

INVESTIGATION OF PARALLEL DATA PROCESSING USING HYBRID HIGH PERFORMANCE CPU+GPU SYSTEMS AND CUDA STREAMS

Paweł Czarnul

*Faculty of Electronics, Telecommunications and Informatics
Gdansk University of Technology, Narutowicza 11/12, 80-233
Gdansk, Poland*

e-mail: pczarnul@eti.pg.edu.pl

Abstract. The paper investigates parallel data processing in a hybrid CPU+GPU(s) system using multiple CUDA streams for overlapping communication and computations. This is crucial for efficient processing of data, in particular incoming data stream processing that would naturally be forwarded using multiple CUDA streams to GPUs. Performance is evaluated for various compute time to host-device communication time ratios, numbers of CUDA streams, for various numbers of threads managing computations on GPUs. Tests also reveal benefits of using CUDA MPS for overlapping communication and computations when using multiple processes. Furthermore, using standard memory allocation on a GPU and Unified Memory versions are compared, the latter including programmer added prefetching. Performance of a hybrid CPU+GPU version as well as scaling across multiple GPUs are demonstrated showing good speed-ups of the approach. Finally, performance per power consumption of selected configurations are presented for various numbers of streams and various relative performances of GPUs and CPUs.

Keywords: GPGPU, overlapping computations and communication, MPS, Unified Memory, performance, power consumption

Mathematics Subject Classification 2010: 68M20, 65Y05, 68N15

1 INTRODUCTION

In today's high performance computing (HPC) systems, several computing devices are typically used – multi- and many-core CPUs, GPUs, FPGAs. All have their advantages

and disadvantages depending on particular types of codes and applications [9]. Most of HPC systems nowadays feature either traditional multicore CPU + accelerator (GPU, Intel Xeon Phi x100) or manycore CPUs (such as Intel Xeon Phi x200 or Sunway manycore CPUs in the Sunway TaihuLight cluster). Selected application examples of applications running on such systems include data encryption and decryption algorithms [32], pattern matching for deep packet inspection [26], RNA secondary structure prediction [27], parallel implementation for a DVB-RCS2 receiver [46], parallelization of large vector similarity computations [14, 11], stitching large scale optical microscopy images [4] etc. For this reason, efficient management of computations among these processors is a key to achieving high throughput, especially for incoming data streams that must be processed under time constraints. GPGPU has become very popular for processing large data sets in the Single Instruction Multiple Threads fashion. As long as processing in threads does not result in too much divergence, one can achieve very high processing throughput. Especially important is also fast delivering of input data from host memory to GPU memory and results back from GPU memory to the host. This can be achieved through overlapping communication and GPU and CPU computations by using multiple streams. This topic is investigated in this paper in detail, in terms of performance for various numbers of streams, threads managing computations in a CPU+GPU setting, using the standard GPU memory management and Unified Memory [34] approaches. Furthermore, as today's HPC is not only about performance, power consumption is also considered, in the context of performance to power consumption ratio of various configurations.

The approach adopted in this paper includes analysis based on a custom built benchmark, described in Section 3, that assumes input data that is composed of multiple data chunks which are fed into CUDA streams to GPUs or processed on multicore host CPUs. The benchmark allows for various compute time to host-device communication time ratios, numbers of streams and threads managing computations and communication and thus, depending on the values of parameters, can be regarded as a template representative of many real world applications.

The objective of this work is to assess performance and selected performance/power characteristics of parallel processing of a data stream which is passed for computations to either GPUs using CUDA streams or to GPUs and CPU cores in a hybrid CPU+GPU approach. The contribution includes assessment of preferred numbers of streams for various GPU architectures, preferred application architecture in terms of the number of host GPU management and computing threads, assessment of performance differences between standard memory management, Unified Memory and Unified Memory with prefetching, all for various compute to communication ratios. Additionally, performance per power consumption is evaluated for selected configurations. Furthermore, scaling from 1 to 4 NVIDIA Tesla V100 GPUs of DGX Station installed at the Faculty of ETI, Gdansk University of Technology is presented.

The outline of the paper is as follows. Section 2 presents existing related work and contributions of this paper in that context, Section 3 the processing model and design of the benchmark used for experiments, Section 4 tests and results including testbed systems, impact of multiple streams on performance using various numbers of threads managing computations, launching computations from multiple processes with and without MPS,

performance with and without Unified Memory, scalability of hybrid CPU+GPU code, scaling across multiple GPUs and performance-power consumption ratios for hybrid configurations. Finally Section 5 presents conclusions and future work.

2 RELATED WORK

2.1 Mechanisms for data management in selected GPU-aware parallel programming APIs

Overlapping computation on the GPU, CPU as well as CPU-GPU and GPU-CPU communication is a well known technique that allows to minimize execution time of an application using GPUs [14, 31, 13, 25]. This approach can be used for both batch processing if the data is already available when the application starts or is incoming to a node in possibly many data streams.

In CUDA, kernel functions are executed in parallel on a GPU by a grid which is composed of thread blocks each of which consists of a number of threads. Blocks within a grid and threads within a block can be lined up in 1, 2 or 3 dimensions. Various operations (out of host-to-device communication, device-to-host communication, kernel execution) submitted to two different streams can potentially be overlapped in H2D, compute and D2H queues. Thus, a larger number of streams can potentially allow better overlapping (so-called n -way in the case of n streams [41]) if there is potential for that in the application and if the GPU and the driver support that. Potentially kernels can also be executed in parallel, depending on their requirements and the GPU. Unified Memory allows allocation and access to data from the host and device sides and page migration, transparent to the user. The contribution of the paper is how a particular configuration (with a given number of streams) for a given GPU (GPUs of various architectures were used) benefits which is otherwise very difficult to predict given these factors.

It should be noted that OpenCL offers a similar programming model to CUDA but targeting systems with both GPUs as well as CPUs [13, 15]. Specifically, a kernel can be executed on a compute device by a structure called NDRange that consists of work groups which in turn consist of work items. Both work groups within the NDRange and work items within a work group can be lined up in 1, 2 or 3 dimensions. A kernel is executed by work items in parallel within a context that is associated with one or more devices. Input and output data are managed through memory objects. Overlapping can be achieved using command queues, similarly to using streams in CUDA. OpenCL version 2.0+ allows to use Shared Virtual Memory which allows codes running on the host and a device to share data. Various modes including coarse-grained or fine-grained with the possibility of accessing locations concurrently if SVM atomic operations are supported. Another high level API allowing to use GPUs in a way similar to OpenMP is OpenACC [13, 15]. OpenACC allows to use directives for instructing parallelization of code regions, specifically loops as well as scoping of data and synchronization. Data related directives allow to specify allocation, releasing memory and rely on the concept of reference counters to data.



