

On the Scalability of Data Reduction Techniques in Current and Upcoming HPC Systems from an Application Perspective

Axel Huebl^{1,2(✉)} , René Widera¹ , Felix Schmitt^{2,3}, Alexander Matthes^{1,2} ,
Norbert Podhorszki⁴ , Jong Youl Choi⁴, Scott Klasky⁴,
and Michael Bussmann^{1(✉)} 

¹ Helmholtz-Zentrum Dresden – Rossendorf, Dresden, Germany
{a.huebl,m.bussmann}@hzdr.de

² Technische Universität Dresden, Dresden, Germany

³ NVIDIA ARC GmbH, Berlin, Germany

⁴ Oak Ridge National Laboratory, Oak Ridge, TN, USA

Abstract. We implement and benchmark parallel I/O methods for the fully-manycore driven particle-in-cell code PIConGPU. Identifying throughput and overall I/O size as a major challenge for applications on today's and future HPC systems, we present a scaling law characterizing performance bottlenecks in state-of-the-art approaches for data reduction. Consequently, we propose, implement and verify multi-threaded data-transformations for the I/O library ADIOS as a feasible way to trade underutilized host-side compute potential on heterogeneous systems for reduced I/O latency.

1 Introduction

Production-scale research simulation codes have been optimized in the last years to achieve maximum compute performance on leadership, heterogeneous computing systems such as the Titan supercomputer at Oak Ridge National Laboratory (ORNL). With close to perfect weak scaling domain scientists can increase spatial and temporal resolution of their simulation and explore systems without reducing dimensionality or feature resolution.

We present the consequences of near-perfect weak-scaling of such a code in terms of I/O demands from an application perspective based on production runs using the particle-in-cell (PIC) code PIConGPU [8, 17]. PIConGPU demonstrates a typical use case in which a PFlops/s-scale, performance portable simulation [29, 30] leads automatically to PByte-scale output even for single runs.

This project has received funding from the European Unions Horizon 2020 research and innovation programme under grant agreement No. 654220. An award of computer time was provided by the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

1.1 PIConGPU

PIConGPU is an electro-magnetic PIC code [5, 14] implemented via abstract, performance portable C++11 kernels on manycore hardware utilizing the Alpaka library [29, 30]. Its applications span from general plasma physics, over laser-matter interaction to laser-plasma based particle accelerator research.

Since its initial open-source release in 2013 with CUDA support, PIConGPU is reportedly the fastest particle-in-cell code in the world in terms of sustained peak Flops/s [8]. We achieved this by not only porting the bottlenecks of the PIC algorithm to new compute hardware but the complete code, thus minimizing data transfer. PIConGPU data structures are tiled and swapping of frequently updated data residing on device memory over low-bandwidth bottlenecks such as the PCI bus is avoided [6].

The overall simulation is spatially domain decomposed and only nearby border areas need to be communicated across compute nodes (and accelerators) inbetween iterations. Iterations in PIConGPU are performed with a frequency of about 10 Hz on current accelerator architectures (GPUs) when simulating 3D spatial domains and up to 60 Hz for two-dimensional domains. Each iteration updates electro-magnetic fields and plasma particles, which together constitute the simulation’s state.

1.2 Physical Observables

We will define primary observables as variables directly accessible and iterated within the simulation. In terms of an electro-magnetic PIC code these are electric field, magnetic field and plasma particles’ properties such as position, momentum, charge to mass ratio and weighting. Primary observables are convenient for the domain expert for exploration, of limited use for theories and models and nearly always inaccessible directly in experiments.

We define secondary observables as computable on-the-fly, PIC examples being the electric current density, position-filtered energy histograms or projected phase space distributions. In practice, analysis of a specific setup needs multiple additional, study-specific derivations from already derived observables which we summarize as tertiary observables. Examples in the domain of plasma physics are integrals over phase-space trajectories, time-averaged fields, sample trajectories or particle distributions in gradients of fields, flux over time, growth-rates, etc. Usually, observables accessible by experiments fall in this last category and can be compared to theoretical model predictions.

1.3 Two Example Workflows to Explore Complex Systems

In daily modeling work we usually iterate between two operational modes while investigating a new physical system. We start with an exploratory phase guided by initial hypotheses, looking at primary observables via visualizations or utilizing existing analysis pipelines to iterate over the result of strongly reduced

secondary and tertiary observables. During this phase, we develop new study-specific analysis steps and working hypotheses.

The second phase continues with a high-resolution, high-throughput scan of an identified regime of the physical system to prove or falsify our working hypotheses. Due to higher resolution and full physical modeling, new observations will emerge from that step. Research is then about iterating both steps in a refined manner until a system is well understood and a model is found to describe the complex processes of interest.

