Exploiting dynamic sparse matrices for performance portable linear algebra operations

Christodoulos Stylianou
EPCC, The University of Edinburgh
Edinburgh, UK
c.stylianou@ed.ac.uk

Michèle Weiland
EPCC, The University of Edinburgh
Edinburgh, UK
m.weiland@epcc.ed.ac.uk

Abstract—**Sparse matrices and linear algebra are at the heart of scientific simulations. More than 70 sparse matrix storage formats have been developed over the years, targeting a wide range of hardware architectures and matrix types. Each format is developed to exploit the particular strengths of an architecture, or the specific sparsity patterns of matrices, and the choice of the right format can be crucial in order to achieve optimal performance. The adoption of dynamic sparse matrices that can change the underlying data-structure to match the computation at runtime without introducing prohibitive overheads has the potential of optimizing performance through dynamic format selection.**

In this paper, we introduce *Morpheus*, a library that provides an efficient abstraction for dynamic sparse matrices. The adoption of dynamic matrices aims to improve the productivity of developers and end-users who do not need to know and understand the implementation specifics of the different formats available, but still want to take advantage of the optimization opportunity to improve the performance of their applications. We demonstrate that by porting HPCG to use *Morpheus*, and without further code changes, 1) HPCG can now target heterogeneous environments and 2) the performance of the Sparse Matrix-Vector Multiplication (SpMV) kernel is improved up to $2.5\times$ and $7\times$ on CPUs and GPUs respectively, through runtime selection of the best format on each MPI process.

Index Terms—**sparse matrix storage formats, generic programming, dynamic matrices, performance portability, productivity**

I. Introduction

Sparse matrices (i.e. matrices that mostly consist of zeros) are an essential concept in computational science and engineering. [1] The computational and memory savings that can be gained from exploiting the sparsity structure of a matrix have driven the development of software specific to sparse linear algebra. As the majority of elements in a sparse matrix are zeros, a special focus was given to defining formats (i.e. data structures) to enable the efficient storage of, and access to, all non-zero elements without the need to explicitly store the zeros as well. Sparse matrix storage formats reduce the memory footprint of the matrix, eliminate redundant computations and allow for larger problems to be processed.

According to Filippone et al. [2], more than 70 such sparse matrix storage formats have been developed over the years to address not only the various types of matrices that result from different discretization methods, but also the evolution of hardware towards multi- and many-core processors and

accelerators. Literature shows that there is no single format that can perform optimally across all different kinds of matrices and types of hardware [2]–[6]. Compared to their dense counterparts, operations on sparse matrices are generally known to be memory-bandwidth bound because of the need to retrieve matrix values via the index information, often resulting in indirect memory accesses and poor cache reuse.

In many numerical applications, computing the solution to linear systems (through performing sparse matrix-vector multiplications) dominates the runtime. The (iterative) solvers can be optimized through the use of storage formats that exploit hardware capabilities given the sparsity pattern (i.e. the distribution of non-zero entries) of the matrix and the operation to be performed. To facilitate the selection of the best format for a given matrix and target architecture, auto-tuners have been developed [3], [7], [8]. These auto-tuners demonstrate the impact of selecting a suitable format at runtime to optimise the performance of iterative solvers on single node multi- or many-core systems and they provide mechanisms to automate the process.

In this work, we develop an abstraction for sparse matrices that enables switching to different storage formats at runtime. This abstraction provides optimization opportunities, by adapting the data-structure given the operation, target architecture and sparsity pattern of the matrix, but without introducing noticeable overheads. We introduce *Morpheus* [9], a library that supports the runtime-switching of matrix storage formats by implementing sparse matrices using a single dynamic "abstract" format and providing a transparent mechanism to switch between different implementations. A detailed description of *Morpheus* is given in Section III. We use the High Performance Conjugate Gradients (HPCG) benchmark [10] as a test case to represent a common numerical workload, and we extend it to support dynamic sparse matrix storage format switching and multiple hardware backends that makes it possible to target CPU or GPU architectures without any further code modifications.

In summary, our contributions are:
- We show that using the abstract sparse matrix format representation provided by *Morpheus* [9], our library of sparse matrix storage formats, incurs no significant runtime overheads. Indeed, without any further code changes and by transparently switching to a different

format, performance of a solver can often improve (see Section V-C).

- We describe a general process for incremental porting of an application to use *Morpheus* (see Section IV-A) and, as validation, we provide an example of the process for HPCG (see Sections IV-B).
- We demonstrate that for an optimal storage format, a specific problem in a shared memory setting can introduce irregularities in the sparsity pattern and perform significantly worse for the same format in a distributed memory setting. We show that this issue can be solved by splitting the local system matrix into a local (regular) and remote (irregular) part, each with possibly different formats, and investigate the performance benefits from doing so.
- We show that by selecting the optimal format for local and remote matrices per process the application can achieve and retain the best possible runtime performance. After enabling HPCG to use *Morpheus*, we are able to perform the SpMV routine up to $2.5\times$ faster on a multi-core CPU system and up to $7\times$ faster on CPU+GPU nodes, using the same source code and without any further modifications.

II. BACKGROUND AND MOTIVATION

A. Sparse Matrix Storage Formats

Dense matrices store all the coefficients of a matrix explicitly in memory, usually by using a two-dimensional linear mapping. The coefficient $A(i, j)$ of a dense $M \times N$ matrix is stored at position $(i \times N) + j$ of a linear array (assuming row-major storage ordering). Sparse matrix formats on the other hand exploit the property that the majority of sparse matrix coefficients are zeros by not explicitly storing those values. As a result, the direct link between the index pair (i, j) and the position of the coefficient in memory is lost. A sparse matrix storage format aims to rebuild this link using auxiliary index information. The cost of calculating this indirection has an impact on the performance of sparse matrix computations.

