object segmentation and feature computation categories (or stages), because these are the most expensive stages in this application. We describe the operations and present their data access and processing patterns.

2.1. Microscopy Image Analysis Application

The microscopy image analysis application is developed primarily to support in silico studies of brain cancer (22; 15; 23). These studies seek more accurate classifications of tumors by leveraging complementary datasets that include high resolution whole tissue slide images (WSIs), clinical data and gene expression data. WSIs are high resolution images with up to 120K×120K pixels per image. Analysis of a WSI typically involves: 1) preprocessing to reduce effects of image acquisition parameters (e.g., reduce color intensity variations due to different stainings via color normalization); 2) segmentation of microanatomic objects, such as nuclei and their cytoplasm regions; 3) feature computation to generate a set of shape and texture properties for each object identified in the segmentation stage; and 4) clustering of patients (from whom the WSIs have been obtained) via machine learning algorithms. The most time consuming stages are the segmentation of objects and the computation of features, because the clustering stage usually is performed on a significantly reduced dataset, which is realized by aggregating and averaging object-level features from each image to compile patient-level morphology profiles.

The segmentation and feature computation stages are built by composing several finer-grain operations as shown in Figure 1. The operations in segmentation include of red blood cells for calculation of potential object seeds; morphological open is used to remove noise and fill small holes in seeds; morphological reconstruction (ReconToNuclei) identifies candidate objects; area threshold filters out objects that are not within a given size range; fill holes fills holes in objects; distance transform operation is used to separate overlapping objects; and connected component labeling (CCL) creates a labeled mask in which the pixels belonging to a segmented object have the same label. The feature computation stage receives the labeled masks, object bounding boxes, and the image as input. It executes a set transformations on the input color image, such as Gradient, Sobel, and Color Deconvolution, which are then used to derive an array of quantitative attributes for each object identified in the segmentation stage.

2.2. Data Access and Computation Characteristics of Core Operations

The operations in Figure 1 are categorized in Table 3 according to data access patterns, computation intensity, and the approach for parallel execution. This classification is intended to facilitate extrapolation of results from our performance evaluations beyond the scope of the example application. We expect that operations sharing these characteristics in other applications will have similar performance trends. The last column (CPU Source) of the table indicates the source of single-core CPU implementations of the respective operations. We have used implementations from popular libraries (OpenCV (24)) and from other research groups wherever possible to ensure that our parallel implementations are compared to efficient baseline implementations.

The operations are first classified according to their data access patterns: 1) regular operations access and process contiguous elements or chunks of data in the domain space, such as scanning pixels in an input image; 2) irregular operations access data elements irregularly distributed in the data domain, and data elements to be accessed are determined at runtime as the computation progresses and dependencies are resolved. Within each of these classes, the operations can further be categorized based on how the value of a data element is computed or updated: 1) local when the computation involves only the value of the elements being processed; 2) multi-channel local if multiple data elements with the same spatial location across multiple layers of the domain are accessed (e.g, the same index in multiple image channels); 3) neighborhood if computation on a data element involves values from data elements in a spatial neighborhood of the data element. The neighborhood could be defined using a structure centered at the element being operated on. Examples of these structures are 4-/8-connected components or discs (defined with a radius from the current element); and 4) areas within bounding-boxes if computations on data elements involve data elements within a region defined by a bounding-box. This type of computation is primarily a characteristic of the operations in the feature computation stage that process data in regions defined by the minimum bounding boxes of segmented objects.

Implementations of the operations for parallel execution can employ a variety of parallelism strategies: 1) Data parallelism; 2) Object parallelism; 3) MapReduce (29) or generalized reduction; 4) Irregular wavefront propagation pattern (IWPP) (25); and 5) Union-find (30). Data parallelism is suitable for operations that carry out computation of elements in the data domain independently. Object parallelism corresponds to parallel and independent computation of data objects. The Generalized Reduction or MapReduce pattern is used in the “Area Threshold” operation. The first step of Area Threshold computes a mapping of foreground pixels based on their labels – the label of a pixel indicates which candidate object it belongs to; pixels that are not part of a candidate object are labeled with 0 (zero). The second step executes a reduction to count the number of pixels with the same label value. The count indicates the area of the corresponding candidate object. In the last step, objects whose areas are not within a given range are removed from further consideration.

The IWPP pattern executes independent wavefront propagations that may start from multiple source elements in the domain. Elements located in the front of these waves, called active elements, serve as sources of propagation to their neighbor elements. The value of an element may be modified by a propagation from neighbor active elements in a wavefront.

The IWPP pattern is presented in Algorithm 1. A set of (active) elements from a multi-dimensional grid space (D) is selected to initialize the wavefront (S). The algorithm then enters the wavefront propagation phase in which an element (e_i) from S is removed for computation. That element will attempt to propagate its value to neighboring elements (N_G) in the G structure. Whenever the propagation condition between e_i and each neighbor ($e_j \in Q$) is satisfied, e_j receives the propagation and is added to the set of active elements in the wavefront. Since e_j has its value modified, it may now be able to change the value of one of its neighbors. This process is similar to the process that takes place in flood-fill algorithms. The wavefront propagations are repeated until the set S is empty.

Because the management of the wavefront is a core part of the IWPP algorithm, the algorithm's performance (in both sequential execution and parallel execution) depends on using an efficient means (e.g., using a queue or a set) to track active elements and, as a consequence, to avoid computations by data elements that do not contribute to the output.

Algorithm 1 Irregular Wavefront Propagation Pattern (IWPP).

```

1:  $D \leftarrow$  data elements in a multi-dimensional space
2: {Initialization Phase}
3:  $S \leftarrow$  select active elements from  $D$ 
4: {Wavefront Propagation Phase}
5: while  $S \neq \emptyset$  do
6:   Extract  $e_i$  from  $S$ 
7:    $Q \leftarrow N_G(e_i)$ 
8:   while  $Q \neq \emptyset$  do
9:     Extract  $e_j$  from  $Q$ 
10:    if  $\text{PropagationCondition}(D(e_i), D(e_j)) = \text{true}$  then
11:       $D(e_j) \leftarrow \text{Update}(D(e_j))$ 
12:      Insert  $e_j$  into  $S$ 

```

The union-find pattern (30) is used for manipulating disjoint-set data structures that may be merged efficiently. Rooted trees are used to represent sets; an arbitrary member of a set is used as its representative or root of the tree. The non-root element points to its parent and root points to itself. The union-find is implemented using the following four operations: (1) **MakeSet(x)**: creates an elementary set containing only x , which is not a member of any other set; (2) **FindRoot(x)**: determines the root of the set in which x is stored; (3) **Union(x,y)**: merges the sets containing elements x and y ; (4) **Criterion(x,y)**: determines if x and y belong to the same set. The connected components labeling (CCL) operation implements the union-find pattern. The CCL first creates a forest in which each element (pixel) from the input image is an independent tree. It iteratively merges trees from adjacent elements in the data domain such that one tree becomes a branch in another tree. The criterion for merging trees is that the neighbor elements must be foreground pixels in the original image. When merging two trees (Union), the label values of the roots of the two trees are compared, and the root with the smaller value is selected as the root of the merged tree. After this process has been done for all the pixels in the image, each connected component is assigned to a single tree.