1.4 Structure of This Paper

As our guiding example, we describe the Titan and Summit systems at ORNL and their I/O bandwidth hierarchies from the special perspective of a fully GPU-driven, massively parallel PIC code. We then evaluate the performance of PIConGPU's I/O implementation, the overhead it introduces and mitigation strategies via on-the-fly data reductions. We address issues in current state-of-the-art compression schemes for our application and compare them to self-implemented compression schemes that make optimal use of underutilized hardware components. More specifically, we integrated the meta compression library *blosc* [2] into ADIOS, thereby for the first time enabling multi-threaded compression within ADIOS.

2 ORNL Titan and Summit Systems

With the launch of the Titan supercomputer to the public in 2013, manycore powered supercomputing finally became accessible on large-scale installations. Since then, the share of accelerator hardware in the TOP 100 systems has risen to one third [24]. Such heterogeneous systems concentrate their compute performance in the accelerator component, usually outnumbering the host system's compute potential by an order of magnitude, a trend that seems to continue on upcoming systems such as Summit.

2.1 I/O Limitations in State-of-the-Art Systems

The parallel file system Atlas at ORNL, partitioned in two islands of 14 PBytes each, provides an overall design parallel bandwidth of $B_{\text{parallel}} = 1 \text{ TByte/s}$. It is worth noting that *if* a hypothetical application would be constantly writing at this maximum parallel bandwidth, Atlas would run out of disk memory in less than 9 hour. We managed to write within each 8 hour production run of our plasma simulation code PIConGPU about 1 PByte of (zlib) compressed data, sampling the full system state every 2000 iteration steps. PIConGPU thus presents a realistic use case that can consume a significant fraction of those resources. With the upper limit of shared storage in mind, it is clear that data reduction comes with great value. Additionally, fast migration to and from tape

storage and a strictly imposed short data lifetime on Atlas also encourage users to avoid occupying disk memory for too long.

An equally severe limitation for I/O besides maximum data size is the overall time $t_{I/O}$ for file I/O compared to one iteration of the simulation, including data preparation time t_{prep} . Compared to the time $t_{\text{without I/O}}$ one iteration takes without I/O this $t_{I/O}$ introduces an overhead to the application run time, so that the single iteration runtime with I/O becomes $t_{\text{with I/O}} = t_{\text{without I/O}} + t_{I/O}$. When considering applications scaling to the full Titan system, reaching TByte/s overall throughput results in a maximum node-average throughput of 55 MByte/s. Applications with near perfect scaling can generate GPU data at two-digit Hz levels amounting to data rates as high as 10×6 GiByte/s (device global memory) on a node-local level, outnumbering the file system performance by three orders of magnitude. Asynchronous I/O lowers this dramatic gap temporarily, but still throttles the application at least to 1/10th of the bandwidth of the CPU-GPU interconnect, not accounting for data reorganization from tiled GPU memory to per-node contiguous memory as expected by parallel I/O APIs.

2.2 Staging, Burst Buffers and I/O Backlog

Even at moderate data rates, asynchronous writing can quickly overlap with the next consecutive write period. Staging [1, 10], if operating off-node, can reduce that data pressure but is similarly limited by another order of magnitude gap in throughput as soon as the interconnect is accessed.

Systems such as NERSC’s Cori recently introduced so-called burst-buffers [4]. Located either off-node similar to I/O nodes or in-node as with the upcoming Summit system, overall size of those burst buffers is usually similar to that of the global host RAM with access bandwidth ranging between network-interconnect and parallel filesystem bandwidth.

Burst buffers provide an interesting mean for temporary checkpointing and error-recovery. Coupled applications that only act as either a data sink or a source for the main application are also major beneficiaries of burst buffers. A prominent example in HPC are in situ visualizations copying on demand snapshots [3, 28] or accessing the primary observables directly [21, 23, 25].

Nevertheless, with the current absolute sizes of burst-buffers it is close to impossible to keep data between application lifetime and parallel filesystem data lifetime, simply because they cannot store a useful multiple of primary observables. As soon as a single stage in the I/O hierarchy is not drained as fast as it is filled, a backlog throughout all previous stages is inevitable even when buffers are used.

3 I/O Measurements

PICongPU implements I/O for outputs and checkpoints within its plugin system. Plugins are tightly coupled algorithms that can register within the main

application for execution after selected iterations. They share full access to primary observables (read and write) of the application.

I/O modules implemented are parallel HDF5 [27] and ADIOS (1.10.0) [20]. In order to tailor domain-specific needs for particle-mesh algorithms, libSplash is used as an abstraction layer [16]. Data objects are described by the meta-data standard openPMD [15] in human- and machine-readable markup, allowing for cross-application exchangeability as needed in post- and pre-processing workflows.

3.1 Preparation of PIConGPU Primary Observables for I/O

In preparation of GPU device data for I/O libraries, PIConGPU field data are copied from device to host via CUDA 3D memory copies while plasma particle attributes stored in tiled data structures are copied via the mallocMC [11] heap manager. Subsequently, scalar particle attributes are concatenated in preparation for efficient parallel I/O in a parallelized manner using OpenMP. The single GPU data size needed for saving a complete system state is typically $S = 4$ GiByte (assuming $2/3^{\text{rd}}$ of device global memory for primary observables). The overall time for preparing these 4 GiByte of data for one GPU is typically $t_{\text{prep}} = 1$ s on the systems considered in this publication.