A large number of sparse matrix storage formats have been developed over the years, each with different storage requirements, computational characteristics, and methods of modifying the entries of the matrix. [11] As a result, determining which format might perform best given an operation is not a trivial task as multiple factors determine the overall performance. No single storage format is able to exploit the matrix structure and perform optimally across multiple operations, or indeed across multiple target architectures. [2]

The most basic and well-known formats are Coordinate (COO) and Compressed Sparse Row (CSR). Both are considered *general purpose* formats, suitable for a broad range of matrices of arbitrary sparsity patterns and target architectures. COO uses three arrays, where each non-zero element is explicitly stored together with its column and row indices, with no guarantees imposed in the ordering of the elements. CSR also explicitly stores the column indices and non-zero values, but it uses an array of pointers to mark the boundaries

of each row, reducing the memory footprint of the format by essentially compressing the row indices. As the row pointers are used to represent the position of the first non-zero element in each row, and the last entry shows the total number of non-zeros in the matrix, CSR naturally also imposes an ordering across rows, but not within each row.

Specific purpose formats aim to address the characteristics of specific classes of matrices and are usually designed to perform optimally with a target architecture in mind. For example, the Diagonal (DIA) format was originally designed for vector processors and is suitable for regular sparsity patterns. DIA uses a two-dimensional array, where each column holds the coefficients of a diagonal of the matrix, and an integer offset array keeping track of where each diagonal starts. Therefore, DIA format is suitable for matrices with structures that dominate along the diagonals, such as banded matrices resulting from discretization methods like Finite Differences Method (FDM), and vector-like architectures, such as GPUs.

Note that the formats above, or a combination of them, constitute a basis from where many other formats are derived. However, it is clear that the underlying data structure across formats can vary significantly. As a result, accessing and manipulating entries in each format can result in different memory access patterns, costs and interfaces amongst formats.

B. HPCG

HPCG [10] is a benchmark that measures the performance of HPC systems by solving the Poisson differential equation on a regular 3D grid, discretized with a 27-point stencil. It uses the Preconditioned Conjugate Gradient (PCG) algorithm with a symmetric Gauss-Seidel [12] as a preconditioner, and includes the following computations: sparse matrix-vector multiplications (SpMVs); vector updates; global dot products; a local symmetric Gauss-Seidel smoother (including a sparse triangular solve); and multi-grid (MG) preconditioned solvers. The benchmark compares and validates the performance of a user-tunable optimized version of the PCG algorithm against a reference in the following phases:

1) *Problem setup*: Constructs the synthetic problem by creating the geometry and linear system.
2) *Reference timing*: Measures the time taken to run the SpMV and MG reference implementations and the time to solution for the reference Conjugate Gradient (CG) solver.
3) *Problem Optimization setup*: Configures the user defined data structures to be used in the optimized problem.
4) *Validation and Verification*: Checks that the optimized problem has returned the expected results.
5) *Optimized problem timing*: Measures the time to solution for the optimized CG solver.

We choose HPCG here because, not only it is a widely accepted and well-understood benchmark, but also because it is designed to use optimized implementations (in our case provided by *Morpheus*) of the linear algebra computations and compare them against a reference baseline. Note that the reference implementation of HPCG can be built with MPI

for distributed parallelism and/or OpenMP for shared memory parallelism.

The algorithm chosen for solving the system of equations is not only limited by the floating point performance but also heavily relies on the performance of the memory system and to some extent the underlying network. Note that similar computations and data access patterns are found in many scientific calculations. Therefore, HPCG provides a good performance assessment that is representative of the performance of many scientific applications. The performance bottleneck in HPCG boils down to the sparse operations that are carried out at every step of the iterative solver, i.e the SpMV and the Gauss-Seidel smoother.

C. Motivation

Scientific codes that use sparse matrices are often built around a single general-purpose storage format. However different applications and problems, including those from new application domains, will exhibit different matrix sparsity patterns. In addition, the continuous evolution of hardware introduces architectures with new characteristics. For software to remain performant throughout its lifetime, algorithms must be adjusted to exploit these new architectures and represent sparsity patterns efficiently. Previous work (such as [3], [7], [8]) has demonstrated the importance of choosing the correct storage format in order to achieve optimal performance. However, only the impact on performance at the shared-memory was considered, without providing performance results at a distributed level. In a distributed environment, the global matrix is partitioned into smaller local matrices that are each assigned to a different process. As a result, the sparsity pattern of each local matrix can differ significantly. In this case, choosing a single storage format based on metrics that arise from the global matrix could result in noticeable load imbalances. This observation motivates a *per-process* selection of the optimum format, with the possibility of using different formats across processes. *Morpheus* supports both single and multi-format runtime configurations and as such can be used to compare both strategies and assess their performance and viability.

III. MORPHEUS - A LIBRARY FOR DYNAMIC SPARSE MATRICES

Adopting an abstraction that can change the underlying sparse matrix storage format efficiently at runtime has the potential of offering performance benefits without invasive code modifications. To achieve this, we have developed *Morpheus*[1], a C++ library that offers transparent, efficient and low overhead run-time switching between different sparse matrix storage formats across different target architectures. By abstracting the different formats under a single *dynamic* format and encapsulating the internal implementation details, end-users are left with a simple and intuitive interface that abstracts away the complexities. At the same time, performance can be

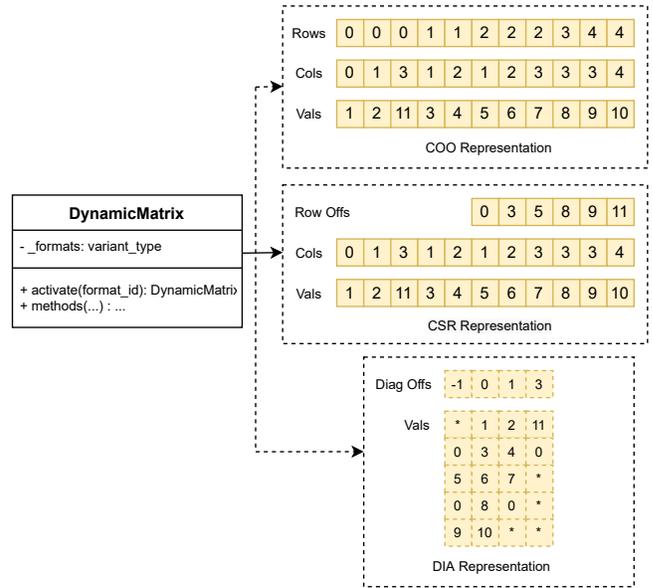[1]Available at: https://github.com/morpheus-org/morpheus.git