Assessment of benefits and performance of Unified Memory was done previously in [22] but for batch type input data for applications such as verification of Goldbach's conjecture, 2D heat transfer analysis and adaptive numerical integration. That research was then extended with evaluation of not only the basic Unified Memory code against the standard approach but also Unified Memory with prefetching [23]. Results were presented for four applications: Sobel and image rotation filters as well as stream image processing and computational fluid dynamic simulation. Tests were performed on Pascal and Volta architecture GPUs, specifically NVIDIA GTX 1080 and NVIDIA V100 cards. Furthermore, evaluation of Unified Memory oversubscription over the standard manual management approach was provided, generally showing slight benefits of the latter, if implemented efficiently. In those contexts, the contribution of this paper is assessment of impact of the number of streams with Unified Memory, assessment of NVIDIA MPS's performance and consideration of power consumption with the number of streams in parallel processing with CUDA.

2.2 Selected works on efficiency of using multiple streams using GPUs

There are studies in the literature on efficiency of using multiple streams using GPUs. For instance, paper [20] investigates the impact of using various numbers of streams on performance of such an application with a theoretical formula for the best number of streams. It was considered in terms of the number of iterations of a loop within a kernel. Tests were performed for GTX 280 and GTX 480 cards which are not used widely anymore. GPU architectures have also changed considerably since then. In paper [12], the author analyzed and compared performance of processing on a GPU using 1, 2 and 4 streams for modern GPUs: mobile NVIDIA GeForce 940MX, desktop GTX 1060, server Tesla K20m and Tesla V100. Tests were performed for various compute time to host-device communication time ratios proving large benefits of using 2 or 4 streams for overlapping communication and computations and showing relative performances of the tested GPUs. Compared to [20] this paper contributes by analysis on newer GPUs, consideration of Unified Memory approach and performance to power consumption analysis. Compared to [12] this paper brings testing using more streams, multi-threaded and single-threaded applications, MPS as well as performance to power consumption considerations. Apart from multiple streams, concurrent kernel execution is also possible on GPUs. Paper [45] investigates approaches such as context switching, manual context funneling and automatic CUDA context funneling but tests were performed on older CUDA 4 and earlier versions and demonstrated that automatic CUDA context funneling (sharing a context among process threads) is very efficient. Work [29] proposes a detailed computation-bound single kernel performance model for understanding the resource scheduling system with CUDA streams and focuses on multi-kernel concurrency. Similarly, paper [8] investigates conditions needed for concurrent execution of kernels simultaneously.

2.3 Using multiple streams for various applications

Deployment of multi-stream processing for GPU based systems for particular applications has been analyzed in the literature. In paper [43] authors focus on performance improvement through more effective overlapping of communication and computations using OpenMP as well as multiple CUDA threads. Many threads control each GPU and the authors have launched 4 CUDA streams for each pair of neighboring GPUs to overlap communication and computation of inner domain points. A 3D stencil use case was used to demonstrate benefits over previous solutions. Tests were performed on Kepler and Fermi cards. Compared to that work, this paper considers a model with independent input data chunks rather than geometric Single Program Multiple Data paradigm [13], considers more streams, Unified Memory and power consumption for a more recent Pascal card. In paper [19] authors focus on improvement of performance of Sparse matrix-vector multiplication (SpMV) code using many GPUs installed within a node. Optimization is performed using multiple OpenMP threads that control particular GPUs as well as multiple CUDA streams for overlapping. Benefits of such improved approach using 2 GPUs are shown against a naive 1 GPU system implementation for a variety of sparse matrices. Compared to that approach, this paper considers hybrid CPU+GPU processing, investigates multiple streams, Unified Memory and performance to power consumption ratios. Paper [35] proposes a multi-stream implementation of stereo disparity estimation and anaglyph video frame generation using GPUs. Specifically, multiple threads are started using Pthreads, each of which manages a certain number of streams. Performance is presented for a thread count between 1 and 8 and the number of streams between 1 and 8 showing considerable speed-ups of the solution with 100 frames per second for 1024x1024 color images. GeForce GTX780 cards were used for experiments. Paper [38] contributes by proposal of a parallel CPU+GPU code for image formation in scanning transmission electron microscopy. Similarly to this work, an algorithm for parallelization using multicore CPUs and GPUs are provided, with assessment of benefits from using multiple CUDA streams. In that context, this paper contributes by analysis of various numbers of streams, Unified Memory and performance to power consumption ratios for similar computations. Utilization of CUDA streams for parallel implementation of a genetic algorithm is presented in paper [39]. Data stream processing accelerated using GPUs in the context of DBMSes is discussed in [36] for data representation better matching the GPU architecture. Similarly, this paper contributes by consideration of various stream and thread CPU+GPU configurations, Unified Memory and performance to power consumption ratios.

2.4 Selected frameworks and environments for processing data using GPUs

Paper [24] provides analysis of programming environments for processing large amounts of data efficiently. Specifically, the work investigates programmability vs performance such that programs can increase their performance at the cost of decreasing programmability. Java and Stream API, C/C++ and OpenMP, C/C++ and CUDA (with and with-



out CUDA streams) are compared. Power-aware computations for data processing is also an important research topic considered today [16]. There exist frameworks that provide higher than OpenMP, CUDA and MPI programming abstractions to processing data streams using GPUs, good performance and relatively easy-to-use programming models. Available solutions for data streaming include, in particular, Spark [47], Storm [28, 21], Storm working in a geographically distributed and highly variable environment [6], FastFlow [2], extension of FastFlow for a network of multi-core workstations [1], Flink [5, 18], PiCo [33], Thrill [3]. Paper [48] describes GStream that is a scalable framework suited for a cluster of GPUs with GStream API over CUDA, Pthreads and MPI. It is demonstrated for benchmarks such as FIR, MM, FFT, IS and LAMMPS that it offers very good speed-ups, only slightly worse than raw CUDA. For this and the following high level approaches, the contributions of this paper can be used for improvement of performance of lower level building blocks and mapping computations onto GPUs and CPUs as well as optimization of CPU-GPU communication. Another general data processing platform utilizing GPUs is G-Storm [7] which can be used for various applications and data types and provides a high level programming approach. It handles data transfers and resource allocation automatically. If data is to be further used on the same GPU in subsequent operations, it will not be copied back and forth between the host and the GPU. G-Storm very much relies on CUDA MPS that allows to create a single context that can be used from many processes on the host. It should be noted that this paper evaluates gains from MPS and shows benefits of multi-threaded and CUDA multi-stream approach for even better performance and such can be used to improve existing systems. Paper [40] proposes an efficient real-time system for processing large amounts of high frequency data such as video and text. The approach integrates Hadoop for parallel processing, Spark for the real-time component and GPUs for processing. Matrix type data is processed on GPUs similarly to MapReduce. The authors conclude that the proposed solution is faster than CPU MapReduce. Such a system could also benefit from low level optimization between host and GPUs presented in this paper. Work [44] proposes a CPU+GPU system for processing a large number of incoming data streams with hard real-time constraints. A scheduler running on the CPU side distributes streams among CPUs and GPUs for high utilization of the system in order to meet the constraints. The solution was evaluated using an AES-CBC encryption kernel on thousands of streams proving over 80% more data processing rate than a single GPU system. Paper [42] presents KernelHive that can be used to optimize scheduling and execution of processing using a stream of multiple independent data chunks on hybrid CPU+GPU systems. Efficient multithreaded data stream processing in a workflow management system called BeesyCluster, either within a high performance workstation or even spanning multiple clusters, is presented in paper [10]. In that context, the contribution of this paper is optimization of internal building blocks for efficient GPU management and consideration of power consumption as well.