3.2 I/O Performance in a Realistic Production Scenario

Measurements of the I/O performance are based on one of the default benchmarks implemented in PIConGPU, a simulation of the relativistic Kelvin-Helmholtz Instability [8, 12]. Starting from two spatially homogeneous, counterpropagating neutral plasma streams, a shear flow instability develops. This scenario shows good load-balancing due to nearly homogeneous data distribution across all GPUs with data size per output and GPU of $S = 4$ GiByte. We thus assume in our following analysis for sake of simplicity that indeed each node has the same output size, the same bandwidth and I/O operations have the same impact on all N nodes of a system.

Our benchmark systems are Titan (ORNL) and the K20 queue of Hypnos (HZDR), see Table 1. We choose the second system intentionally, since it has roughly the same age, similar ratio of Flops/s between CPU host and GPU device, multiple GPUs per node as in upcoming systems, even less CPU cores per GPU and an even higher single node average filesystem bandwidth compared to Titan. All measurement input and results of the following sections are available in the supplementary materials [18] and all software used is open source.

Most relevant from an application point of view is the absolute overhead $t_{\text{I/O}}$ in seconds caused by enabling I/O since it equals ‘wasted’ computing time that could be otherwise spent to iterate the problem further or in higher resolution. We define the effective parallel I/O throughput T_{eff} in GiByte/s as

$$T_{\text{eff}} \equiv N \times \frac{S}{t_{\text{with I/O}} - t_{\text{without I/O}}} = \frac{N \times S}{t_{\text{I/O}}} \quad (1)$$

Table 1. PIConGPU I/O benchmark systems, both commissioned in 2012/13: relevant system characteristics and single node average filesystem throughput T_{FS} , defined as the design parallel bandwidth B_{parallel} divided by N nodes

| | Titan | Hypnos (queue: ‘k20’) |
|--|---------------------|-----------------------|
| GPUs/node | 1× K20x | 4× K20m |
| CPUs/node | 1× AMD Opteron 6274 | 2× Intel Xeon E5-2609 |
| CPU-cores/GPU | 16 (8 FP) | 2 |
| GPU/CPU Flop/s (DP) | 9.3 : 1 | 7.6 : 1 |
| File system | Spider/Lustre | GPFS |
| $B_{\text{parallel}} = T_{\text{FS}} * N$ [GiByte/s] | 1000 | 20 |
| T_{FS} [GiByte/s] | 0.055 | 1.25 |
| CPU T_{memcpy} [GiByte/s] | 6.0 | 6.1 |
| Maximum number of nodes N_{max} | 18000 | 16 |

with the number of nodes N , the data size per node S and the difference between execution time with I/O $t_{\text{with I/O}}$ and without I/O $t_{\text{without I/O}}$ as $t_{\text{I/O}}$. Besides the (included) correction for intrinsic overheads in scaling the application, all measurements are performed as a weak scaling of PIConGPU, which is near-perfect up to the full size of Titan [8]. We average over 11 outputs within 2000 iterations with an average application iteration frequency of one Hertz.

In the following we model the I/O time per node by

$$t_{\text{I/O}}^{\text{simple}} \equiv t_{\text{prep}} + t_{\text{off RAM}} = t_{\text{prep}} + \frac{S}{T_{\text{FS}}} \quad (2)$$

defining t_{prep} as the time to concatenate data into large, I/O-API compatible chunks and $t_{\text{off RAM}} \equiv S/T_{\text{FS}}$ as the time to synchronously send the data off RAM. This preparation time can potentially be lowered by reorganizing data on the accelerator, where RAM is usually in full utilization from the application alone, while asynchronous (non-blocking) writes that hide data transfer latency require large enough temporary buffers to avoid backlog (see discussion in Sect. 2.2) and I/O library support. It is thus S/T_{FS} that will dominate overhead compared to iterations without I/O.

Figure 1 shows the achieved effective parallel I/O throughput T_{eff} on Titan. We noticed HDF5 I/O overhead getting prohibitively large for production runs as its parallelism is currently limited by the number of allocatable Lustre OSTs (≤ 160) on which one global file needs to be strided over. After optimizing HDF5 performance with MPI-I/O and HDF5 hints, first manually via best-practices and later using T3PIO [22], we turned down the strategy of parallel output in one global file (June 2014) and started adopting ADIOS aggregators, which enable transparently striding on subgroups of processes over a limited number of OSTs




Fig. 1. PIconGPU I/O weak scaling on Titan from 1 to 16384 K20x GPUs (nodes). Zlib was only supported serially with compression mode fast. MPI.Info hints for parallel HDF5 set via T3PIO (v2.3). For ADIOS, labels denote number of *OSTs*|*aggregators*, resulting for $N \geq 32$ in a striping of each aggregated process group over four OSTs. Lustre filesystem limits enforced 160 OSTs for (single-file) parallel HDF5 writes.