Fig. 1: `DynamicMatrix` Container with CSR representation as its current active type (solid arrow). Other possible active types are represented via dashed arrows and can be switched to during run-time using the *activate()* method.

optimised by transparently switching to the "best" format for each operation, sparsity pattern and target hardware, without any code modifications.

The *Morpheus* library follows a functional design and separates the data structures (containers) from the functions (algorithms), with the algorithms acting on the containers. To enable support for various hardware platforms and memory hierarchies, we are adopting the three core abstractions concepts offered by Kokkos [13]:

- *Execution Space*: Specifies *where* the code will be executed. Examples include the CPU and GPU cores.
- *Memory Space*: Specifies *where* the data will reside in memory, enabling exploitation of the different characteristics of various types of memory, such as non-volatile memory or High Bandwidth Memory (HBM).
- *Memory Layout*: Specifies *how* the data will reside in memory, enabling efficient data access patterns in algorithm and architecture dependent optimisations.

The adoption of these abstractions facilitates the formulation of generic algorithms and containers that can be efficiently mapped to the different types of architectures and memory hierarchies that are currently available, but also enable the addition of future developments.

A. Defining the data structures

Currently, *Morpheus* supports three containers for sparse matrix storage formats (*CooMatrix*, *CsrMatrix* and *DiaMatrix*), hereafter also called concrete formats, and two containers representing dense formats (*DenseMatrix* and *DenseVector*). Each container can be seen as a different type, uniformly

parameterized by the type of the values and indices it holds, as well as where and how it resides in memory. Note that all the aforementioned containers are resolved at compile-time, and they can potentially each have a different interface.

To enable dynamic switching between formats we have introduced the concept of a *DynamicMatrix*, a container acting as a composition of all the supported sparse containers. The *DynamicMatrix* follows the same semantics and is parameterized in the same way as the concrete formats. A high-level overview of the *DynamicMatrix* container is shown in Figure 1. Through its own interface and the use of the *State* pattern [14], which allows an object to alter its behaviour when its internal state changes, the *DynamicMatrix* provides a unified way of interacting with all the supported storage formats. In other words, at any given moment it can hold any of the available storage formats participating in the composition and can switch to a different one by changing its active state. We allow the active state to be changed through the `activate()` member function by selecting one of the available *enum* values associated with each of the available storage formats or their equivalent index.

B. Data management

Each container is responsible for acquiring and releasing its own resources when it goes out of scope in order to avoid memory leaks. Communication and data exchange between containers is achieved using three different types of *copy* concepts, each with a different set of requirements to be met and each with a different associated cost.

Shallow Copy implies that no actual copy or data transfer is performed, but that instead containers share resources, with deallocation taking place once all containers sharing those same resource go out of scope. *Shallow* copy has the lowest cost of the three copy concepts we support, however to succeed it requires that the two containers are of the same type (i.e the same format, type of values and indices, memory space and layout) otherwise a compile-time error is triggered.

Deep Copy performs a bit-wise copy of the data from a source to a destination container, and no shared state is maintained between the containers. The cost of such an operation is much higher compared to a shallow copy and increases as the size of the data increases, therefore such an operation has to be as efficient as possible. As a consequence, for a deep copy to be issued the source and destination containers must be compatible. In other words, the two containers must have the same format, type of values and indices with identical memory layout and alignment although the memory space can differ ensuring that the deep copy is transformed into a *memcopy* operation. At the same time, this also holds for data transfers across different spaces. We extend the mirroring interface of *Kokkos* to support sparse matrices, making sure that we can create and allocate compatible containers.

Convert performs an element-wise copy between two containers. This operation requires only that the two containers are in the same memory space, without imposing any further restrictions on the format, the type of values and indices or

the layout and alignment. Therefore, this enables conversions from one sparse format to another, but also between containers of the same format that are not compatible, although at a higher cost compared to a deep copy operation. To avoid the need to provide one conversion implementation for each pair of supported storage formats, which could rapidly become intractable, we use a proxy format policy where COO format acts as an intermediate format (i.e. all conversions involve going from the existing format *to* COO and *from* COO to a new format). Note that this is simply a pragmatic choice in order to control the complexity of the *Morpheus* code base; creating specific optimised conversion operations (say from DIA directly to CSR) is of course possible.

It is also worth pointing out that in the case of *shallow* and *deep* copy operations, the same semantics apply to both the concrete and dynamic containers. The only difference is that when at least one of the containers participating in the operation is a *DynamicMatrix*, its active state must also match the other container's state.

C. Supporting heterogeneous platforms

On heterogeneous systems, it might be necessary to perform data transfers across memory spaces. We assume that we have compute units responsible for general housekeeping (host) and compute units for performing the core computation (device). For example, on a CPU+GPU system, the CPU acts as the host and is responsible for preparing and transfering the data to the device (in this case the GPU), which in turn will perform the computation.

We adopt a host-device model to support heterogeneous platforms. Containers in *Morpheus* are by default assumed to live on the device space and offer a *HostMirror* type that represents an equivalent and compatible mirror container residing in the *HostSpace*. Through the mirroring interface we can allocate host containers to match the size of their device counterpart and the data transfers from and to the two types of containers are managed performing deep copies. However, in the case where the device container already resides on *HostSpace* (for example if only the CPU backend is enabled) both the device and host containers are the same hence performing deep copies between the two would be a redundant and expensive operation. To avoid such overheads, *Morpheus* is intelligent enough to transform the deep copy operations in shallow copies, without user having to do any code modifications.