3. Implementations of Operations on multi-core CPU, MIC and GPU

Programming tools and languages for code development for a MIC are similar to those used for CPUs. This is a significant advantage as compared to GPUs. It alleviates code migration overheads. The MIC supports several parallel programming languages and models, e.g., OpenMP, Intel Threading Building Blocks (TBB), and Intel Cilk Plus. In this work, we have implemented the parallel versions of the operations on the MIC and CPU using OpenMP, because it is a well known parallelization framework that only requires annotation of the code. The implementations for the GPU are done with CUDA¹. The MIC supports two execution modes: native and offload. In the native mode the application runs entirely within the co-processor, because MICs run a specialized Linux kernel that provides the necessary services and interfaces to applications. The offload mode allows for the CPU to execute regions of the application code with a MIC. These regions are defined using pragma tags and include directives for transferring data. The offload mode also supports conditional offload directives, which can be used to decide at runtime whether a region should be offloaded to the co-processor or should be executed on the CPU. This feature is used in our dynamic task assignment strategy to use the CPU and MIC cooperatively for application execution. We present the details of the MIC and GPU implementations only, because the CPU runs the same code as the MIC.

The data parallel operations are trivial to implement, since computing threads may be assigned for independent computation of elements from the input data domain. However, we had to analyze the results of the auto vectorization performed by the compiler for the MIC, since it could not vectorize some of the loops with complex pointer manipulations. We annotated those codes with (#pragma simd) directives to guide the vectorization where appropriate.

The implementation of an operation with the IWPP pattern makes use of efficient parallel containers to store the wavefront elements. The parallel computation of elements in the wavefront requires those elements be atomically updated, as multiple elements may concurrently update an element e_j . We have developed a complex hierarchical parallel queue to store wavefront elements for the GPU implementation (25). The parallel queue exploits the multiple GPU memory levels and is implemented in a thread block basis, such that each block of threads has an independent instance of the queue to avoid synchronization among blocks. The implementation of the IWPP for the MIC utilizes the standard C++ queue data structure as the container. The IWPP code instantiates one copy of this container per computing thread. Each computing thread independently carries out propagations of a subset of wavefront elements. Atomic operations are used to update memory during a propagation to avoid race conditions. As a result, the MIC vectorization is not possible since vector atomic instructions are not supported.

The MapReduce-style pattern is employed in object area calculations. The MIC and GPU implementations use a vector with an entry per object to accumulate its area. Threads concurrently scan pixels in the input data domain to atomically increment the corresponding

¹<http://nvidia.com/cuda/>.

entry in the reduction vector. Because the number of objects may be very high, it is not feasible to create a copy of this vector per thread and eliminate the need for atomic instructions.

In the Union-find pattern a forest is created in the input image, such that each pixel stores its neighbor parent pixel or itself when it is a tree root. For the parallelization of this pattern, the input image data is divided into tiles that may be independently processed in parallel. A second phase is executed to merge trees that cross tile boundaries. The MIC implementation assigns a single tile per thread, which eliminates the need for atomic instructions in the first phase. The GPU implementation, on the other hand, computes each tile using a thread block. Since the threads computing a tile are in the same block, they can take advantage of the fast shared-memory atomic instructions. The implementation of the second phase of Union-find is similar for both the MIC and the GPU. It uses atomic updates to guarantee consistency during tree merges across tile boundaries. The object parallel operations are concurrently executed with each instance processing a segmented object. A single thread in the MIC or a block of threads in the GPU is assigned for the computation of features related to each object. All the operations with this type of parallelism are fully vectorized.

4. Execution on Hybrid Cluster Systems

Execution of the microscopy image analysis application on a hybrid cluster system is supported by a runtime system designed for dataflow applications and developed in our previous work (31; 32). This runtime system expresses an application as a hierarchical dataflow graph with multiple levels, which allows for a computation task in one level to be described as another dataflow of finer-grained operations. The hierarchical representation allows for more flexibility in scheduling of application stages and operations. For example, different scheduling strategies can be used in each level. The microscopy image analysis application is represented in this runtime system as a two-level computation graph. The first level consists of coarse-grain computing stages: segmentation and feature computation. Each of these stages is represented as a graph of finer-grain operations which forms the second level. This is a natural representation of the application, which, as presented in Figure 1, has a hierarchical organization of computation.

4.1. Execution on a Distributed-memory Computation Cluster

Our parallel execution strategy is based on a Manager-Worker model that combines a bag-of-tasks style execution with coarse-grain dataflow and makes use of function variants (see Figure 2) – a function variant represents multiple implementations of a function with the same signature; in our case, the function variants of an operation are the CPU, GPU, and MIC implementations. Each image in a dataset is partitioned and processed in tiles. Image tiles can be processed concurrently. The parallel execution strategy assigns the first level (coarse-grain) stage tasks and image tiles to the nodes of the cluster in a demand-driven fashion. Within each node, it employs another task scheduling method (see Section 4.2) to map the second level (finer-grain) tasks to computing devices.

One Manager process is instantiated on the head node and each computation node is designated as a Worker. The Manager creates tasks of the form (input image tile, processing

stage), where processing stage is either the segmentation stage or the feature computation. Each of these tasks is referred to as a stage task. The Manager also builds the dependencies between stage task instances to enforce correct execution. The stage tasks are scheduled to the Workers using a demand-driven approach, and a Worker may request and compute multiple stage tasks in parallel to use all the computing devices on a node. A local Worker Resource Manager (WRM) is created on each Worker to manage CPU cores and co-processors (GPUs or MICs) available. The WRM creates a device manager thread to control each accelerator or CPU core. When the Worker receives a stage task, it creates fine-grain tasks that implement the stage task and dispatch them for execution with the WRM. The assignment of fine-grain tasks to computing devices in a node is carried out in a demand-driven basis. Whenever a device manager thread becomes idle, it requests a task for execution with the WRM. The WRM then selects a task among those with dependencies resolved using either a First Come, First Served (FCFS) strategy or one of the performance-aware scheduling algorithms proposed in Section 4.2. After all fine-grain tasks of a given coarse-grain stage have been executed, the WRM notifies the Manager process and requests more coarse-grain tasks. This process continues until all coarse-grain or stage tasks have been processed.

4.2. Performance-Aware Task Scheduling

We propose new scheduling strategies for execution of tasks on a node with multi-core CPUs and multiple co-processor types or models. In our earlier work (21), we used a Performance-Aware Task Scheduling (PATS) strategy to efficiently assign computation for cooperative execution on a node with CPUs and one type or model of co-processor. PATS assigns tasks to CPU cores or co-processors based on the load of the computing devices and an estimate of each task's speedup on the co-processor. When a co-processor requests a task, PATS assigns the task with estimated higher speedup to the co-processor. If the device available for computation is a CPU, the task with the estimated lower speedup is chosen. PATS was built targeting systems with CPUs and a single type of co-processor. It is not able to handle scheduling in environments with CPUs, GPUs, and MICs. This limitation of PATS motivated the development of new performance-aware scheduling policies presented in the remaining of this section. In these policies, assignment of tasks to devices is performed on a demand-driven basis as computing devices become available.