3 PROCESSING MODEL AND DESIGN OF BENCHMARK

This section presents the custom-developed application benchmark that is representative of various applications run on GPUs or in a hybrid CPU+GPU environment. Many variables have been considered and can be changed in the proposed processing model and as such were used for subsequent tests. Design of the benchmark application is shown in Figure 1. It is assumed that the application processes a sequence of input data packets such that two data packets serve as input to a processing function that produces output data. This general assumption corresponds to many real life applications, depending on relative sizes of output and input data, e.g. multiplication, addition or other operations on matrices that are important computational steps in various artificial intelligence applications such as deep neural network training. Parallelization involves the following elements and ideas:

1. At a high level of parallelism, OpenMP threads are spawned – one thread per each GPU and additionally one thread managing computations on a multi-core CPU(s). These threads fetch input data from memory in a critical section and pass for computations either to a GPU or the CPU(s). This scheme, working in a loop, effectively supports dynamic load balancing among compute devices.
2. Nested OpenMP parallelism is used for parallelization with many threads on the CPU(s).
3. Input data can be stored in regular RAM from which it can be sent to GPU's global memory explicitly or stored in previously allocated space in Unified Memory. In the latter case, prefetching can be turned on for enabling overlapping computations with host-device communication. In the case of the Unified Memory based version, streams are still used for maximum concurrency of operations [34].

The benchmark allows to set various modes and parameters and correspondingly allows to mimic behavior of various applications following the assumed processing pattern:

1. memory mode – several modes are possible:
 - (a) allocation of host memory *std* using `cudaHostAlloc()` with flag `cudaHostAllocPortable` that does allow subsequent overlapping computations and communication in various streams,
 - (b) allocation of memory *UM* using Unified Memory (UM) by calling `cudaMallocManaged()` that allows to use the same pointer from both host threads to write input data, from a kernel to read input data and write output results as well as from the host to read output.
 - (c) allocation of memory *UMprefetch* using Unified Memory with data prefetching through `cudaMemPrefetchAsync()` for streams to be used in subsequent steps,
2. compute time to host-device communication time ratio that corresponds to the computational time on a given input data chunk divided by the communication time of this data chunk (CPU-GPU-CPU),






Figure 1: Proposed processing framework

3. output-input ratio that denotes the ratio of the size of output data to the size of input data,
4. stream count – the number of streams per one GPU used,
5. host thread count – the number of threads among which computations are scheduled on CPU(s) cores,
6. gpu count <ids of gpus> – the number and ids of GPU(s) to be used for computations.

In each experiment, unless otherwise noted, data chunk was 256KB in size and 1.6 GBs of data was processed. In the test we assumed 1024 threads per block and the total number



of threads was 262144. In the GPU kernel function, a thread fetches its unique index in a grid and processes data from two input arrays into a result stored in its own location (depending on its index) in an output array. Specifically, it computes averages of selected vector elements of the two input arrays and computes a distance between the averages which is added to the final output. All arrays are stored in global memory and the kernel uses 3 variables as temporary indices and one variable as a loop counter. Compute time to host-device communication time ratio is configured with a proper number of iterations of the aforementioned loop.

4 EXPERIMENTS AND TESTS

4.1 Testbed systems

For experiments, we used the benchmark described in Section 3 run on three modern multicore CPU(s) + GPUs workstations. Specifications of the systems are listed in Table 1. Testbeds 1 and 2 feature 2 Intel Xeon CPUs + 2 NVIDIA GPUs, of various generations while testbed 3 an Intel Xeon CPU + 4 NVIDIA Tesla V100 cards used for testing scaling across multiple GPUs.

For each particular configuration, unless otherwise noted, 10 tests were performed and the average value is presented.

4.2 Impact of multiple streams on performance

The purpose of following experiments is to determine the impact of using multiple streams for overlapping computations and communication and finally execution time of a GPU enabled application.

The following tests have been performed for several values of compute time to host-device communication time ratio, for several GPU cards and for the number of streams between 1 and 32. Additionally, two different ways of launching computations on a GPU are presented and compared:

1. A: one thread per GPU managing computations through one or more streams. In this case, the thread launches CPU-GPU communication, kernel and GPU-CPU communication asynchronously through streams one after another.
2. B: as many threads as the number of streams are launched per GPU, each of which launches CPU-GPU communication, kernel, GPU-CPU communication in a separate stream. Threads need to synchronize while fetching new input data packets.

Figure 2 presents results for these versions for particular numbers of threads and streams used for testbed 1 while Figure 3 does so for testbed 2. It can be seen that, in general, best results were obtained using one dedicated host thread per GPU launching communication and computations to multiple streams with 2 streams for testbed 1. For testbed 2 the same implementation offers best results with 2+ streams with small differences between the number of streams larger than 2-4. At the same time, we can see very



Table 1: Testbed configurations

Testbed	1	2	3
CPUs	2 x Intel Xeon CPU E5-2620v4 @ 2.10GHz	2 x Intel(R) Xeon(R) CPU E5-2640 @ 2.50GHz	Intel(R) Xeon(R) CPU E5-2698 v4 @ 2.20GHz
CPUs – total number of physical/logical cores	16/32	12/24	20/40
System memory size (RAM) [GB]	128	64	256
GPUs	2 x NVIDIA GTX 1070 (Pascal)	2 x NVIDIA Tesla K20m (Kepler)	4 x NVIDIA Tesla V100 (Volta)
GPUs – total number of CUDA cores	2 x 2048	2 x 2496	4 x 5120
GPU Compute capability	6.1	3.5	7.0
GPU memory size [MB]	2 x 8192	2 x 5120	4 x 16384
Operating system	Ubuntu Linux version 4.15.0-36-generic	CentOS Linux version 3.10.0-862.9.1.el7.x86_64	Ubuntu Linux version 4.4.0-83-generic
Compiler/version	Cuda compilation tools, release 9.1, V9.1.85, gcc 7.3.0	Cuda compilation tools, release 9.1, V9.1.85, gcc 4.8.5	Cuda compilation tools, release 9.0, V9.0.176, gcc 5.4.0

small deviations between runs (10 measured) for testbed 2 (default affinity values are presented). For testbed 1, we can observe larger deviations for configurations with multiple host threads launching operations on the GPU (we present threads/close affinity values). These differences might stem from various operating system settings and compiler version as the CUDA versions were the same.