D. Supporting different algorithms and programming models

The algorithms supported in *Morpheus* are exposed to the user through a generic high-level interface. Each algorithm has a unified interface such that it can be used in the same way across different containers. To specify where the algorithm will be executed, an *ExecSpace* parameter must be provided, which must be a valid execution space provided by *Kokkos*. Using compile-time introspection, *Morpheus* selects and dispatches the appropriate algorithm implementation for the containers and execution space provided by the user.

In the case where at least one of the containers is a *DynamicMatrix*, the same high-level interface is still used as *DynamicMatrix* follows the same semantics as the concrete formats. However, changing the active state of the *DynamicMatrix* at run-time would result in a different algorithm implementation being dispatched under the hood. We achieve this by using the *Visitor* pattern [14], where a "visitor"operation is responsible for inspecting the current state of the variant and for dispatching the correct routine. Note that because *DynamicMatrix* knows all the active states it can switch to at compile-time, all the versions of the algorithm are generated by the compiler *a priori* resulting in efficient run-time dispatches.

Currently, we support algorithms for serial, OpenMP and CUDA backend environments, with each algorithm kernel explicitly implemented for each backend. We also provide a *GenericSpace* wrapper that will dispatch the generic version of the algorithm written in *Kokkos*. For each matrix container we support basic operations such as `SpMV`, `diagonal update` and `extraction`. In addition, we provide algorithms such as `dot`, `WAXPBY`, `reduction` and `scan` for the *DenseVector* container. The current algorithms we support are commonly found in scientific applications that use iterative solvers and represent a minimum set of functionalities needed for us to evaluate *Morpheus*. For more complicated workloads, for example those using preconditioners, further algorithms will need to be implemented, although all will follow the same interface and design principles as discussed above.

IV. INTEGRATING MORPHEUS INTO AN EXISTING APPLICATION

The main goal of *Morpheus* is to increase the end-user's productivity and the performance of their applications through efficient format switching and selection without the user having to dive into the specifics of each supported storage format. For that to be possible, *Morpheus* needs to be able to be integrated in another application following a straightforward and incremental porting process, which is described below. For completeness, we also describe our approach to adding *Morpheus* to *HPCG* and provide details on the challenges and issues faced as well as how our library can help future-proof applications and benchmarks.

A. Porting process

Integrating *Morpheus* into an existing application can be done incrementally in 3 steps, as described below.

Step 1: Converting user-defined data structures. In order to use the algorithms provided by *Morpheus* we must first convert any user-defined data structures to the containers supported by *Morpheus*. By default, a container is responsible for managing (i.e. allocating) its own resources, however it can also be specified to be "unmanaged", i.e the allocation has to be provided by the user during construction. Consequently, we can convert user-defined data structures to *Morpheus* containers by passing an allocation that is used by the structure to an unmanaged container. As a result, both share the same

resources and the container will be aliasing the original data structure. Any updates will be directly reflected on both.

For array-like data structures allocated by the user creating an unmanaged *Morpheus* container is particularly useful. However, in the case of sparse matrices, the user-defined structure potentially has a completely different representation in memory from what *Morpheus* supports. As a result, constructing a *Morpheus* container will be only possible through an element-wise conversion provided by the user.

After the user-defined data structures have been converted into *Morpheus* containers, the *Morpheus* algorithms can now be used. However, at this stage only the CPU backends can be used as the GPU backend must first be enabled as it requires extra data management (as described in Step 2).

Step 2: Enabling GPU support. *Morpheus* will not handle any data transfers between two different memory spaces, therefore it is up to the user to manage data transfers from the CPU and GPU. However, *Morpheus* provides a mirroring interface and copy semantics as described in Section III-C that allows the user to target heterogeneous platforms and handle data transfers in a constructive way.

During this step, the user has to decide which of the containers will be used in an algorithm and which will be used for general housekeeping. Consequently, any container that will be used in an algorithm will be assumed to reside in the device space and any access to its data outside of the algorithm will have to be made through its equivalent *HostMirror* container. After the user handles the data transfers between device containers and *HostMirror* containers, the application code can now run both on CPU and GPU without any further modifications and using a single source code.

Step 3: Enabling dynamic switching. At this point, the application is able to run with *Morpheus* both on CPU and GPU backends, however no dynamic format switching functionality is available yet. In order to enable run-time polymorphism to switch to different formats the user needs to convert the concrete sparse matrix container into a *DynamicMatrix*. This is facilitated through the *convert* routine, which will switch and convert to the desirable storage format representation. No further changes are required as algorithms for both concrete and dynamic matrices have the same high-level interface. The run time switching functionality therefore allows the application to not only be future proof in terms of the hardware it can target, but also in terms of the sparsity patterns it can efficiently process as more storage formats are added to *Morpheus*.

B. Example: HPCG with Morpheus

The HPCG benchmark, described in detail in Section II-B, was chosen as an exemplar application to test *Morpheus* as it is a representative (albeit simplified) version of a real-life workload, with well understood performance characteristics. The goal is to create a *Morpheus*-enabled version of HPCG[2] to 1) confirm the soundness of the porting process, 2) identify

---

[2]Available at: https://github.com/morpheus-org/morpheus-hpcg.git

any obstacles and pitfalls and 3) assess the performance impact of using *Morpheus* as part of HPCG. Note that we concentrate on the main computational bottleneck of the application, the SpMV operation, and therefore for simplicity we disable the preconditioner step.

The reference implementation of HPCG benchmark runs on CPUs only and can be configured to use MPI and OpenMP. The benchmark has four main data structures: *Vector*, *Sparse-Matrix*, *CGData* and *MGData*. Here, we are interested in the first three (as preconditioning is disabled). The sections below describe the code changes that are required to port the reference HPCG implementation to support *Morpheus*.