(latest benchmark: September 2015). When using ADIOS in this manner, we were able to reach an overall application throughput close to 280 GByte/s, see Fig. 2. We are not aware of substantial changes in the Atlas filesystem during this period of time, expecting both benchmarks to be comparable.

It is important to note that measuring the I/O throughput indirectly via introduced overhead masquerades the actual filesystem bandwidth B_{FS} which is always higher than the previously defined effective parallel throughput T_{eff} for raw, untransformed data as seen by the application. This is very important to keep in mind as the effective parallel throughput determines the application performance in most realistic scenarios.

As mentioned in Sect. 2.1, absolute I/O size during production runs quickly becomes a show-stopper. Compressing data streams on the fly seems to suggest itself as data reduction technique, either lossless or lossy, depending on application needs. In ADIOS, compression schemes are implemented transparently for the user as so-called data transforms. One would not only expect a reduction in data size but also an increase in effective bandwidth since the size of the compressed data S_C written to the filesystem is lowered by a compression ratio $f_C \equiv S_C/S \leq 1$ compared to the initial size S . We observed that this expectation could not be fulfilled using even the fastest compression algorithm implemented at the time in ADIOS, zlib, see Fig. 1.




Fig. 2. Actual filesystem throughput as seen by Atlas 2 (ORNL) during run no. 2489794 (Sep 23rd, 2015) on 16384 nodes according to user support (data: DDNTool, Splunk).

We therefore expanded our model to account for the time $t_{\text{reduce}} \equiv S/T_C$ it takes to reduce the data by compression or other means and copy it from an application-side buffer to an I/O library buffer. Up to now, data transforms in ADIOS are performed before starting to send the data off-node, while parallel HDF5 does not yet support data compression¹. In order to account for data reduction, Eq. (2) needs to be extended to add synchronous reduction overhead by

$$\begin{aligned}
 t_{\text{I/O}}^{\text{reduce}}(t_{\text{reduce}}) &\equiv t_{\text{prep}} + t_{\text{reduce}} + f_C \times t_{\text{offRAM}} \\
 &= t_{\text{prep}} + \frac{S}{T_C} + \frac{f_C \times S}{T_{\text{FS}}} \\
 &= t_{\text{prep}} + \frac{S}{T_C \times T_{\text{memcpy}}} + \frac{f_C \times S}{T_{\text{FS}} \times T_{\text{memcpy}}} \\
 f_C &\equiv \frac{S_C}{S} \quad T_C \equiv \frac{T_C}{T_{\text{memcpy}}} \quad T_{\text{FS}} \equiv \frac{T_{\text{FS}}}{T_{\text{memcpy}}}.
 \end{aligned} \tag{3}$$

T_C and T_{FS} characterize throughput for compression and filesystem writes, respectively, normalized to in-node memory copy throughput T_{memcpy} . We acknowledge that $t_{\text{reduce}} + f_C \times t_{\text{offRAM}}$ could in principle be lowered by copying the data to an I/O stage immediately and performing compression there, again within the limits of the discussion in Sect. 2.2.

Consequently, for a given normalized per-node filesystem throughput T_{FS} any data reduction algorithm C needs to fulfill the relation

$$\frac{T_C \times (1 - f_C)}{1 - T_C} > T_{\text{FS}} \tag{4}$$

in order to not only reduce data size by f_C but also perceived write time. This inequality arises from Eq. (3) assuming a reduce operation that is as fast as

¹ An experimental development preview with compression support in parallel HDF5 was announced after our measurements in February 2017.




Fig. 3. Weak scaling of PIconGPU with implemented I/O methods on Hypnos from 4 to 64 K20m GPUs (16 nodes). In contrast to Titan and Summit nodes, on Hypnos only two physical CPU cores are available per GPU, resulting in I/O performance with zlib and zstd [9] below the untransformed output.

possible by setting the second term of the sum $t_{\text{reduce}} \equiv t_{\text{memcpy}}$ and thus comparing $t_{\text{I/O}}^{\text{reduce}}(t_{\text{reduce}}) < t_{\text{I/O}}^{\text{reduce}}(t_{\text{memcpy}})$. The left-hand side of Eq. (4), which we call the *break-even threshold* for a given data transform algorithm and single (parallel) I/O stage, is discussed in greater detail in the following section.

In order to confirm this observation, we measured I/O performance on the K20 queue of the HZDR compute cluster Hypnos, see Fig. 3 (data points ‘no transform’ and ‘zlib’). Following Eq. (4) it should be even harder for a compression algorithm, lossless or lossy, to fulfill the requirements for break-even on Hypnos. Therefore, an improvement in the latter case will be automatically favorable for Titan or a Summit-like system.

3.3 Measurement of Compression Performance

In the interest of exploring feasible compression methods for PIconGPU data, we performed ex situ benchmarks on generated data. Visualized in Fig. 4, such a measurement directly allows a prediction for individual systems and user data when comparing to our model, Eq. 4.