4.3 Launching computations from multiple processes using MPS

In case there is no dedicated parallel application available for parallel processing of incoming data streams to a computer node, it is probable that several processes working in parallel will try to submit work for processing on a GPU that will be shared in such a




Figure 2: Comparison of implementations with various numbers of threads and streams on a GPU, testbed 1, bars represent standard deviation

case. This may lead to inefficiency of processing. One solution would involve writing a dedicated multi-stream application as analyzed in this paper. An alternative approach has been made available by NVIDIA through Multi Process Service (MPS) that tries to overlap CPU-GPU communication and processing on a GPU from various contexts. It does not require code modifications which is a considerable advantage. Details of its usage can be found in [13]. Figures 4 and 5 present results of using the MPS enabled configuration vs the standard configuration for testbed 1 and testbed 2 respectively. Five tests were performed for each configuration and the average value is presented. Results really indi-






Figure 3: Comparison of implementations with various numbers of threads and streams on a GPU, testbed 2, bars represent standard deviation

cate that the solution improves execution time visibly, except for smaller compute time to host-device communication time ratio for testbed 1. In these tests, two different processes were launched in parallel on the number of data chunks half the sizes of the cases shown in Figures 2 and 3. It should be noted that the best results obtained for 2 streams shown in Figure 2 still offer better execution times while the ones shown in Figure 3 show practically the same or marginally better performance compared to the one with MPS.






Figure 4: Comparison of performance with and without NVIDIA MPS, testbed 1

4.4 Performance with Unified Memory

Since the latest cards and CUDA versions offer the benefit of easier programming with Unified Memory, this experiment is to show performance of Unified Memory based implementation compared to previous best cases. The test involves setting input data on the host and launching a kernel that processes data packets on the GPU. Subsequently, results are read from the host side in order to find the maximum of results and display to the user.

The basic UM enabled version was further optimized using data prefetching (we denote this version by `UMprefetch`). Specifically, the data packet to be processed in a subsequent step in a given stream is prefetched using a call to function `cudaMemPrefetchAsync(...)` on the two input buffers.

Figure 6 presents comparison between `std`, `UM` and `UMprefetch` versions for 1 GPU on testbed 1 while Figure 7 presents comparison between `std`, `UM` and `UMprefetch` versions for 2 GPUs on testbed 1, data size proportionally smaller than in the previous tests. It can be seen that prefetching really improves performance but still the standard memory optimized multi-stream version offers the best performance. This is in line with some previous works comparing performance of Unified Memory to standard based versions showing generally similar or worse performance in [22], [30] and [37] in return for an easier programming model. This paper confirms it for various compute time to host-device communication time ratios, numbers of streams and 1 and 2 GPUs.




Figure 5: Comparison of performance with and without NVIDIA MPS, testbed 2

4.5 Scalability of hybrid CPU+GPU code

The purpose of the following experiments (using standard memory management) is to show scalability of the hybrid parallel code on the two testbeds with 1 GPU, 2 GPUs as well as host threads engaged for computations, for various GPU/CPU performance ratios. The latter can vary depending on an application. In this case, 2 streams per GPU were used for testbed 1 and 4 streams per GPU for testbed 2. The same thread affinities and binding as in Section 4.2 were used.

The results presented in Figure 8 for testbed 1 and in Figure 9 for testbed 2 allow to assess GPU/CPU performances for which adding host threads for computations brings visible savings in execution times. It can be noticed that 2 GPUs configurations achieve relatively better performance than proportional scaling from 1 GPU configurations, apparently due to using one of the GPUs for display as well. Scaling from 1 to 2 GPUs is clearly visible. Increasing the number of host threads decreases application execution time at rates very much depending on GPU to CPU performances, with practically no gains when using 2 GPUs and GPU/CPU performance ratio around 30 for testbed 1. It should be kept in mind that in case some CPU cores are used for computations, still as many threads as the number of GPUs are used for management of computations on the GPUs. Furthermore, the threads managing computations on the GPUs and the CPUs fetch next data packets synchronizing on an OpenMP critical section which also decreases potential speed-ups.




Figure 6: Comparison of standard memory (std), Unified Memory (UM) and optimized Unified Memory (UMprefetch) implementations – testbed 1, 1 GPU, bars represent standard deviation

4.6 Performance-power consumption ratio

In today's high performance computing systems, power consumption has become an important topic. It is considered in designs of future clusters for which the total power consumption is suggested not to exceed 20MW for 1 Exaflop/s [17]. In this context, we analyze performance to power consumption for the various configurations analyzed in this paper, specifically for:






Figure 7: Comparison of standard memory (std), Unified Memory (UM) and optimized Unified Memory (Umprefetch) implementations – testbed 1, 2 GPUs, bars represent standard deviation

1. various numbers of streams involved when using 1 GPU,
2. GPU + CPU configurations with various numbers of host threads involved in computations.

GPU performance was calculated as the inverse of the sum of data chunk CPU-GPU communication, processing and GPU-CPU result transfer times. CPU performance was calculated as the inverse of data chunk processing time on the CPU(s). Average power consumption of various configurations was measured using a hardware meter within a




Figure 8: Performance of a hybrid GPU+CPU implementation for various GPU/CPU performances and numbers of host threads, testbed 1




Figure 9: Performance of a hybrid GPU+CPU implementation for various GPU/CPU performances and numbers of host threads, testbed 2

10 minute period for each configuration. A bash script was used to run a particular configuration. Figure 10 shows normalized performance computed as inverse of execution time divided by average power consumption through application run for 1 GPU and various numbers of streams. Normalization of performance was done by dividing results by quotients of compute time to host-device communication time ratios of various configurations. It can be seen that normalized performance per power consumption has its maxima depending on compute time to host-device communication time ratio. It is interesting to note that for 2+ numbers of streams the best normalized ratios are observed for compute time to host-device communication time ratio 9.98 and lower for the other ratios.