Performance-Aware Dequeue Adapted Scheduler (PADAS): is implemented using a single queue to hold tasks ready for execution. This queue maintains tasks in sorted order based on the expected speedup of each task on the co-processors. Insertion of tasks in this queue may start from the head or tail of the queue. If the expected acceleration of a task on a MIC is higher than that on a GPU, the task is inserted into the queue from the head. The task is placed in the queue before the first task with a smaller speedup on the MIC. If the expected speedup on a GPU of the task is higher than that on the MIC, the task is inserted from the end of the queue. The task is placed in the queue before the first task with a smaller expected speedup on the GPU. The reasoning with this strategy is to keep tasks that benefit the most from the MIC in one end of the queue and tasks more suitable for the GPU in the other. The use of a single queue leads to reduced overheads and allows for quick retrieval of the most appropriate task to the type of processor available. When a MIC requests a task, PADAS will

retrieve the first task from the head of the queue. When a GPU is available, the task in the tail is chosen. The CPU cores consume tasks from the middle point between tasks inserted from the head (MIC) and end (GPU) of the queue (see Figure 3).

Performance-Aware Multiqueue Scheduler (PAMS): uses multiple queues; one queue for each type of computing device. The task queues are maintained sorted according to the expected speedup of tasks on the target device. The GPU and MIC queues are sorted in descending order, whereas the CPU queue uses the $\max(GPU\ Speedup, MIC\ Speedup)$ as a metric to order its tasks in ascending order. This strategy aims to use CPU cores for tasks that are not efficiently executed by co-processors. Task insertions are performed in all queues. A pointer to an inserted task exists in each queue. The pointer is used to efficiently delete a task from the other queues once it is selected for execution from a queue. The process of scheduling a task for computation is simple. When a device becomes idle, the task on the head of its queue is selected for execution, and the pointers built during the insertion are used to quickly delete that task from the other queues. Figure 3 presents a scheme of this scheduling strategy.

The performance-aware scheduling strategies we proposed only rely speedups to maintain order of tasks in queues. Therefore, if errors in estimate of speedups are not sufficient to change order of tasks, the performance of our policies will not be affected. In fact, our performance evaluation shows that scheduling strategies based on execution times, such as Heterogeneous Earliest Finish Time (HEFT), are more sensitive to errors in estimates. As a consequence, they require more accurate input to compute an efficient mapping of tasks to processors.

Task scheduling in hybrid machines has been investigated in a number of projects (33; 34; 35; 36; 37; 38; 39; 40). Scheduling strategies developed in prior work partition tasks among computing devices (i.e., CPUs and GPUs) for applications in which application tasks/operations attain the same level of speedup on an accelerator. However, applications that apply multiple, different transformations to input data are composed of operations which accelerate at different speeds on an accelerator. The scheduling strategies we proposed take these speedup variations into account in order to better use hybrid system, whereas the other approaches will commonly only try to minimize load imbalance among computing devices.

5. Experimental Performance Evaluation

We conducted a series of experiments to compare the performance impact of co-processors on image analysis operations and to evaluate the scheduling algorithms for efficient execution on a computational cluster with co-processors. We used a distributed-memory Linux cluster, called Stampede, for this purpose. Stampede is one of the high performance computing resources provided by XSEDE²; the machine is hosted at Texas Advanced Computing Center³. Each compute node of Stampede has dual socket Intel Xeon E5-2680 processors, an Intel Xeon Phi SE10P co-processor and 32GB main memory. The cluster also

²<https://www.xsede.org/home>

³<https://portal.xsede.org/tacc-stampede>

has 128 nodes each with one SE10P and one NVIDIA K20 GPU. The hardware details of the Intel Phi and GPU used in our evaluation are presented in Table 4. All the nodes in the cluster are connected via Mellanox FDR InfiniBand switches.

In all of the experiments the MIC based codes were executed using the offload mode. One MIC computing core was reserved for running the offload daemon. The MIC and multi-core CPU codes for the image analysis operations were implemented using C++ and OpenMP, while the GPU codes were developed using CUDA. We measured processing performance using a set of images which had been collected from The Cancer Genome Atlas repository for brain tumor studies (15). Each image is partitioned into 4K×4K tiles, which can be processed concurrently for image analysis.

While we used a microscopy image analysis application for performance evaluation, we believe our findings will be applicable in other applications that have data access and processing characteristics similar to those of the microscopy image analysis application. These characteristics can be summarized as follows: (1) spatial data domain is partitioned into chunks (or tiles), and tiles are processed concurrently; (2) processing of each tile is expressed as a pipeline of operations; (3) the pipeline operations have a variety of data access and processing patterns (see segmentation and feature computation categories in Tables 1 and 2), which lead to variable speedups on co-processors; (4) the variability in speedups also arise from the fact that some pipeline operations are data intensive while others are computationally more expensive. We observed similar processing characteristics in applications that analyze satellite imagery and that process telescope datasets.

5.1. Performance and Scalability of Operations on MIC

Thread affinity (i.e., how process threads are mapped to the physical cores of a MIC) can affect the performance of operations running on a MIC. In this set of experiments we evaluate the performance impact of three thread affinity strategies (Figure 4) on the core operations described in Section 2. The first strategy, called Compact, binds threads to the next free thread context. That is, all four contexts in a physical MIC core are used before threads are placed in the contexts of another core. The second strategy, called Balanced, allocates threads to new cores before all the contexts in the same core are used. Threads are balanced among cores and subsequent thread IDs are assigned to neighbor contexts. The last strategy, Scatter, also balances the computation, but threads with subsequent IDs are assigned to different physical cores.

The experimental evaluation was performed using two operations with different data access and computation patterns: (1) Morphological Open has a regular data access pattern and low computation intensity; (2) Distance Transform performs irregular data accesses and has moderate computation intensity. The codes employed OpenMP static loop scheduling, because it led to better performance.

As is shown in Figure 5, the scalability of the operations varies significantly under different thread affinity strategies. In addition, speedup values peak when the number of threads is a multiple of the number of computation cores (60, 120, 180, and 240 threads). This is expected because executing the same number of threads as the number of cores reduces

computational imbalance among the cores. Morphological Open and Distance Transform achieve the best performances with 120 and 240 threads, respectively. The performance of Morphological Open decreases after 120 threads because it is a memory-intensive operation. As is reported by other researchers (41; 6), the peak memory bandwidth of MIC for regular data access (STREAM benchmark (41)) is attained with 1 or 2 threads per core. If more cores are used in data intensive operations, the bandwidth decreases because of overload on the memory subsystem.

Distance Transform scales well until 240 threads are dispatched. The performance of this operation also depends on the memory bandwidth of the MIC. The difference is that this operation's data access pattern is irregular. To better understand the performance behavior of Distance Transform and because memory bandwidth figures for irregular data access are not provided in the hardware specifications, we created a micro-benchmark to measure the memory bandwidth of the MIC for random data accesses. The micro-benchmark program reads and writes data elements from/to random locations in a matrix in parallel. Locations to be accessed in the matrix are pre-computed and stored in a secondary array. Our experiments showed that the benchmark program scaled until 240 threads were executed. This finding correlates with the observed performance of Distance Transform.