1) The Vector and CGData data structures: The first step is to transform all the vector-vector operations in HPCG to the equivalent algorithms from *Morpheus*. To do so, we need to transform the *Vector* data structure to also hold the *DenseVector* container. Because the memory layout of both *Vector* and *DenseVector* is the same, we can construct an unmanaged *DenseVector* using the existing memory allocations. Therefore, during the `Problem Optimization Setup` phase of HPCG, for every vector (i.e. b, x and $xexact$) used in the `Optimized problem timing` phase, we initialize the unmanaged *DenseVector* by passing it the allocation to the data.

Since both *Vector* and *DenseVector* are now aliasing the same memory, they can be used interchangeably throughout the application without any need for further data management. We can now update HPCG's `ComputeDot` and `ComputeWAXPBY` functions to invoke the *Morpheus* algorithms for `dot` and `WAXPBY`.

2) The SparseMatrix data structure: Changing the *Sparse-Matrix* data structure to also hold one of the sparse matrix containers offered by *Morpheus* requires an element-wise conversion between HPCG's format of choice and the format of choice from *Morpheus*, which can be any of the formats that are supported as we are not limited to a single format, as is the case for other approaches like SMAT [7]. However, since HPCG uses a variation of CSR format it makes sense that the format of choice for *Morpheus* container is *CsrMatrix*.

During the `Problem Optimization Setup` phase we perform the conversion between the *SparseMatrix* and the *Morpheus* container for the system matrix A. Note that the two matrix representations live in different allocations and the user needs to ensure data consistency. However, since HPCG builds the system matrix in the `Problem Setup` phase and there is no change to it after that, both matrices will have the same data during the `Optimized problem timing` phase.

At this point, the *ComputeSPMV* routine in HPCG can be modified to invoke the *multiply* algorithm from *Morpheus*. Note that although the computation in the iterative solver returns the same results in the *Morpheus*-enabled HPCG, as it stands the benchmark will report invalid results. HPCG performs tests to ensure the solver works as expected and one of tests (`TestCG`) involves modifying the elements of the matrix diagonal. To address this issue, we have to update

the diagonal of the *Morpheus* container during the test by invoking the equivalent algorithm.

3) Enabling GPU support: During the `Problem Optimization Setup` phase, for every *DenseVector* and *CsrMatrix* container the equivalent *HostMirror* container needs to be created through the *Morpheus* mirroring interface, which will be initialized using the data from the HPCG data structures. The data will be loaded onto the device via a deep copy from the *HostMirror* container.

Since the benchmark repeats multiple *CG* runs, and within each run updates the vector data, these data transfers need to be setup manually by the user using the *Morpheus* copy interface, i.e. any `CopyVector` call needs to be substituted with `Morpheus::copy`. In the same way, every time a vector is zeroed using `ZeroVector` operation, the `Morpheus::assign` routine must be used instead. This will ensure that irrespective of which backend is enabled, the data will be handled properly for the end results to be valid and all from the same source code.

If MPI is enabled, HPCG assumes that for the `ComputeSPMV` operation the source vector is distributed across processes, therefore every time prior to performing the operation the remote vector parts have to be exchanged. When the GPU backend is enabled each *MPI* process is responsible for a GPU and the vector is distributed across the different GPUs. This means that the `ExchangeHalo` routine must be adapted to copy the data from one GPU to another via the CPU again using the *Morpheus* copy interface.

Once this step has been completed, the benchmark is able to work with any backend that is supported now (or will be in the future) that follows the *Host-Device Model* without any further source code modifications.

4) Enabling Dynamic Switching: Up until this point, the benchmark is using the *CsrMatrix* container directly. If we want to use the dynamic capabilities of *Morpheus*, the *DynamicMatrix* container needs to be used instead. We change the container held by the *SparseMatrix* to be the *DynamicMatrix* and during the `Problem Optimization Setup` we assign the converted *CsrMatrix* to the *DynamicMatrix*. This will result in a *DynamicMatrix* with its active state being the CSR format. To change the active state (i.e. switch to a different format) we can now apply an in-place conversion using the *convert* interface from *Morpheus* and selecting the format to use at run time, either through the command line or an input file. This *Morpheus*-enabled version of HPCG will be able to take advantage of any new storage formats to optimise the performance without any code modifications.

V. RESULTS AND EVALUATION

A. Experiments

Using the dynamic capabilities of *Morpheus* as part of an application should not introduce significant overheads. Note that *Morpheus* knows *a priori* which format representations are supported and as a result, the compiler can generate a complete set of the possible algorithms that can be efficiently dispatched at runtime. To assess and quantify the overheads introduced
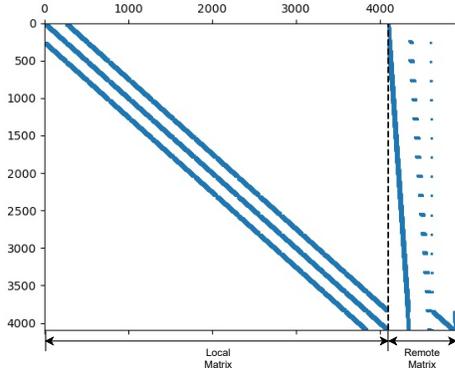
Fig. 2: Image representation of the sparsity pattern of a 4096x4096 local system matrix at rank 0. The dashed line indicates which part of the system matrix belongs to the local matrix and which to the remote, demonstrating how the matrix will be divided into the two parts. Note that when MPI is disabled the remote matrix part disappears and the sparsity pattern of the local system matrix is regular and concentrated around the diagonals.

to HPCG after adding support for *Morpheus*, we compare the runtime of the original HPCG with respect to the *Morpheus*-enabled HPCG with the active state of the *DynamicMatrix* set to CSR. We perform this comparison for a set of compilers.