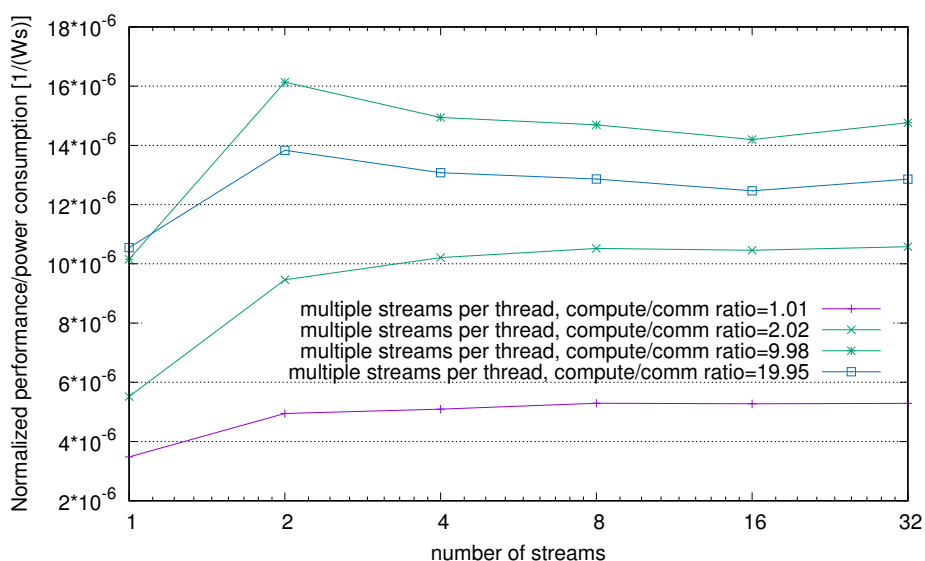


Figure 10: Normalized performance by power consumption for 1 GPU and various numbers of streams, testbed 1

Furthermore, performance by power consumption is shown for 1 and 2 GPU configurations with addition of various numbers of host threads used for computations using testbed 1. Results for the GPU/CPU performance ratio of around 30 are shown in Figure 11. It can be seen that, while execution times slightly decrease as shown in Figure 8 before, the performance-power consumption ratio goes down due to too little improvement of execution times thanks to CPU compared to its power consumption. Had the computational power been better compared to GPUs, the ratio would have been better for higher numbers of host threads. Such a simulation was performed and its results are shown in Figure 12 for a smaller GPU/CPU relative performance ratio. It can be seen from the tests that for GPU/CPU performance equal to 2 using more host threads offers benefits in terms of performance/power consumption.



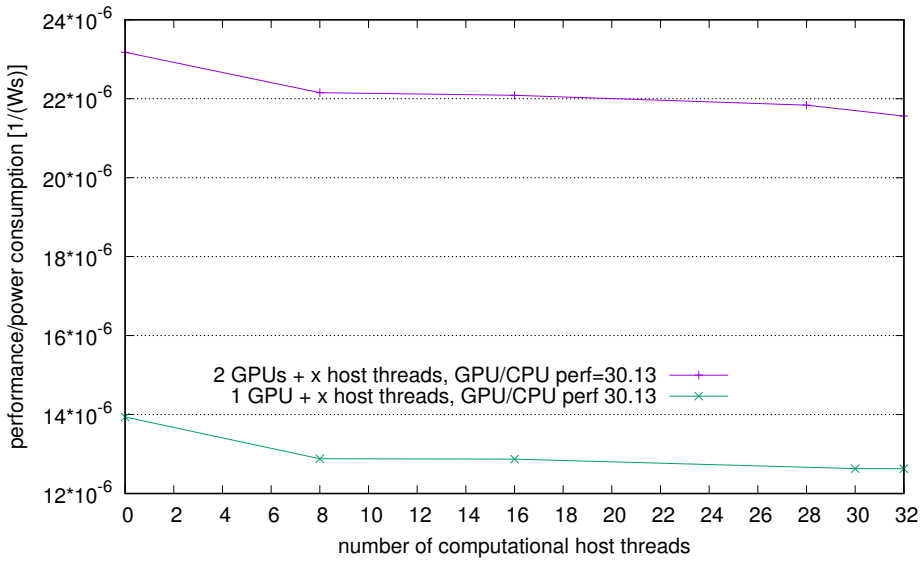


Figure 11: Performance by power consumption for 1 and 2 GPU configurations and various numbers of host threads, GPU/CPU performance=30.13, testbed 1

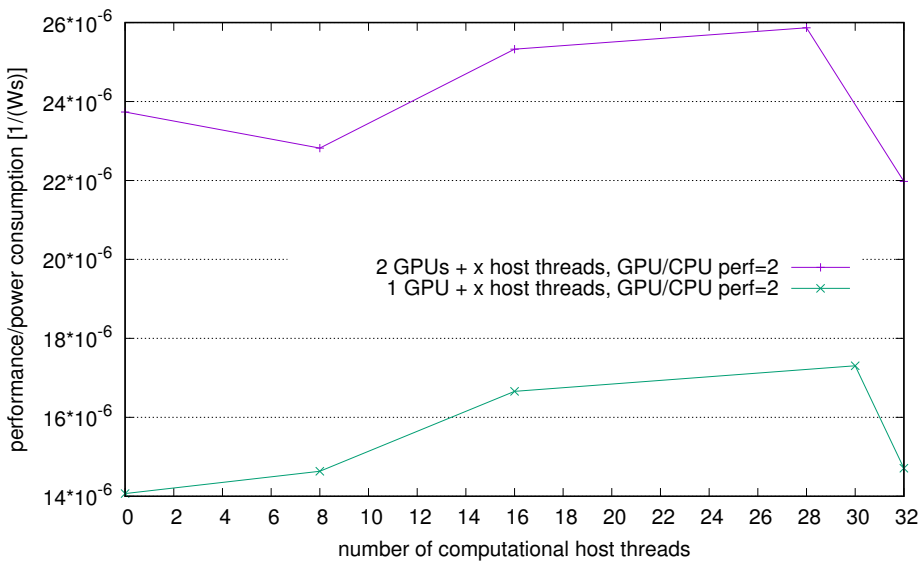


Figure 12: Performance by power consumption for 1 and 2 GPU configurations and various numbers of host threads, GPU/CPU performance=2, testbed 1

4.7 Scaling across multiple GPUs




Figure 13: Execution time vs number of streams for various numbers of GPUs, largest case from Figure 2, testbed 3

The following experiments demonstrate how the code scales across GPUs in testbed 3 with 4 NVIDIA Tesla V100 Volta series GPU cards. Firstly, Figure 13 presents how the numbers of streams affect performance for the largest configuration shown in Figure 2. Then, assuming 8 streams per GPU which (the configuration which already gives small execution times on the flat parts of the chart), execution times are shown for the standard memory management, UM and Umprefetch as in the previous cases. It can be seen in Figure 14 that, again, the UM version offers visible overhead over the standard memory management version. Umprefetch, thanks to manual prefetching, offers performance half-way between these two versions for 1 GPU. For 2 and 4 GPUs, it is worse than the standard memory management version by about 30% of the difference between the other two versions.