The performances of the Scatter and Balanced thread affinity strategies were similar for both operations. The only significant discrepancy among the strategies occurred for Morphological Open and the Compact affinity strategy. This is because Compact utilizes the 60 physical cores only when 240 threads are dispatched. However, the Morphological Open achieves its best performance when 120 threads are created (2-way hyperthreading) with 60 cores.

5.2. Performance Analysis of Core Operations on MIC, GPU and CPU

These experiments compare the speedups of core operations on MIC, GPU and CPU. We computed the speedup values for an operation using the single-core CPU execution as the baseline performance. We used single-core implementations from well-known libraries and implementations published by other research groups (see Section 2.2). We have classified the core operations into three groups based on data access and computation patterns in order to better understand and characterize the speedup values obtained in the experiments. The groups are: (1) Operations with regular data access: RGB2Gray, Morphological Open, Color Deconvolution, Pixel Stats, Gradient Stats, and Sobel Edge; (2) Operations with irregular data access: Morphological Reconstruction, FillHoles, and Distance Transform; (3) Operations that heavily use atomic instructions: Area Threshold and Connected Component Labeling (CCL). We draw from the Roofline model proposed by Williams et al. (42) in our evaluation.

Regular Operations—The performance of an operation in this class heavily depends on the memory bandwidth of a computing device. We measured the memory bandwidths of the GPU, CPU and MIC using the STREAM benchmark (41) which are 148GB/s, 78GB/s (total bandwidth of two CPUs) and 160GB/s (one or two threads per core achieve similar performance), respectively. The peak computation capacities of the K20 GPU and MIC for

double precision numbers are about 1 TFLOPS each, while 2 CPUs together achieve 345 MFLOPS.

The Morphological Open, RGB2Gray and Color Deconvolution operations are memory bound with low arithmetic-instruction-to-byte ratio. Their speedups are low on the CPU because they quickly saturate the memory bus. Their performances on the GPU are nearly $1.25\times$ higher than on the MIC as is shown in Figure 6. The remaining regular operations (Pixel Stats, Gradient Stats and Sobel Edge) are compute bound due to their higher computation intensity. These operations achieve better scalability on all the devices. The GPU and the MIC result in similar performance improvements for these operations. Their execution times improve by about $1.9\times$ over the multi-core CPU executions. This group of compute intensive operations achieve the best performance on the MIC using 120 threads. Using more threads does not improve performance because the MIC threads can launch a vector instruction every two cycles.

Irregular Operations—The operations with irregular data access patterns include Morphological Reconstruction, FillHoles and Distance Transform. We measured the bandwidth of each device for random data accesses using the micro-benchmark program described in Section 5.1. The program executed 10 million random reads and writes on a $4K\times 4K$ matrix of integer values. The results are presented in Table 5. The GPU has much higher bandwidth as compared to the other devices. The performance of the MIC is low, especially for writes. This is a result of heavy data traffic on the ring bus connection in the MIC which is necessary to maintain consistency.

Given the micro-benchmark results, it is not surprising that Distance Transform is much faster (about two times faster) on the GPU than on the MIC or CPU as is shown in Figure 6. All phases of this operation perform only irregular data accesses with significantly more data reading than writing. Thus, the random data access bandwidths of the devices have significant impact on the operation's overall performance. Morphological Reconstruction and Fill Holes, on the other hand, have different performance behaviors than Distance Transform, because their data access patterns are not purely irregular. These operations have two computation phases. They first perform regular raster-/anti-raster passes on data. Then, they carry out irregular computations that use a queue to execute propagations. As such, the operations benefit more from the MIC processor, which is efficient in their regular phase. Additionally, they may execute the regular phase a number of times before the second phase. This value was tuned in the experiments for each device. The regular phase was executed a large number of times on the MIC, before the second phase was started. Hence, for these operations, the MIC compares better with the GPU; the GPU is only $1.33\times$ faster than the MIC.

Operations that Rely on Atomic Instructions—Atomic instructions are necessary for correct execution in some operations; Area Threshold and CCL intensively use atomic instructions. We developed another micro-benchmark to measure device bandwidths when atomic instructions are used. The benchmark program implements two scenarios. In the first scenario, a single variable is updated by all threads. This is a worst-case scenario because all threads access the same memory address and cause a lot of congestion on atomic instruction

executions. In the second scenario, the elements of an array are updated such that an element is updated by one thread only; however, a thread may update multiple array elements and multiple array elements may be updated concurrently. The second scenario does not cause congestion on atomic instruction executions, although each thread executes an atomic instruction when accessing an array element. Less congestion is expected to yield better performance on all the devices.

The results of the benchmark runs are presented in Table 6. The bandwidth of the GPU is the highest in both of the scenarios. However, the GPU is also the device that suffers the highest performance degradation from the second scenario to the first (worst-case) scenario. This large decrease in performance is a consequence of the fact that warps of threads on a GPU must execute the same instructions. It is expected that all threads executing in a warp will cause a conflict when updating the memory in the single variable case, and, consequently, threads will be serialized. This problem is intensified in the GPU because it supports a much higher level of concurrency than the other devices. The CPU is almost 2.4× faster than the MIC in both scenarios. The MIC relies on the use of vectorized instructions for good performance, since it is equipped with a simpler computing core. However, it does not support atomic vectorized instructions, such as those proposed by Kumar et. al. (43) for other multiprocessors.

As is expected from the benchmark results, Area Threshold and CCL achieve higher performance on the GPU than on the other two devices, with executions on the CPU being faster than those on the MIC.

5.3. Cooperative Execution on Hybrid Machines

Comparing Performances of Schedulers—This experiment compares the performance of proposed and baseline schedulers. The experiment uses 1 GPU, 1 MIC, and 14 CPU cores for computation; the remaining 2 CPU cores are used to manage the co-processors. The input data contains 800 4K×4K image tiles. Because the decision space may affect the performances of the schedulers, we varied the number of coarse-grained application stages that were concurrently executed in a Worker. This value is referred to as the window size. The window size was varied from 16 (the number of computing slots) to 70 – values larger than 70 had no significant effect on performance. The speedup and execution time estimates used by the schedulers were determined from previous runs.

The execution times of the image analysis application using FCFS, PADAS, HEFT and PAMS are shown in Figure 7. FCFS and PADAS result in similar execution times. Window size has little impact on their performance. HEFT results in the highest execution times with small window sizes, but performs better than FCFS and PADAS with window sizes larger than 45. HEFT performs better with larger decision spaces because it is able to balance well the work among processors and minimize miss-assignment of tasks. The PAMS scheduler achieves the best performance among all the schedulers for all window sizes. It is at least 1.15× faster than HEFT, but its gains are up to 1.39× for configurations with small window sizes. This scheduler is able to perform very well even with small window sizes and is less sensitive to window size. This is an important attribute of a scheduler, because smaller

window sizes in Workers reduces load imbalance among nodes on a distributed-memory machine.