PIconGPU currently only utilizes one host thread per GPU, so we decided to implement and explore compression throughput for blosc as an example for a multi-threaded algorithm and compare it to other, previously implemented compression algorithms. Blosc provides several bitshuffle pre-conditioners, which we found of great importance for floating-point compression performance in agreement with recent studies [7]. Further benchmarks with four threads on Hypnos’ K20 queue, limited to two host threads per GPU without oversubscription, indicated that on Hypnos application throughput would benefit from more physical CPU cores per GPU since the recent filesystem upgrade to GPFS.

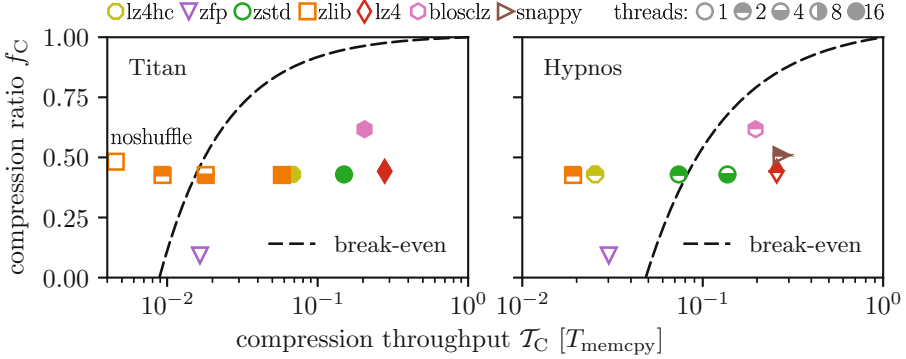


Fig. 4. Compression throughput T_C and ratio f_C measured on PIconGPU particle data (32 Bit floating point and integers). Lower f_C and higher T_C is better. All operations performed on contiguous, aligned, none-page-locked memory. The blosc [2] compression level is 1 (fast). From available pre-conditioners (none, shuffle, bitshuffle), the latter is shown due to the observed positive influence on f_C with small impact on T_C for floating point data which otherwise could not be compressed with LZ4 (v1.7.5) and snappy [13]. Zfp (v0.5.1) was used in fixed-precision mode with three uncompressed bits per scalar [19].

4 Analysis

Fully accelerator driven applications can use ‘the last 10% of system performance’ on the host side in order to trade compute performance for I/O latency. The Titan system provides up to 16 physical CPU cores per GPU and Summit is expected to allow for an order of magnitude higher parallelization on the host. This section explores the limits to data reduction methods in terms of data reduction ratio and throughput for an individual I/O stage independently of the method of data reduction and only exemplified for compression methods.

4.1 Overhead of Compression in Parallel I/O

From Eq. (3) the relative I/O performance ratio Γ when using data reduction instead of direct pass-through in an I/O stage follows as:

$$\Gamma \equiv \frac{t_{\text{I/O}}^{\text{reduce}}(t_{\text{reduce}})}{t_{\text{I/O}}^{\text{reduce}}(t_{\text{memcpy}})} = \frac{C_{\text{prep}} + \frac{f_C}{T_{\text{FS}}} + T_C^{-1}}{C_{\text{prep}} + T_{\text{FS}}^{-1} + 1} \quad (5)$$

$$C_{\text{prep}} \equiv \frac{t_{\text{prep}}}{S} \times T_{\text{memcpy}}$$

where we assume that the time for reducing the data $t_{\text{reduce}} \geq t_{\text{memcpy}}$ at minimum is as long as for copying data from node RAM to I/O buffer. It is clear that in terms of I/O throughput reduction algorithms are beneficial if $\Gamma < 1$ compared to I/O without reduction. Cases of $\Gamma \geq 1$ and $f_C < 1$ can still be

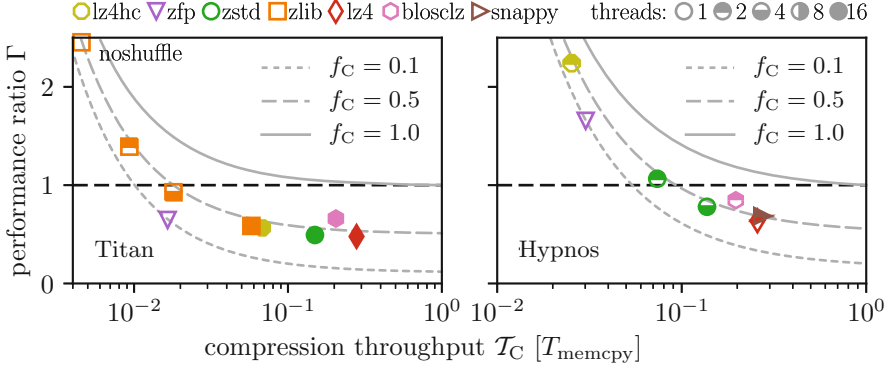


Fig. 5. Visualization of Eq. (5) predicting the relative I/O overhead $\Gamma > 1$ (gain $\Gamma \leq 1$) of compression during parallel I/O. The break-even threshold discriminates between feasible and overhead-adding compression algorithms at $\Gamma = 1$ (dashed line). Lower Γ and higher \mathcal{T}_C is better. Iso-compression lines for user data are plotted for individual systems (see Table 1) and compared to measured ex situ compression performance on PIConGPU user data (see Fig. 4).

relevant in case of limited disk space. Note, that decreasing C_{prep} would increase the gradient of Γ , but not affect the position \mathcal{T}_C for which we expect break-even.