Furthermore, since HPCG solves a Poisson differential equation on a regular 3D grid discretized with a 27-point stencil, the system matrix is expected to be regular with non-zeros concentrated around the diagonals. For that type of sparsity pattern the performance of the SpMV kernel is expected to be improved by switching to DIA format instead of CSR, as it will enable contiguous accesses to vector x. However, once MPI is enabled and in order to optimize the communication of the remote vector elements, on each process the local system matrix is conceptually divided into a local and a remote part, as shown in Figure 2. With the remote part containing the matrix elements that interact with the remote part of the vector, the system matrix is effectively transformed to an irregular matrix and the selection of DIA format would result in excessive zero-padding and potentially run out of memory. In addition to changing the regularity of the matrix, note that from square matrix it is now rectangular.

We measure the single node performance (i.e. with MPI disabled) of each supported format, over a number of problem sizes, systems and architectures in order to determine the performance benefit from switching to the best available format (see Section V-D). In addition, in Section V-E we split the local and remote matrix parts and by switching to the available format combinations for each part we evaluate the scaling performance application with respect to the performance of the reference HPCG implementation.

B. Compute node architectures

All experiments were carried out on the ARCHER2 and Cirrus supercomputers; their compute node architectures are

described in Table I. Note that Cirrus has both CPU-only as well as CPU+GPU nodes. Also note that for each experiment we explicitly state the compilers used, however the compiler optimization flag `-O3` is always used. For each experiment build, run and data processing scripts for the machines used are available online.[3]

PLATFORM	CIRRUS (GPU NODE)	CIRRUS (CPU NODE)	ARCHER2
CPU	INTEL XEON GOLD 6248 (x2)	INTEL XEON E5-2695 (x2)	AMD EPYC 7742 (x2)
GPU	NVIDIA TESLA V100 SXM2-16GB (x4)	N/A	N/A

TABLE I: Node configurations for the systems used in the experiments.

C. Overhead Comparison

It is vital for an abstraction framework to not introduce overheads that might negate any performance gains. In order to quantify any overheads introduced by *Morpheus* we compare the performance of SpMV as used by the *Morpheus*-HPCG (with *DynamicMatrix* set to CSR) to the original HPCG implementation, both using the same backends (Serial/OpenMP) on ARCHER2. The experiment is repeated over a set of per-core problem sizes. Note that for the OpenMP backend we are limiting the experiment to a single NUMA region (16 cores on AMD EPYC 7742) to avoid any NUMA effects.

The efficiency of the dynamic dispatch mechanism used by *Morpheus* is implementation specific and overheads might vary across compilers. We therefore repeated the experiments across all available compilers on ARCHER2: GNU 9.3.0, 10.3.0 and 11.2.0; AOCC 2.2.0 and 3.0.0; and CRAY-clang 11.0.4 and 12.0.3.

Figure 3 shows that the SpMV runtime ratio of the *Morpheus*-HPCG and the the original implementation is concentrated slightly below 1. This indicates that no significant overheads are introduced by *Morpheus* for both the Serial and OpenMP backends. On the contrary, the runtime is slightly improved ($\approx 5\%$). The slight improvement in performance can be attributed to the fact that the original implementation uses a variation of CSR as the format of choice, where the values and column indices are accessed by traversing pointers-of-pointers, whilst the *Morpheus*-enabled version uses the traditional CSR.

D. Single Node Performance

By integrating *Morpheus* into an application users can take advantage of its dynamic switching capabilities and target heterogeneous environments. In this experiment, we show the performance gains that can be achieved by changing the active type of the *DynamicMatrix*.

In order to quantify the performance gain, we run the *Morpheus*-HPCG with MPI disabled and measure the SpMV runtime. For each run, we change the active state of the

---

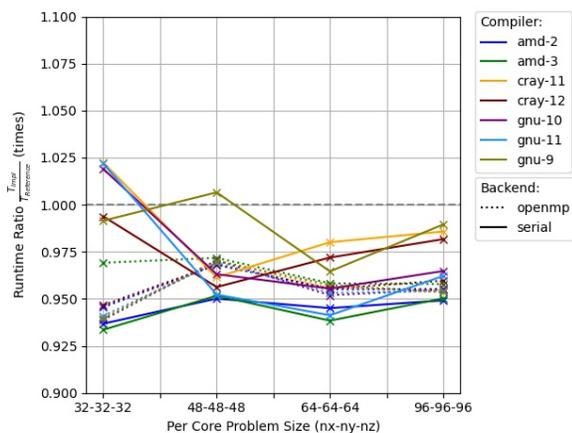[3]Available at: https://github.com/morpheus-org/morpheus-benchmarks.git

Fig. 3: Overheads resulting from adding *Morpheus* to HPCG. We compute the runtime ratio of SpMV as computed by the *Morpheus*-HPCG (*DynamicMatrix* set to CSR) and the original HPCG implementation on ARCHER2. A ratio above 1 indicates that overheads are introduced, a ratio below 1 indicates a performance improvement.

DynamicMatrix to one of the supported formats (i.e. COO, CSR, DIA). The runtime ratio is computed by dividing the runtime of the reference state (CSR) with the actual state we are measuring for. The experiment is repeated for a set of problem sizes, available compilers and backends, both on ARCHER2 and Cirrus.

On ARCHER2, one compiler from each vendor was used (GNU 10.3.0, AOCC 3.0.0 and CRAY-clang 12.0.3) and the backend of choice was OpenMP. To avoid NUMA effects, only a single chiplet (16 cores) was used. On Cirrus the experiment was repeated on both CPU and GPU nodes. On the CPU nodes, the compiler of choice was GNU 8.2.0 with the OpenMP backend enabled, and a single CPU (18 cores) was used, again to avoid NUMA effects. On the GPU nodes the device code was compiled using NVCC 11.6 and the host code with GNU 8.2.0, with CUDA backend enabled and the experiment was run on a single GPU.