To better understand the relative performance of the schedulers, we looked at the task scheduling computed by a scheduler with respect to the speedups of the operations on the devices and the percent contribution of the operations in the overall application execution time (Figure 6). The percentage of operations executed by each computing device for a given scheduling strategy is presented in Figure 8. The number of operations that FCFS assigns to each device is almost the same across all the operations. PADAS computes a scheduling in which operations are preferably assigned to certain types of processing devices. Nevertheless, it is not able to significantly improve on the performance of FCFS. PADAS fails when it assigns a substantial number of Distance Transform tasks to CPU cores although Distance Transform is more efficient on co-processors. The Sobel operation illustrates well what we believe to be the major problem with this policy. Sobel achieves the highest speedups on the GPU, but it is also efficient on the MIC. Because it is faster on the GPU, it is inserted in the queue from the GPU end. This significantly reduces the chances of the MIC getting selected for this operation. As a consequence, the MIC is used with other operations for which it is less efficient.

The analysis of the HEFT scheduling (Figure 8(c)) shows that most of the tasks with short execution times were executed by the CPU cores, regardless of the speedup values. HEFT also assigns all of the Distance Transform operations to the co-processors, but this operation is 2× more efficient on the GPU than on the MIC. Since no other device is able to efficiently perform this computation, the GPU should be used with higher percentages than what was observed in the experiments. Moreover, because this operation is responsible for nearly 40% of the whole application execution time, its correct scheduling is critical for high performance.

The profile of PAMS scheduling shows a better distribution of tasks. It schedules most of the Distance Transform operations on the GPU and utilizes the CPU cores for operations (such as Area Threshold and CCL) that do not attain high performance on the co-processors. It also schedules operations (such as ColorDevonv, FillHoles, Gradient, Pixel Stats and Sobel Edge) to the MIC that benefit from this device.

Evaluating Different Configurations of Hybrid System—This section investigates how well the scheduling strategies perform when different combinations of CPU and co-processors are used as the hybrid system. In this work we looked at three hardware configurations: CPU-GPU (15 CPU cores and 1 GPU), CPU-MIC (15 CPU cores and 1 MIC) and CPU-GPU-MIC (14 CPU cores, 1 GPU, and 1 MIC).

The experimental results presented in Figure 9 show that FCFS has the worst performance in all hardware configurations. In the scenarios with two device types (CPU-MIC and CPU-GPU), PADAS and PAMS result in nearly the same application execution times. This is expected because these schedulers will compute the same scheduling, since PAMS is reduced to PADAS with an implementation using two queues instead of one. PADAS and PAMS outperform HEFT in all of the configurations. The application executes faster in the

configuration with three devices with PAMS. The configuration with CPU-GPU is about $1.27\times$ faster than the one with CPU-MIC. This is because of the fact that neither CPU nor MIC is efficient with irregular patterns and atomic operations, which account for a large amount of the application execution time.

Sensitivity to Inaccurate Estimates—This section empirically evaluates the ability of the schedulers to perform well in cases in which estimations of the metrics (execution time and speedup) they use in computing a schedule are inaccurate. To carry out these experiments, we introduced errors in the execution times and, as a consequence, in the speedups in a controlled manner. The amount of error was varied as 0%, 10%, 25%, 50%, 75% and 100% of the execution times of the operations, with equal probability that the error was an increase or a decrease in execution time.

As presented in Figure 10, the performances of PADAS and PAMS are not significantly affected as the amount of error is varied. On the other hand, HEFT, which uses execution times to calculate schedules, is strongly impacted by the errors and is even less efficient than FCFS when the amount of error is higher than 50%. The experimental results show that the scheduling strategies based on the speedup metric are much less sensitive to inaccurate estimates and perform well even with large error in estimates. These schedulers use the input metric only to order tasks in a queue. The amount of error in estimated speedup has to be large enough to significantly alter the positions of tasks in the queue in order to adversely impact computed schedules and overall application performance.

Scalability Results—These experiments assessed the performance of the image analysis application on a distributed memory cluster. The scheduling strategies used the speedup values obtained in the experiments in Section 5.2 (see Figure 6). Figure 11 presents the performance results of both the CPU-only version of the application (16-CPU), which uses only the 16 CPU cores on each node, and the hybrid version, which uses the CPU, GPU and MIC in a coordinated manner via FCFS, HEFT or PAMS. The input data for this experiment consists of 6,379 4K $\times$ 4K image tiles. Both versions of the application achieved good scalability with efficiency higher than 90%. The hybrid version with PAMS results in a performance improvement of about $2.2\times$ on top of the CPU-only version. The performance improvement by PAMS on top of HEFT is maintained as the application scales.

We also executed the application on a large number of computing nodes equipped with two CPUs and one MIC only, since Stampede has more nodes with this configuration. We experimented with five versions of the application: CPU-only refers to the multi-core CPU version that uses all the CPU cores available but no co-processors; MIC-only uses only the MIC in each node to perform computations; CPU-MIC uses all the CPU cores and the MIC in coordination and distributes tasks to devices in each node via FCFS, PAMS, or HEFT as the scheduling strategy.

Figure 12 shows the weak scaling performance of the application when the dataset size and the number of nodes increase proportionally. For the configuration with 192 nodes, the input data contains 68,284 4K $\times$ 4K image tiles for a total of 3.27 TB of uncompressed data. The application scaled well with about 84% of efficiency on 192 nodes. The main factor limiting

scalability is the increasing cost of reading the input image tiles concurrently from storage as the number of nodes (and processes) grows. The CPU-MIC versions resulted in improvements of up to 2.06× on top of the MIC-only version. The CPU-MIC PAMS version results in the lowest execution times and is on average 1.29× faster than the CPU-MIC FCFS version.

6. Related Work

Cooperative use of CPU and GPU has been target of a number of projects (33; 34; 35; 36; 37; 38; 39; 40; 44; 45; 46; 47; 48; 49; 50; 51; 52). The Mars (33) and Merge (34) systems are seminal works, which target the acceleration of MapReduce applications. Mars has evaluated the static partition strategies for dividing MapReduce tasks between CPU and GPU. Merge, on the other hand, proposed the use of a dynamic work partition scheme during application execution. Qilin (35) presented an automatic approach for distributing tasks for computation in hybrid systems. PTask (46) provides OS abstractions for task based applications on GPU equipped systems. These systems are designed for shared memory machines.

Runtime systems for execution on distributed memory systems have also been developed (38; 39; 50; 44; 45; 47; 48). DAGuE (38) and StarPU (36; 48) have been successfully applied in regular linear algebra applications. These systems describe applications as directed acyclic graphs of operations and implement multiple scheduling policies, e.g., including those that prioritize computation of critical paths in the dependency graph in order to maximize parallelism. Ravi (39) and Hartley (50) proposed runtime techniques to auto-tune work partitioning among CPUs and GPUs. OmpSs (45) supports execution of dataflow applications created via compilation of annotated codes. The performance-aware scheduling strategies employed proposed in this work use a different priority metric (relative performance or speedup) of operations on CPUs or accelerators for mapping of work to processors. As presented, it results into significant performance improvement on top of HEFT, which is widely used in previous works and is known as a very efficient scheduling policy.

The release of the new Intel Xeon Phi has motivated a number of research efforts (2; 3; 4; 5; 6). Hamidouche et al.(53) have developed a very interesting platform that allows for off-loading computations to a MIC on a remote machine. In this way, applications running on multiple nodes may share a MIC to maximize its utilization. The use of OpenMP to develop code targeting MICs has also been studied (4). The authors analyzed the overheads of important OpenMP directives for thread creation and synchronization. They also evaluated the processor performance with respect to memory bandwidth. Saule et al. (6) implemented optimized sparse matrix multiplication kernels for MICs and provided a comparison of MICs and GPUs for this operation.