5 CONCLUSIONS AND FUTURE WORK

In the paper we analyzed performance and performance to power consumption ratio of multi-stream data processing on modern multicore CPU+ GPU systems. Using a benchmark that allows to set up various compute time to host-device communication time ratios, number of streams, number of threads managing computations it was possible to assess performance of various configurations on modern testbeds with Intel Xeon CPUs






Figure 14: Execution time vs memory management solutions for various numbers of GPUs, largest case from Figure 2, 8 streams per GPU, testbed 3

and NVIDIA Tesla K20m, GTX 1070 Pascal and Tesla V100 Volta series cards. Benefits of using a properly implemented multi-stream code were shown compared to GPU computations managed by various threads or processes for various numbers of streams. Furthermore, benefits of such code compared to standard Unified Memory and Unified Memory with prefetching were shown showing performance gains at the cost of increased programming effort. Additionally, gains from using NVIDIA Multi Process Service have been presented for multi-process configurations. Performance to power consumption ratios have been shown for various numbers of streams and compute time to host-device communication time ratios as well as for hybrid CPU+GPU configurations for various numbers of computational threads on the host and relative GPU and CPU performances. Scalability of the code was presented between 1 and 4 GPUs using NVIDIA Tesla V100 cards.

The results can be generalized as follows. For the considered data stream processing application and various compute to communication ratios, using multiple streams, at least 2 offered visible benefits, with best results using one dedicated host thread per GPU launching communication and computations to multiple streams. Some configurations result in best execution times for 2, 4, 8 or even 16 streams but we can note that benefits over 4 streams, if any, are very small. Secondly, we confirmed that using NVIDIA MPS gives visible benefits especially for larger compute to communication ratio. Furthermore, for various compute to communication ratios we confirmed that Unified Memory brings visible overhead over the standard memory management implementation while a Unified Memory version with manual prefetching ranks between the two. For CPU+GPU codes,



increasing the number of computational host threads up to the number of available logical processors decreases application execution time at rates very much depending on GPU to CPU performances with considerable gains with CPU performance in the same order as the one of the GPU. It has been shown that observed performance per power consumption varies with the number of streams, GPU to CPU performance ratio and the number of computational host threads.

These results can be used as guidelines for best performance implementations for various applications as the tests are of generic nature and, depending on values of particular aforementioned parameters, are representative of many applications. Specifically, obtained results can be used for implementation of building blocks for data stream frameworks using multi-core CPUs and GPUs, especially multi CUDA stream communication optimization.

Future work includes extending the scope of the conducted tests performed on systems with NVIDIA Tesla V100, specifically regarding various CPU+GPU configurations, tests for various compute/communication ratios, as well as extending tests to larger V100 based systems such as NVIDIA DGX-1 featuring 8 V100 GPUs. More experiments with thread affinities will be conducted, with research of their impact for particular codes. Additionally, we plan to incorporate of the outcomes of this work into higher level frameworks such as KernelHive [42] and possibly others and investigate the impact of Unified Memory oversubscription compared to the traditional implementation model.

6 ACKNOWLEDGMENTS

The research in the paper has been partially supported by Statutory Funds of Electronics, Telecommunications and Informatics Faculty, Gdansk University of Technology, Poland. Additionally, the author would like to express his gratitude to Aleksandra Preiss from Gdansk University of Technology.

References

- [1] Marco Aldinucci, Sonia Campa, Marco Danelutto, Peter Kilpatrick, and Massimo Torquati. Targeting distributed systems in fastflow. In Ioannis Caragiannis, Michael Alexander, Rosa Maria Badia, Mario Cannataro, Alexandru Costan, Marco Danelutto, Frédéric Desprez, Bettina Krammer, Julio Sahuquillo, Stephen L. Scott, and Josef Weidendorfer, editors, *Euro-Par 2012: Parallel Processing Workshops*, pages 47–56, Berlin, Heidelberg, 2013. Springer Berlin Heidelberg.
- [2] Marco Aldinucci, Marco Danelutto, Peter Kilpatrick, and Massimo Torquati. *Fastflow: High-Level and Efficient Streaming on Multicore*, chapter 13, pages 261–280. Wiley-Blackwell.
- [3] T. BINGMANN, M. AXTMANN, E. JÖBSTL, S. LAMM, H. C. NGUYEN, A. NOE, S. SCHLAG, M. STUMPP, T. STURM, AND P. SANDERS. THRILL: High-performance algorithmic distributed batch data processing with c++. In *2016 IEEE International Conference on Big Data (Big Data)*, pages 172–183, Dec 2016.