Figure 4 confirms our hypothesis that the DIA format would perform best both on CPU and GPU, giving improvements of $3.5\times$ on ARCHER2 and up to $4.5\times$ on a Cirrus GPU node compared to the equivalent CSR implementations using the same backend. Since the matrix is highly regular with non-zeros around the diagonals, the DIA format exploits this by offering direct and contiguous access to both the values of the matrix A and the vector x, eliminating the need for indirect accesses. However for smaller problem sizes (up to $32\times32\times32$ - see Figures 4a and 4b), CSR is either performing better or on par with DIA as the system matrix is small enough such that the costs from indirection are less than the ones from the zero-padding introduced by DIA. In addition, even though COO underperforms compared to CSR, it is worth pointing out that for a problem size of $256\times256\times256$ on ARCHER2, COO now beats CSR performance. This observation further motivates the use of dynamic matrices and shows the importance of having

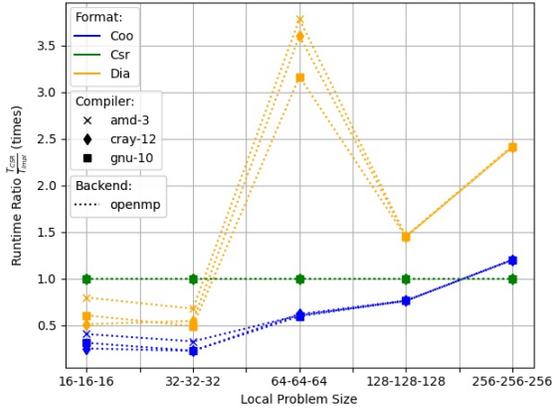the ability to adapt the data-structure at runtime.

E. Multi-node Performance

Enabling MPI results in an unstructured local system matrix with a sparsity pattern similar to Figure 2. As a result, DIA introduces excessive zero padding and will cause out-of-memory errors for large problems. Instead of conceptually dividing the local system matrix into local and remote parts, we can actually split it into two parts and convert each part to *DynamicMatrix*. By exploiting the runtime switching functionality supported by *Morpheus* we can change the active state of each matrix to potentially different formats based on the sparsity pattern, and also extend this principle across processes.

In this experiment we investigate the multi-node scaling performance of the different versions of *Morpheus*-enabled HPCG. Scaling performance is evaluated both in terms of strong scaling (increasing number of processing units whilst keeping the same global problem) and weak scaling (maintaining same problem size per node whilst increasing number of processing units). The reference version is the original MPI-enabled HPCG and the *Morpheus*-enabled versions are: 1) *Morpheus*, where the storage format of the local matrix varies for all processes in the same way and the remote matrix is set to CSR 2) *Ghost*, where the remote storage format varies for all processes and local matrix is set to CSR and 3) *Multi-Format*, where the formats of both local and remote matrices vary and can be different across processes. To obtain the runtime ratio we divide the SpMV runtime of the reference MPI-enabled HPCG by the SpMV runtime of each version. A ratio above 1 indicates a speed-up over the performance achieved with the original HPCG implementation.
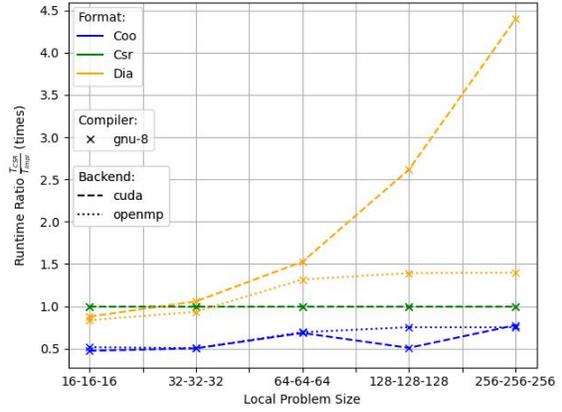
The experiments on ARCHER2 for the Serial and OpenMP backends are each configured with 128 processes and 8 processes each with 16 threads respectively. The global problem size for strong scaling is $512 \times 512 \times 256$ and the same problem size is used per node for the weak scaling. For Cirrus we use the GPU nodes to benchmark the CUDA backend, which is configured as one GPU per process. For comparison, we evaluate the MPI-only reference HPCG on the CPU nodes. The global problem size for strong scaling is set to $384 \times 256 \times 128$, and we duplicate this per node during the weak scaling experiment. Note that the CPU cores on Cirrus are underpopulated (32 processes out of the available 36) to simplify the computation process of problem size as we vary resources.

In order to select the best format for each process and matrix in the *Multi-Format* version, a naive auto-tuner is used. For each combination of formats, profiling runs are performed and the per-process runtime of the SpMV kernel is recorded. The auto-tuner then selects the best performing format combination for each process.