In our work, we perform a comparative performance evaluation of MICs, multi-core CPUs, and GPUs using an important class of operations. These operations employ diverse computation and data access patterns and several parallelization strategies. The comparative performance analysis correlates the performance of operations with co-processors characteristics using co-processor specifications or performance measured using micro-

kernels. This evaluation provides a methodology and insights towards a better understanding of the efficacy of co-processors for a class of applications. We also investigate coordinated use of MICs, GPUs and CPUs on a distributed memory machine and its impact on application performance. Our approach takes into account performance variability of operations to make smart task assignments.

7. Conclusions

Development of application optimizations to maximize performance benefits from hybrid systems is a very challenging problem. Besides optimizations for load balancing and minimizing data transfers between devices, device hardware and execution characteristics need to be taken into account. With increasing availability of high-end machines with multiple co-processors with different architectures, application developers have to understand the effects of different devices on the performance of their applications. However, little information about the relative performance of modern co-processors for different computation patterns is available.

In this work, we aim to address this issue by conducting a systematic performance comparison of CPUs, GPUs and MICs using a set of operations with diverse data access patterns (regular and irregular), computation intensity, and types of parallelization strategies from an important class of applications. We also propose and evaluate new performance-aware scheduling techniques for efficiently executing complex applications in systems equipped with CPUs and multiple accelerators. Our scheduling techniques take into account performance variabilities across computation tasks or operations to better utilize hybrid systems.

The experimental results show that different types of co-processors are more appropriate for specific data access patterns and types of parallelism, as expected. The MIC's performance compares well with that of the GPU when regular operations and computation patterns are used. However, the GPU is significantly more efficient for operations that perform irregular data access and heavily employ atomic operations. Our work also shows that a strong performance variability exists among different operations, as a result of their computation patterns. This performance variability needs to be taken into account to efficiently execute pipelines of operations using co-processors and CPUs in coordination.

We compare the proposed scheduling strategies (PADAS and PAMS) with two well known scheduling policies: FCFS and HEFT on a distributed-memory cluster of nodes with multiple types of co-processors. The results show that PAMS achieves better performance than other policies, being at least $1.15\times$ faster than its closest competitor: HEFT. We also experimentally show that performance-aware strategies perform well even when a high level of inaccuracy exists in the input information (speedups) used to compute the scheduling. On the other hand, the HEFT time based scheduler suffers dramatically as the estimation error in the execution times of operations increases. The example application achieves good scalability in distributed memory settings via co-operative use of CPUs and co-processors.

Acknowledgments

This work was supported in part by HHSN261200800001E and 1U24CA180924-01A1 from the NCI, R24HL085343 from the NHLBI, R01LM011119-01 and R01LM009239 from the NLM, RC4MD005964 from the NIH, National Institutes of Health (NIH) K25CA181503 and by CNPq, CAPES, FINEP, Fapemig, and INWEB. This research used resources provided by the XSEDE Science Gateways program and the Keeneland Computing Facility at the Georgia Institute of Technology, which is supported by the NSF under Contract OCI-0910735.

References

1. NVIDIA. GPU Accelerated Applications. 2012
2. Joó B, Kalamkar D, Vaidyanathan K, Smelyanskiy M, Pamnany K, Lee V, et al. Lattice QCD on Intel Xeon Phi Coprocessors. *Supercomputing*. 2013; 7905:40–54.LNCS
3. Heinecke A, Vaidyanathan K, Smelyanskiy M, et al. Design and Implementation of the Linpack Benchmark for Single and Multi-node Systems Based on Intel Xeon Phi Coprocessor. *The 27th IEEE Int. Symposium on Parallel Distributed Processing*. 2013
4. Cramer T, Schmidl D, Klemm M, an Mey D. OpenMP Programming on Intel Xeon Phi Coprocessors: An Early Performance Comparison. *Many-core Applications Research Community Symposium*. 2012
5. Eisenlohr J, Hudak DE, Tomko K, Prince TC. Dense Linear Algebra Factorization in OpenMP and Cilk Plus on Intel's MIC Architecture: Development Experiences and Performance Analysis. *TACC-Intel Highly Parallel Computing Symp*. 2012
6. Saule E, Kaya K, Çatalyürek ÜV. Performance Evaluation of Sparse Matrix Multiplication Kernels on Intel Xeon Phi. *CoRR*. 2013 abs/1302.1078.
7. Sharp R, Ridgway R, Iyengar S, Gulacy A, Wenzel P, de Bruin A, et al. Registration and 3D Visualization of Large Microscopy Images. *Proceedings of the SPIE Annual Medical Imaging Meetings*. 2006:923–934.
8. Cooper LA, Kong J, Gutman DA, Wang F, Gao J, Appin C, et al. Integrated morphologic analysis for the identification and characterization of disease subtypes. *J Am Med Inform Assoc*. 2012; (2): 317–323.
9. Wan Z, Dozier J. A generalized split-window algorithm for retrieving land-surface temperature from space. *IEEE Transactions on Geoscience and Remote Sensing*. 1996; 34(4):892–905.
10. Chandola V, Vatsavai RR. A scalable gaussian process analysis algorithm for biomass monitoring. *Stat Anal Data Min*. 2011 Aug; 4(4):430–445.
11. Vatsavai R, Bhaduri B. A hybrid classification scheme for mining multisource geospatial data. *GeoInformatica*. 2011; 15(1):29–47.
12. Parashar M, Muralidhar R, Lee W, Arnold DC, Dongarra J, Wheeler MF. Enabling interactive and collaborative oil reservoir simulations on the Grid. *Concurrency – Practice and Experience*. 2005; 17:1387–1414.
13. Saltz J, Teodoro G, Pan T, Cooper L, Kong J, Klasky S, et al. Feature-based analysis of large-scale spatiotemporal sensor data on hybrid architectures. *International Journal of High Performance Computing Applications*. 2013; 27(3):263–272.
14. Cooper LAD, Kong J, Gutman DA, Wang F, Cholleti SR, Pan TC, et al. An Integrative Approach for In Silico Glioma Research. *IEEE Transactions on Biomedical Engineering*. 2010 Oct.
15. Kong J, Cooper LA, Wang F, Gao J, Teodoro G, Scarpace L, et al. Machine-Based Morphologic Analysis of Glioblastoma Using Whole-Slide Pathology Images Uncovers Clinically Relevant Molecular Correlates. *PloS one*. 2013; 8(11):e81049. [PubMed: 24236209]
16. Vatsavai R, Tuttle M, Bhaduri B, Bright E, Cheriyaat A, Chandola V, et al. Rapid Damage Assessment Using High-Resolution Remote Sensing Imagery: Tools and Techniques. 2011 *IEEE International Geoscience and Remote Sensing Symposium (Igarss)*. 2011:1445–1448.
17. Vatsavai R, Bhaduri B. A hybrid classification scheme for mining multisource geospatial data. *Geoinformatica*. 2011; 15:29–47.
18. Vatsavai R, Shekhar S, Bhaduri BL. A Learning Scheme for Recognizing Sub-classes from Model Trained on Aggregate Classes. *SSPR/SPR*. 2008:967–976.