- [4] T. Blattner, W. Keyrouz, J. Chalfoun, B. Stivalet, M. Brady, and S. Zhou. A hybrid cpu-gpu system for stitching large scale optical microscopy images. In *2014 43rd International Conference on Parallel Processing*, pages 1–9, Sep. 2014.
- [5] PARIS CARBONE, ASTERIOS KATSIFODIMOS, STEPHAN EWEN, VOLKER MARKL, SEIF HARIDI, AND KOSTAS TZOUMAS. APACHE FLINKTM: Stream and batch processing in a single engine. *IEEE Data Eng. Bull.*, 38:28–38, 2015.
- [6] V. Cardellini, V. Grassi, F. L. Presti, and M. Nardelli. On qos-aware scheduling of data stream applications over fog computing infrastructures. In *2015 IEEE Symposium on Computers and Communication (ISCC)*, pages 271–276, July 2015.
- [7] Z. CHEN, J. XU, J. TANG, K. A. KWIAT, C. A. KAMHOUA, AND C. WANG. GPU-ACCELERATED HIGH-THROUGHPUT ONLINE STREAM DATA PROCESSING. *IEEE Transactions on Big Data*, 4(2):191–202, June 2018.
- [8] Rommel Cruz, Lucia Drummond, Esteban Clua, and Cristiana Bentes. Analyzing and estimating the performance of concurrent kernels execution on gpus. *Simpósio em Sistemas Computacionais de Alto Desempenho (WSCAD)*, 18(1/2017), 2017.
- [9] C. CULLINAN, C. WYANT, T. FRATTESI, AND X. HUANG. COMPUTING PERFORMANCE BENCHMARKS AMONG CPU, GPU, AND FPGA. 2012. WORCESTER POLYTECHNIC INSTITUTE, E-PROJECT-030212-123508, [HTTPS://web.wpi.edu/Pubs/E-project/Available/E-project-030212-123508/unrestricted/Benchmarking_Final.pdf](https://web.wpi.edu/Pubs/E-project/Available/E-project-030212-123508/unrestricted/Benchmarking_Final.pdf).
- [10] PAWEŁ CZARNUL. A MODEL, DESIGN, AND IMPLEMENTATION OF AN EFFICIENT MULTITHREADED WORKFLOW EXECUTION ENGINE WITH DATA STREAMING, CACHING, AND STORAGE CONSTRAINTS. *The Journal of Supercomputing*, 63(3):919–945, Mar 2013.
- [11] PAWEŁ CZARNUL. BENCHMARKING PERFORMANCE OF A HYBRID INTEL XEON/X-EON PHI SYSTEM FOR PARALLEL COMPUTATION OF SIMILARITY MEASURES BETWEEN LARGE VECTORS. *International Journal of Parallel Programming*, 45(5):1091–1107, Oct 2017.
- [12] Paweł Czarnul. Benchmarking overlapping communication and computations with multiple streams for modern gpus. In Maria Ganzha, Leszek A. Maciaszek, and Marcin Paprzycki, editors, *Communication Papers of the 2018 Federated Conference on Computer Science and Information Systems, FedCSIS 2018, Poznań, Poland, September 9-12, 2018.*, pages 105–110, 2018.
- [13] Paweł Czarnul. *Parallel Programming for Modern High Performance Computing Systems*. Chapman and Hall/CRC, Taylor&Francis, 1st edition, 2018. ISBN 9781138305953.
- [14] PAWEŁ CZARNUL. PARALLELIZATION OF LARGE VECTOR SIMILARITY COMPUTATIONS IN A HYBRID CPU+GPU ENVIRONMENT. *The Journal of Supercomputing*, 74(2):768–786, Feb 2018.
- [15] Paweł Czarnul, Jerzy Proficz, and Krzysztof Drypczewski. Survey of methodologies, approaches and challenges in parallel programming using high performance computing systems. *Scientific Programming*, 2020. in press.
- [16] MARCO DANELUTTO, DANIELE DE SENSI, AND MASSIMO TORQUATI. A POWER-AWARE, SELF-ADAPTIVE MACRO DATA FLOW FRAMEWORK. *Parallel Processing Letters*, 27(01):1740004, 2017.
- [17] JACK DONGARRA. CHALLENGES FOR EXASCALE COMPUTING, JUNE 2010. PARA 2010, [HTTP://www.netlib.org/utk/people/JackDongarra/SLIDES/para-06102.pdf](http://www.netlib.org/utk/people/JackDongarra/SLIDES/para-06102.pdf).



- [18] Ellen Friedman and Kostas Tzoumas. *Introduction to Apache Flink: Stream Processing for Real Time and Beyond*. O'Reilly Media, Inc., 1st edition, 2016.
- [19] P. Guo and C. Zhang. Performance optimization for spmv on multi-gpu systems using threads and multiple streams. In *2016 International Symposium on Computer Architecture and High Performance Computing Workshops (SBAC-PADW)*, pages 67–72, Oct 2016.
- [20] JUAN GÓMEZ-LUNA, JOSÉ MARÍA GONZÁLEZ-LINARES, JOSÉ IGNACIO BENAVIDES, AND NICOLÁS GUIL. PERFORMANCE MODELS FOR ASYNCHRONOUS DATA TRANSFERS ON CONSUMER GRAPHICS PROCESSING UNITS. *Journal of Parallel and Distributed Computing*, 72(9):1117 – 1126, 2012. Accelerators for High-Performance Computing.
- [21] Ankit Jain. *Mastering Apache Storm: Real-time Big Data Streaming Using Kafka, Hbase and Redis*. Packt Publishing, 2017.
- [22] LUKASZ JARZABEK AND PAWEŁ CZARNUL. PERFORMANCE EVALUATION OF UNIFIED MEMORY AND DYNAMIC PARALLELISM FOR SELECTED PARALLEL CUDA APPLICATIONS. *The Journal of Supercomputing*, 73(12):5378–5401, Dec 2017.
- [23] Marcin Knap and Paweł Czarnul. Performance evaluation of unified memory with prefetching and oversubscription for selected parallel cuda applications on nvidia pascal and volta gpus. *The Journal of Supercomputing*, Aug 2019.
- [24] B. KO, S. HAN, Y. PARK, M. JEON, AND B. LEE. A COMPARATIVE STUDY OF PROGRAMMING ENVIRONMENTS EXPLOITING HETEROGENEOUS SYSTEMS. *IEEE Access*, 5:10081–10092, 2017.
- [25] Jochen Kreutz. Cuda streams, events and asynchronous memory copies, April 2017. GPU Programming@Jülich Supercomputing Centre, https://www.fz-juelich.de/SharedDocs/Downloads/IAS/JSC/EN/slides/cuda/09-cuda-streams-events.pdf?__blob=publicationFile.
- [26] CHUN-LIANG LEE, YI-SHAN LIN, AND YAW-CHUNG CHEN. A HYBRID CPU/GPU PATTERN-MATCHING ALGORITHM FOR DEEP PACKET INSPECTION. *PLOS ONE*, 10(10):1–22, 10 2015.
- [27] GUO-QING LEI, YONG DOU, WEN WAN, FEI XIA, RONGCHUN LI, MENG MA, AND DAN ZOU. CPU-GPU HYBRID ACCELERATING THE ZUKER ALGORITHM FOR RNA SECONDARY STRUCTURE PREDICTION APPLICATIONS. *BMC Genomics*, 13(S-1):S14, 2012.
- [28] Jonathan Leibiusky, Gabriel Eisbruch, and Dario Simonassi. *Getting Started with Storm*. O'Reilly Media, Inc., 2012.
- [29] H. LI, D. YU, A. KUMAR, AND Y. TU. PERFORMANCE MODELING IN CUDA STREAMS - A MEANS FOR HIGH-THROUGHPUT DATA PROCESSING. IN *IEEE International Conference on Big Data. Proceedings*, PAGES 301–310, 2014. [HTTP://doi.org/10.1109/BigData.2014.7004245](http://doi.org/10.1109/BigData.2014.7004245).
- [30] W. Li, G. Jin, X. Cui, and S. See. An evaluation of unified memory technology on nvidia gpus. In *2015 15th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, pages 1092–1098, May 2015.
- [31] JUSTIN LUITJENS. CUDA STREAMS. BEST PRACTICES AND COMMON PITFALLS, 2014. NVIDIA. GPU TECHNOLOGY CONFERENCE, [HTTP://on-demand.gputechconf.com/gtc/2014/presentations/S4158-cuda-streams-best-practices-common-pitfalls.pdf](http://on-demand.gputechconf.com/gtc/2014/presentations/S4158-cuda-streams-best-practices-common-pitfalls.pdf).