Figure 5 shows the effect of threaded compression, keeping the compression ratio along iso-compression lines. Following the graph to the right, the higher the throughput of a compression algorithm the less importance it has on Γ compared to the compression ratio f_C . Thus, an important limit to Γ is the high-throughput limit $\mathcal{T}_C \rightarrow 1$ for fast compression algorithms below the break-even threshold. For such, the performance ratio over non-compressed I/O can barely be improved further via throughput but solely by compression ratio.

Exactly the opposite is true for any reduction algorithm with low throughput \mathcal{T}_C , to the left of the graph. Above the break-even threshold (dashed line at $\Gamma = 1$), data reduction quickly becomes impractical for medium to high-throughput tasks for a specific system, as the relatively wasted computing time never reaches $\Gamma \leq 1$ even for small f_C .

Following the last argument one can further derive from Eq. (4) with ‘perfect reduction’ $f_C \rightarrow 0$: For *any* given I/O stage with write and reduction throughput $\mathcal{T}_{\text{FS}}, \mathcal{T}_C$ the effective time an application spends in (synchronous) I/O can only be reduced, if the data reduction operation provides at least a throughput of

$$\frac{\mathcal{T}_C}{1 - \mathcal{T}_C} > \mathcal{T}_{\text{FS}}. \quad (6)$$

5 Summary and Outlook

We implemented and benchmarked parallel I/O methods on top of state-of-the-art I/O libraries for the massively parallel, fully-manycore driven, open source

PIC code PIConGPU. We outlined performance bottlenecks for medium to high-throughput applications in general and the possibility to overcome these with general data reduction techniques such as compression. We then derived and verified a scaling law that gives limits to expected application speed up when using data reduction schemes for medium- to high-throughput applications. With this we were able to derive a system- and application-specific break-even threshold that allows for predicting when reducing data is beneficial in terms of I/O throughput compared to I/O without reduction.

5.1 Compression Algorithms

For the special case of compression algorithms, future designs to soften I/O bottlenecks first and foremost need to improve throughput for floating point data. Even for a relatively large gap between local memory and filesystem throughput as on the current Titan system, many single-threaded compression algorithms that are still in use today do not fulfill the break-even threshold in Eq. (4).

Existing high-throughput compression algorithms would benefit from research improving the compression ratio f_C instead of throughput \mathcal{T}_C [7, 26]. This case is of importance since, due to high entropy in HPC applications' primary observables (e.g. floating point), only lossy compression algorithms are likely to bridge the upcoming throughput gaps between node-local high-bandwidth memory and storage accessible longer than application lifetime.

For ADIOS we proposed, implemented and benchmarked for the first time host-side multi-threaded transform methods as a feasible step to reach the break-even threshold. With that, we successfully traded unused compute performance within a heterogeneous application for overall I/O performance.

5.2 I/O Libraries

Burst-buffers are identified as enablers to reduce blocking time of the application caused by synchronous transformations within I/O libraries, but are vulnerable to backlog. Nonetheless, burst-buffers alone cannot cover the gap that will arise between expected I/O on system today and in the future. Further applications of burst-buffers are coupled multi-scale simulations, in situ processing and checkpointing and not in the scope of this paper.

Nevertheless, for both explorative-qualitative and medium- to high-throughput quantitative studies I/O libraries need to act now to provide transparent and easily programmable means for multi-stage I/O. For any practical application, the first I/O stage should immediately start with a maximum-throughput memory copy from user RAM to I/O buffer, ideally asynchronously, while later stages need to follow fully asynchronously. Copied memory (in unutilized RAM or burst-buffers) will need several off-node user-programmable transformations which are finally staged transparently through a subsequent non-blocking data reduction (compression) pipeline. In each I/O stage, the break-even threshold derived in this paper needs to be fulfilled or backlog will occur for successive outputs and the overall application will be throttled by that specific bottleneck. With deeper

memory hierarchies, user-programmability of stages will be a human bottle-neck and needs to be addressed with easy and fast turnaround APIs to design application- and study-specific stages, e.g. via Python/Numba.

In conclusion, introducing data reduction for I/O will be necessary because of limited medium to long term storage size expected for future systems. Our analysis and measurements show that even today one should however not expect I/O performance gains when using reduction. Parallelization of reduction algorithms is one way to gain overall I/O performance but requires compute resources in addition to those used by the application. Even for fully GPU accelerated applications one should not assume resources to be ‘free’ for I/O and analysis tasks, since loosely coupled application workflows and models that depend heavily on hardly-parallelizable aspects such as atomic data lookups will in the future be more widespread and compete for the exact same resources.