19. Upadhyay G, Ray SS, Panigrahy S. Derivation of crop phenological parameters using multi-date SPOT-VGT-NDVI data: A case study for Punjab. *Journal of the Indian Society of Remote Sensing*. 2008; 36:37–50.
20. Showalter PS. Remote sensing in disaster research: A review. *Disaster Prevention and Management*. 2001; 10:21–29.
21. Teodoro, G., Kurc, T., Kong, J., Cooper, L., Saltz, J. Comparative Performance Analysis of Intel (R) Xeon Phi (TM), GPU, and CPU: A Case Study from Microscopy Image Analysis; 2014 IEEE 28th International Parallel and Distributed Processing Symposium. 2014. p. 1063-1072. IPDPS '14 Available from: <http://dx.doi.org/10.1109/IPDPS.2014.111>
22. Cooper L, Kong J, Gutman D, Wang F, et al. An Integrative Approach for In Silico Glioma Research. *IEEE Trans Biomed Eng*. 2010; 57(10):2617–2621. [PubMed: 20656651]
23. Kong, J., Wang, F., Teodoro, G., Cooper, LAD., Moreno, CS., Kurç, TM., et al. High-performance computational analysis of glioblastoma pathology images with database support identifies molecular and survival correlates. 2013 IEEE International Conference on Bioinformatics and Biomedicine; Shanghai, China. December 18–21, 2013; 2013. p. 229-236.
24. Bradski G. The OpenCV Library Dr Dobb's Journal of Software Tools. 2000
25. Teodoro G, Pan T, Kurc T, Kong J, Cooper L, Saltz J. Efficient Irregular Wavefront Propagation Algorithms on Hybrid CPU-GPU Machines. *Parallel Computing*. 2013
26. Vincent L. Morphological grayscale reconstruction in image analysis: Applications and efficient algorithms. *IEEE Transactions on Image Processing*. 1993; 2:176–201. [PubMed: 18296207]
27. Oliveira VMA, de Alencar Lotufo R. A Study on Connected Components Labeling algorithms using GPUs. *SIBGRAPI*. 2010
28. Ruifrok AC, Johnston DA. Quantification of histochemical staining by color deconvolution. *Analytical and Quantitative Cytology and Histology*. 2001 Aug; 23(4):291–299. [PubMed: 11531144]
29. Dean J, Ghemawat S. MapReduce: Simplified Data Processing on Large Clusters. *The Sixth Symp. on Operating System Design and Implementation*. 2004:137–150.
30. Tarjan RE. Efficiency of a Good But Not Linear Set Union Algorithm. *J ACM*. 1975 Apr; 22(2): 215–225.
31. Teodoro G, Pan T, Kurc TM, Kong J, Cooper LAD, Podhorszki N, et al. High-throughput Analysis of Large Microscopy Image Datasets on CPU-GPU Cluster Platforms. IPDPS '13: Proceedings of the 2013 IEEE International Symposium on Parallel and Distributed Processing. 2013
32. Teodoro G, Kurc T, Pan T, Cooper L, Kong J, Widener P, et al. Accelerating Large Scale Image Analyses on Parallel, CPU-GPU Equipped Systems. 26th IEEE Int. Parallel and Distributed Processing Symp. 2012:1093–1104.
33. He B, Fang W, Luo Q, Govindaraju NK, Wang T. Mars: A MapReduce Framework on Graphics Processors. *Parallel Architectures and Compilation Techniques*. 2008
34. Linderman MD, Collins JD, Wang H, Meng TH. Merge: a programming model for heterogeneous multi-core systems. *SIGPLAN Not*. 2008; 43(3):287–296.
35. Luk CK, Hong S, Kim H. Qilin: Exploiting Parallelism on Heterogeneous Multiprocessors with Adaptive Mapping. 42nd Int. Symp. on Microarchitecture. 2009
36. Augonnet C, Thibault S, Namyst R, Wacrenier PA. StarPU: A Unified Platform for Task Scheduling on Heterogeneous Multicore Architectures. *Euro-Par '09: Proceedings of the 15th International Euro-Par Conference on Parallel Processing*. 2009:863–874.
37. Damos GF, Yalamanchili S. Harmony: an execution model and runtime for heterogeneous many core systems. *Proc. of the 17th International Symposium on High-Performance Distributed Computing (HPDC)*. 2008
38. Bosilca G, Bouteiller A, Herault T, Lemerminier P, Saengpatana NO, Tomov S, et al. Performance Portability of a GPU Enabled Factorization with the DAGuE Framework. *IEEE Int. Conf. on Cluster Comp*. 2011:395–402.
39. Ravi VT, Ma W, Chiu D, Agrawal G. Compiler and runtime support for enabling generalized reduction computations on heterogeneous parallel configurations. *The 24th ACM Int. Conf. on Supercomputing*. 2010:137–146.

40. Huo X, Ravi VT, Agrawal G. Porting irregular reductions on heterogeneous CPU-GPU configurations. 18th International Conference on High Performance Computing (HiPC). 2011:1–10.
41. McCalpin JD. Memory Bandwidth and Machine Balance in Current High Performance Computers. IEEE Computer Society Technical Committee on Computer Architecture Newsletter. 1995 Dec.: 19–25.
42. Williams S, Waterman A, Patterson D. Roofline: an insightful visual performance model for multicore architectures. Commun ACM. 2009 Apr; 52(4):65–76.
43. Kumar S, Kim D, Smelyanskiy M, Chen YK, Chhugani J, et al. Atomic Vector Operations on Chip Multiprocessors. 35th Int. Symp. on Computer Architecture. 2008
44. Teodoro G, Hartley T, Catalyurek U, Ferreira R. Run-time optimizations for replicated dataflows on heterogeneous environments. The 19th ACM International Symposium on High Performance Distributed Computing. 2010:13–24.
45. Bueno J, Planas J, Duran A, Badia RM, Martorell X, Ayguade E, et al. Productive Programming of GPU Clusters with OmpSs. 2012 IEEE 26th Int. Parallel Distributed Processing Symp. (IPDPS). 2012:557–568.
46. Rossbach CJ, Currey J, Silberstein M, Ray B, Witchel E. PTask: operating system abstractions to manage GPUs as compute devices. The 23rd ACM Symposium on Operating Systems Principles. 2011:233–248.
47. Teodoro G, Hartley TR, Catalyurek U, Ferreira R. Optimizing dataflow applications on heterogeneous environments. Cluster Computing. 2012; 15:125–144.
48. Augonnet C, Aumage O, Furmento N, Namyst R, Thibault S. StarPU-MPI: Task Programming over Clusters of Machines Enhanced with Accelerators. The 19th European MPI Users' Group Meeting (EuroMPI 2012). 2012
49. Teodoro G, Sachetto R, Sertel O, Gurcan M, M W Jr, Catalyurek U, et al. Coordinating the Use of GPU and CPU for Improving Performance of Compute Intensive Applications. IEEE Cluster. 2009:1–10.
50. Hartley T, Saule E, Catalyurek U. Automatic dataflow application tuning for heterogeneous systems. HiPC. IEEE. 2010:1–10.
51. Hartley T, Catalyurek U, Ruiz A, Ujaldon M, Igual F, Mayo R. Biomedical Image Analysis on a Cooperative Cluster of GPUs and Multicores. 22nd ACM Intl. Conference on Supercomputing. 2008
52. Teodoro, G., Valle, E., Mariano, N., Torres, R., Wagner, Meira J., Saltz, J. Approximate similarity search for online multimedia services on distributed CPU&GPU platforms; The VLDB Journal. 2013. p. 1-22. Available from: <http://dx.doi.org/10.1007/s00778-013-0329-7>
53. Hamidouche K, Potluri S, Subramoni H, Kandalla K, Panda DK. MIC-RO: enabling efficient remote offload on heterogeneous many integrated core (MIC) clusters with InfiniBand. 27th Int. ACM International Conference on Supercomputing. 2013:399–408. ICS '13

Fig. 1.