- [32] MICHAŁ MARKS, JAROSŁAW JANTURA, EWA NIEWIADOMSKA-SZYNKIEWICZ, PRZEMYSŁAW STRZELCZYK, AND KRZYSZTOF GOZDZ. HETEROGENEOUS GPU&CPU CLUSTER FOR HIGH PERFORMANCE COMPUTING IN CRYPTOGRAPHY. *Computer Science*, 13(2):63, 2012.
- [33] CLAUDIA MISALE, MAURIZIO DROCCO, GUY TREMBLAY, AND MARCO ALDINUCCI. PICO: A novel approach to stream data analytics. In Dora B. Heras, Luc Bougé, Gabriele Mencagli, Emmanuel Jeannot, Rizos Sakellariou, Rosa M. Badia, Jorge G. Barbosa, Laura Ricci, Stephen L. Scott, Stefan Lankes, and Josef Weidendorfer, editors, *Euro-Par 2017: Parallel Processing Workshops*, pages 118–128, Cham, 2018. Springer International Publishing.
- [34] NVIDIA. CUDA TOOLKIT v10.0.130 PROGRAMMING GUIDE, OCTOBER 2018. [HTTPS://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html](https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html).
- [35] KENIA PICOS, VÍCTOR H. DÍAZ RAMÍREZ, AND JUAN J. TAPIA. REAL-TIME 3D VIDEO PROCESSING USING MULTI-STREAM GPU PARALLEL COMPUTING. *Research in Computing Science*, 80:87–95, 2014.
- [36] Marcus Pinnecke, David Broneske, and Gunter Saake. Toward GPU accelerated data stream processing. In Gunter Saake, David Broneske, Sebastian Dorok, and Andreas Meister, editors, *Proceedings of the 27th GI-Workshop Grundlagen von Datenbanken, Gommern, Germany, May 26-29, 2015.*, volume 1366 of *CEUR Workshop Proceedings*, pages 78–83. CEUR-WS.org, 2015.
- [37] A. PIRJA AND M. PETROSANU. IMPROVING PARALLEL PROGRAMMING IN THE COMPUTE UNIFIED DEVICE ARCHITECTURE USING THE UNIFIED MEMORY FEATURE, 2014. [FTP://ftp.repec.org/opt/ReDIF/RePEc/rau/jisomg/WI14/JISOM-WI14-A14.pdf](ftp://ftp.repec.org/opt/ReDIF/RePEc/rau/jisomg/WI14/JISOM-WI14-A14.pdf).
- [38] ALAN PRYOR, COLIN OPHUS, AND JIANWEI MIAO. A STREAMING MULTI-GPU IMPLEMENTATION OF IMAGE SIMULATION ALGORITHMS FOR SCANNING TRANSMISSION ELECTRON MICROSCOPY. *Advanced Structural and Chemical Imaging*, 3(1):15, Oct 2017.
- [39] DAVID RADFORD AND DAVID CALVERT. A COMPARATIVE ANALYSIS OF THE PERFORMANCE OF SCALABLE PARALLEL PATTERNS APPLIED TO GENETIC ALGORITHMS AND CONFIGURED FOR NVIDIA GPUS. *Procedia Computer Science*, 114:65 – 72, 2017. Complex Adaptive Systems Conference with Theme: Engineering Cyber Physical Systems, CAS October 30 – November 1, 2017, Chicago, Illinois, USA.
- [40] M. MAZHAR RATHORE, HOJAE SON, AWAIS AHMAD, ANAND PAUL, AND GWANGGIL JEON. REAL-TIME BIG DATA STREAM PROCESSING USING GPU WITH SPARK OVER HADOOP ECOSYSTEM. *International Journal of Parallel Programming*, 46(3):630–646, Jun 2018.
- [41] STEVE RENNICH. CUDA C/C++. STREAMS AND CONCURRENCY, 2011. NVIDIA, [HTTP://on-demand.gputechconf.com/gtc-express/2011/presentations/StreamsAndConcurrencyWebinar.pdf](http://on-demand.gputechconf.com/gtc-express/2011/presentations/StreamsAndConcurrencyWebinar.pdf), accessed on 19th July 2017.
- [42] PAWEŁ ROŚCISZEWSKI, PAWEŁ CZARNUL, RAFAŁ LEWANDOWSKI, AND MARCEL SCHALLY-KACPRZAK. KERNELHIVE: a new workflow-based framework for multilevel high performance computing using clusters and workstations with cpus and gpus. *Concurrency and Computation: Practice and Experience*, 28(9):2586–2607.
- [43] M. Sourouri, T. Gillberg, S. B. Baden, and X. Cai. Effective multi-gpu communication using multiple cuda streams and threads. In *2014 20th IEEE International Conference on Parallel*



and Distributed Systems (ICPADS), pages 981–986, Dec 2014.

- [44] URI VERNER, ASSAF SCHUSTER, MARK SILBERSTEIN, AND AVI MENDELSON. SCHEDULING PROCESSING OF REAL-TIME DATA STREAMS ON HETEROGENEOUS MULTI-GPU SYSTEMS. IN *Proceedings of the 5th Annual International Systems and Storage Conference, SYSTOR '12*, PAGES 8:1–8:12, 2012.
- [45] L. Wang, M. Huang, and T. El-Ghazawi. Exploiting concurrent kernel execution on graphic processing units. In *2011 International Conference on High Performance Computing Simulation*, pages 24–32, July 2011.
- [46] YUEQING WANG, FANG WANG, RONGCHUN LI, AND YONG DOU. AN EFFICIENT CPU-GPU HYBRID PARALLEL IMPLEMENTATION FOR DVB-RCS2 RECEIVER. *Concurrency and Computation: Practice and Experience*, 30(19):e4529, 2018. e4529 CPE-17-0020.R2.
- [47] MATEI ZAHARIA, REYNOLD S. XIN, PATRICK WENDELL, TATHAGATA DAS, MICHAEL ARMBRUST, ANKUR DAVE, XIANGRUI MENG, JOSH ROSEN, SHIVARAM VENKATARAMAN, MICHAEL J. FRANKLIN, ALI GHODSI, JOSEPH GONZALEZ, SCOTT SHENKER, AND ION STOICA. APACHE SPARK: A unified engine for big data processing. *Commun. ACM*, 59(11):56–65, October 2016.
- [48] Y. ZHANG AND F. MUELLER. GSTREAM: A general-purpose data streaming framework on gpu clusters. In *2011 International Conference on Parallel Processing(ICPP)*, volume 00, pages 245–254, 2011.



Pawel CZARNUL received his Ph.D. in computer science in 2003 and D.Sc. in computer science in 2015, both from Gdansk University of Technology, Poland. His research interests include: high performance computing, distributed information systems and processing, artificial intelligence. He is author of over 80 publications in the area of parallel and distributed processing, including book entitled *Parallel Programming for Modern High Performance Computing Systems*, Chapman and Hall/CRC, 2018. Currently Head of Computer Architecture Department and Vice-dean of the Faculty of ETI, Gdansk University of Technology, Poland.