References

1. Abbasi, H., Wolf, M., Eisenhauer, G., Klasky, S., Schwan, K., Zheng, F.: Datastager: scalable data staging services for petascale applications. *Clust. Comput.* **13**(3), 277–290 (2010). doi:[10.1007/s10586-010-0135-6](https://doi.org/10.1007/s10586-010-0135-6)
2. Alted, F.: blosc 1.11.4-dev, March 2017. <https://github.com/Blosc/c-blosc>
3. Ayachit, U., Bauer, A., Geveci, B., O’Leary, P., Moreland, K., Fabian, N., Mauldin, J.: ParaView catalyst: enabling in situ data analysis and visualization. In: Proceedings of the First Workshop on In Situ Infrastructures for Enabling Extreme-Scale Analysis and Visualization, ISAV2015, pp. 25–29. ACM (2015). doi:[10.1145/2828612.2828624](https://doi.org/10.1145/2828612.2828624)
4. Bhimji, W., Bard, D., Romanus, M., Paul, D., Ovsyannikov, A., Friesen, B., Bryson, M., Correa, J., Lockwood, G.K., Tsulaia, V., et al.: Accelerating science with the NERSC burst buffer early user program. In: Proceedings of Cray Users Group (2016)
5. Birdsall, C., Langdon, A.: Plasma physics via computer simulation. The Adam Hilger series on plasma physics. McGraw-Hill, New York (1985). ISBN 9780070053717
6. Burau, H., Widera, R., Honig, W., Juckeland, G., Debus, A., Kluge, T., Schramm, U., Cowan, T.E., Sauerbrey, R., Bussmann, M.: PIConGPU: a fully relativistic particle-in-cell code for a gpu cluster. *IEEE Trans. Plasma Sci.* **38**(10), 2831–2839 (2010)
7. Burtscher, M., Mukka, H., Yang, A., Hesaaraki, F.: Real-time synthesis of compression algorithms for scientific data. In: SC16: International Conference for High Performance Computing, Networking, Storage and Analysis, pp. 264–275, November 2016. doi:[10.1109/SC.2016.22](https://doi.org/10.1109/SC.2016.22)
8. Bussmann, M., Burau, H., Cowan, T.E., Debus, A., Huebl, A., Juckeland, G., Kluge, T., Nagel, W.E., Pausch, R., Schmitt, F., Schramm, U., Schuchart, J., Widera, R.: Radiative signatures of the relativistic Kelvin-Helmholtz instability. In: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis, SC 2013, pp. 5:1–5:12. ACM (2013). doi:[10.1145/2503210.2504564](https://doi.org/10.1145/2503210.2504564)
9. Collet, Y., Skibinski, P., Terrell, N., Purcell, S.: Contributors: Zstandard (zstd) 1.1.4 - fast real-time compression algorithm, March 2017. <https://github.com/facebook/zstd>