Detailed representation of the Segmentation and Feature Computation stages of our example microscopy image analysis application. These stages are composed of several compute intensive finer-grain operations, which may be naturally represented as a hierarchical computation dataflow.

Fig. 2.
Runtime system architecture

Fig. 5. Scalability of Morphological Open and Distance Transform under three thread affinity strategies.

Fig. 6.

Speedups of the core operations on multi-core CPU, MIC and GPU, using the single-core CPU versions as baseline. The number above each dash is the number of threads that led to the best performance on the MIC. The number above the graph in each column is the percent contribution of the corresponding operation to the overall application execution time.

Fig. 7.
Performance of the scheduling strategies when the number of concurrent tasks assigned to a Worker (window size) is varied.

Fig. 8.

Task assignment profile for each of the operations executed in the image analysis application under different scheduling strategies. A fixed window size of 80 is used in all runs.

Fig. 9.

Application performance when application tasks are scheduled and mapped to two or more computation devices.

Fig. 10.

Performance of scheduling strategies as errors are introduced in estimated execution times of operations.

Fig. 11.

Strong-scaling results on a distributed-memory cluster in which each node has two 8-core CPUs, one MIC, and one GPU.

Fig. 12.

Weak scaling results on a distributed-memory cluster in which each node has two 8-core CPUs and one MIC.

Table 1

Operation Categories

Operation Category	Microscopy Image Analysis	Weather Prediction	Monitoring and Change Analysis
Data Cleaning and Low Level Transformations	Color normalization. Thresholding of pixel and regional gray scale values.	Remove anomalous measurements and convert spectral intensities to values of interest	Remove unusual readings. Convert signal intensities to color and other values of interest.
Data Subsetting, Filtering, and Subsampling	Selection of regions within an image. Thresholding of pixel values.	Spatial selection/cross match to find portion of a dataset corresponding to a given geographic region.	Spatial selection/cross match to find portion of a dataset corresponding to a given geographic region.
Object Segmentation	Segmentation of nuclei and cells.	Segmentation of regions with similar land surface temperature.	Segmentation of buildings, trees, plants, etc.
Feature Computation	Compute texture and shape features for each cell.	Compute texture and shape features for each region.	Compute texture and shape features for objects.
Spatio-temporal Mapping and Registration	Deformable registration of images to anatomical atlas.	Generation of mosaic of tiles to get complete coverage.	Registering low and high resolution images corresponding to same regions.
Classification	Clustering of nuclei and/or images into groups.	Classification of segmented regions.	Classification of buildings, trees, plants.
Aggregation	Aggregation of object features for per image features.	Time-series calculations on changing land and air conditions.	Aggregation of labeled buildings, trees, plants into residential, industrial, vegetation areas.
Change Detection and Comparison	Spatial queries to compare segmented nuclei and features within and across images.	Spatial and temporal queries on classified regions and aggregation to look for changing weather patterns.	Characterize vegetation changes over time and area.

Table 2

Data access patterns and computational complexity of common operation categories.

Operation Category	Data Access Pattern and Computational Complexity
Data Cleaning and Low Level Transformations	Mainly local and regular data access patterns. Moderate computational complexity.
Data Subsetting, Filtering, and Subsampling	Local data access patterns as well as indexed access. Low to moderate, mainly data intensive computations.
Object Segmentation	Regular and Irregular, but primarily local, data access patterns. High computational complexity.
Feature Computation	Regular and Irregular, but primarily local, data access patterns. High computational complexity.
Spatio-temporal Mapping and Registration	Irregular local and global data access patterns. Moderate to high computational complexity.
Classification	Irregular and global data access patterns. High computational complexity.
Aggregation	Primarily local with a crucial global component for aggregation. Moderate/high computation complexity.
Change Detection and Comparison	Compute and data-intensive computations. Mixture of local and global data access patterns as well as indexed access.

Table 3

Segmentation and feature computation operations from the microscopy image analysis application.

Operations	Description	Data Access Pattern	Computation	Parallelism	CPU Source
Segmentation Stage					
Covert RGB to grayscale	Covert RGB image into grayscale intensity image	Regular, multi-channel local	Moderate	Data	OpenCV
Morphological Open	Remove small objects and fills holes in foreground	Regular, neighborhood (13×13 disk)	Low	Data	OpenCV
Morphological Reconstruction	Flood-fill a marker image that is limited by a mask image.	Irregular, neighborhood (4-/8-connected)	Low	IWPP (25)	Vincent (26)
Area Threshold	Remove small or large objects	Mixed, neighborhood	Low	Reduction	Implemented
FillHoles	Fill holes in objects using a flood-fill	Irregular, neighborhood (4-/8-connected)	Low	IWPP (25)	Vincent (26)
Distance Transform	Compute distance to closest background pixel for each foreground pixel	Irregular, neighborhood (8-connected)	Moderate	IWPP (25)	Implemented
Component Labeling	Label with the same value each component of the mask	Irregular, global	Low	Union-find	Oliveira (27)
Feature Computation Stage					
Color Deconvolution (28)	Separate multi-stained image into different channels	Regular, multi-channel local	Moderate	Data	Implemented
Pixel Statistics	Compute intensity stats (mean, median, max, etc) per object	Regular, access a set of bounding-boxed areas	High	Object	Implemented
Gradient Statistics	Calculate x,y gradient and derive per object stats	Regular, neighborhood and bounding-boxed areas	High	Object	Implemented OpenCV(Grad)
Sobel Edge	Compute Sobel and derive per object stats	Regular, access a set of bounding-boxed areas	High	Object	Implemented OpenCV(Sobel)

Table 4

Hardware specifications of the MIC and GPU used in the performance experiments.

	Intel Xeon Phi SE10P	NVIDIA K20 GPU
# of cores	61	2496
Clock	1.09 GHz	0.706 GHz
L1 Cache	32 KB	64 KB
L2 Cache	512 KB	768 KB
Memory	8 GB	5 GB
Bandwidth	352 GB/s	205 GB/s

Table 5

Device bandwidths for random data accesses (MB/s).

	CPU	MIC	GPU
Reading	305	399	895
Writing	74	16	126

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

Table 6

Measured device bandwidths for the Atomic Add operation (Millions/sec).

	CPU	MIC	GPU
Single Variable	134	55	693
Array	2,200	906	38,630