10. Docan, C., Parashar, M., Klasky, S.: DataSpaces: an interaction and coordination framework or coupled simulation workflows. In: Proceedings of 19th International Symposium on High Performance and Distributed Computing (HPDC 2010), June 2010. doi:[10.1007/s10586-011-0162-y](https://doi.org/10.1007/s10586-011-0162-y)
11. Eckert, C.H.J.: Enhancements of the massively parallel memory allocator scatter-alloc and its adaption to the general interface mallocMC, October 2014. doi:[10.5281/zenodo.34461](https://doi.org/10.5281/zenodo.34461)
12. Grismayer, T., Alves, E., Fonseca, R., Silva, L.: dc-magnetic-field generation in unmagnetized shear flows. *Phys. Rev. Lett.* **111**, 015005 (2013). doi:[10.1103/PhysRevLett.111.015005](https://doi.org/10.1103/PhysRevLett.111.015005)
13. Gunderson, S.H., Evlogimenos, A.: Contributors: Snappy 1.1.1 - a fast compressor/decompressor (2011). <https://github.com/google/snappy>
14. Hockney, R., Eastwood, J.: *Computer Simulation Using Particles*. Taylor & Francis, Bristol (1988). ISBN: 9780852743928
15. Huebl, A., Lehe, R., Vay, J.L., Grote, D.P., Sbalzarini, I., Kuschel, S., Bussmann, M.: openPMD 1.0.0: a meta data standard for particle and mesh based data, November 2015. doi:[10.5281/zenodo.33624](https://doi.org/10.5281/zenodo.33624)
16. Huebl, A., Schmitt, F., Widera, R., Grund, A., Schumann, C., Eckert, C., Bukva, A., Pausch, R.: libSplash: 1.6.0: SerialDataCollector filename API, October 2016. doi:[10.5281/zenodo.163609](https://doi.org/10.5281/zenodo.163609)
17. Huebl, A., Widera, R., Grund, A., Pausch, R., Burau, H., Debus, A., Garten, M., Worpitz, B., Zenker, E., Winkler, F., Eckert, C., Tietze, S., Schneider, B., Knespel, M., Bussmann, M.: PIConGPU 0.2.4: Charge of bound electrons, openPMD axis range, manipulate by position, March 2017. doi:[10.5281/zenodo.346005](https://doi.org/10.5281/zenodo.346005)
18. Huebl, A., et al.: Supplementary materials: On the scalability of data reduction techniques in current and upcoming HPC systems from an application perspective, April 2017. doi:[10.5281/zenodo.545780](https://doi.org/10.5281/zenodo.545780)
19. Lindstrom, P.: Fixed-rate compressed floating-point arrays. *IEEE Trans. Vis. Comput. Graph.* **20**(12), 2674–2683 (2014). doi:[10.1109/TVCG.2014.2346458](https://doi.org/10.1109/TVCG.2014.2346458)
20. Liu, Q., Logan, J., Tian, Y., Abbasi, H., Podhorszki, N., Choi, J.Y., Klasky, S., Tchoua, R., Lofstead, J., Oldfield, R., et al.: Hello ADIOS: the challenges and lessons of developing leadership class I/O frameworks. *Concurr. Comput. Pract. Exp.* **26**(7), 1453–1473 (2014)
21. Matthes, A., Huebl, A., Widera, R., Grottel, S., Gumhold, S., Bussmann, M.: In situ, steerable, hardware-independent and data-structure agnostic visualization with ISAAC. *Supercomputing Frontiers and Innovations* **3**(4) (2016). <http://superfri.org/superfri/article/view/114>
22. McLay, R., James, D., Liu, S., Cazes, J., Barth, W.: A user-friendly approach for tuning parallel file operations. In: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC 2014, pp. 229–236. IEEE Press (2014). doi:[10.1109/SC.2014.24](https://doi.org/10.1109/SC.2014.24), <https://github.com/TACC/t3pio>
23. Meredith, J.S., Ahern, S., Pugmire, D., Sisneros, R.: EAVL: The extreme-scale analysis and visualization library. In: Childs, H., Kuhlen, T., Marton, F. (eds.) *Eurographics Symposium on Parallel Graphics and Visualization*. The Eurographics Association (2012). doi:[10.2312/EGPGV/EGPGV12/021-030](https://doi.org/10.2312/EGPGV/EGPGV12/021-030)
24. Meuer, H.W., Strohmaier, E., Dongarra, J., Simon, H., Meuer, M.: November 2016 — TOP500 Supercomputer Sites, June 2016. <https://www.top500.org/lists/2016/11/>. Accessed 22 Mar 2017
25. Corporation, N.: NVIDIA IndeX 1.4. <https://developer.nvidia.com/index>

26. Tao, D., Sheng, D., Chen, Z., Cappello, F.: Significantly improving lossy compression for scientific data sets based on multidimensional prediction and error-controlled quantization. In: IPDPS 2017: Proceedings of the 31th IEEE International Parallel and Distributed Processing Symposium, May 2017
27. The HDF Group: Hierarchical data format version 5 (C-API: 1.8.14) (2000–2017). <http://www.hdfgroup.org/HDF5>
28. Whitlock, B., Favre, J.M., Meredith, J.S.: Parallel in situ coupling of simulation with a fully featured visualization system. In: Kuhlen, T., Pajarola, R., Zhou, K. (eds.) Eurographics Symposium on Parallel Graphics and Visualization. The Eurographics Association (2011). doi:[10.2312/EGPGV/EGPGV11/101-109](https://doi.org/10.2312/EGPGV/EGPGV11/101-109)
29. Zenker, E., Widera, R., Huebl, A., Juckeland, G., Knüpfer, A., Nagel, W.E., Bussmann, M.: Performance-portable many-core plasma simulations: porting PIconGPU to openpower and beyond. In: Taufer, M., Mohr, B., Kunkel, J.M. (eds.) ISC High Performance 2016. LNCS, vol. 9945, pp. 293–301. Springer, Cham (2016). doi:[10.1007/978-3-319-46079-6_21](https://doi.org/10.1007/978-3-319-46079-6_21)
30. Zenker, E., Worpitz, B., Widera, R., Huebl, A., Juckeland, G., Knüpfer, A., Nagel, W.E., Bussmann, M.: Alpaka-an abstraction library for parallel kernel acceleration. In: 2016 IEEE International on Parallel and Distributed Processing Symposium Workshops, pp. 631–640. IEEE (2016)

High Performance Computing

ISC High Performance 2017 International Workshops,
DRBSD, ExaComm, HCPM, HPC-IODC, IWOPH, IXPUG,
P³MA, VHPC, Visualization at Scale, WOPSSS,
Frankfurt, Germany, June 18-22, 2017, Revised Selected
Papers

Kunkel, J.; Yokota, R.; Taufer, M.; Shalf, J. (Eds.)

2017, XX, 743 p. 266 illus., Softcover

ISBN: 978-3-319-67629-6