Figure 5a shows that on ARCHER2, *Morpheus*-enabled HPCG runs up to $2.5\times$ faster when the active state of the local matrix is set to DIA compared to the MPI-enabled reference HPCG. This remains true until the problem size per
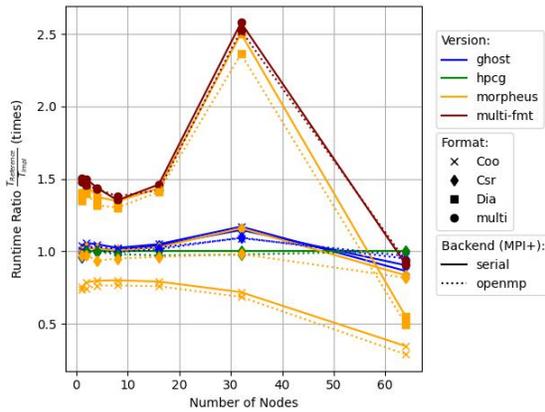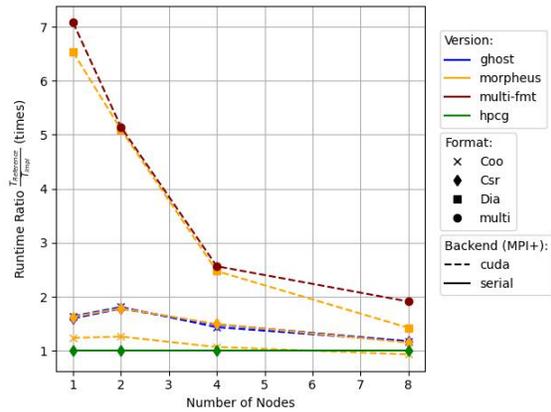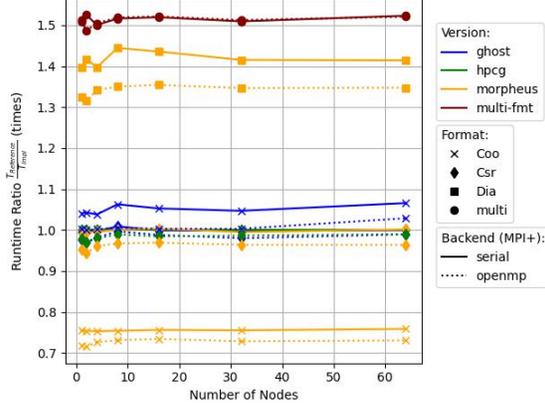
(a) ARCHER2

(b) Cirrus

Fig. 4: Single node performance of the *Morpheus*-HPCG. For each backend and compiler, the performance is measured as the SpMV runtime ratio of the *DynamicMatrix* with active state set to CSR w.r.t the *DynamicMatrix* with an active state to each of the supported storage formats (COO, CSR, DIA) over a set of problem sizes on ARCHER2 and Cirrus. A ratio above 1 indicates a speedup over the performance achieved when using CSR.
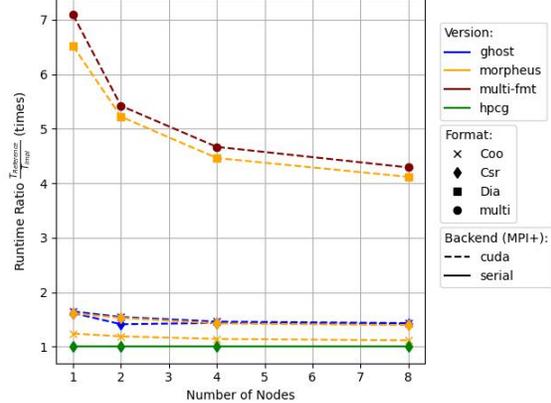

(a) Strong Scaling (ARCHER2)

(b) Strong Scaling (Cirrus)

(c) Weak Scaling (ARCHER2)

(d) Weak Scaling (Cirrus)

Fig. 5: Multi-node strong and weak scaling performance of *Morpheus*-HPCG. The SpMV runtime ratio for each number of nodes(/GPUs) is obtained by evaluating the MPI-only reference HPCG implementation w.r.t to the versions provided by *Morpheus*-enabled HPCG. Available versions are: 1) Original HPCG 2) Morpheus (Local part changes format, remote part in CSR) 3) Ghost (Local part in CSR, remote part changes) 4) Multi-format (Local and remote part change per process to the optimum format).

node becomes very small; this happens at 64 nodes when the system matrix becomes almost dense, and now the CSR format outperforms DIA by $2\times$. In addition, changing the active state of the remote matrix slightly improves runtime with up to $1.2\times$ when it is set to COO. However, by adopting the multi-format approach, as the amount of resources varies and different formats perform best, we are now able to maintain the best possible performance since we are no longer restricted to using a single storage format. Similarly, in Figure 5b the GPU runtime using DIA compared to the MPI-enabled reference HPCG is $6.5\times$ faster on a single GPU and drops to $1.3\times$ on 8 GPUs. This can be attributed to the fact that the problem size per GPU significantly reduces and consequently occupancy on each GPU reduces too. Furthermore, as the number of GPUs increases, the time spent in the `ExchangeHalo` routine increases and approaches the times spent in the SpMV kernel. The multi-format version achieves similar performance as the *Morpheus* version and outperforms it by 25% on 8 GPUs as it sets the active type of the local matrix to DIA and the ghost matrix on half of the processes to COO and the other half to CSR. Note that because now the local and remote matrices on each process are separate, it is also possible to overlap communication and computation in order to reduce communication time and improve scaling, however this is out of the scope of this work.

The weak scaling experiments in Figures 5c and 5d show that for the current problem size per node the best performance is obtained when the active state of the local matrix is set to DIA. Changing the active state of the remote matrix from CSR to COO will marginally improve runtime. For ARCHER2, it is worth pointing out that the optimal setup (which improves runtime by $1.5\times$) is obtained when the active states of local and remote matrices on each process are set to DIA and COO respectively. On Cirrus, weak scaling demonstrates similar trends as strong scaling however because the load per GPU here is larger and the communication times are the same as before, the runtime performance improvement saturates at $4\times$. Note that for both architectures and available backends, in the case where the problem size per node is very small, the optimal format ends up being CSR for both the local and remote matrices.

VI. Related Work

Research efforts into optimizing sparse computations focus on many directions with the most prominent being the creation of novel storage formats, and performance tuning of operations through parameter optimization or automatic format selection.

Storage Formats Many of the new formats that are proposed are derivations or extensions of existing formats aiming to improve performance by addressing known weaknesses. In addition, hybrid approaches that combine multiple formats into a single data structure have been proposed, with the aim to better exploit certain sparsity patterns. Each hybrid format uses mechanisms to determine in which underlying format each portion of the matrix is stored. Examples include Hybrid (HYB) [15] format, that combines ELLPACK (ELL) [2] and

COO, Hybrid DIA/CSR (HDC) [16] and Hybrid ELL/CSR (HEC) [17]. The cocktail [8] format takes the idea a step further and partitions the matrix into many sub-matrices by collecting matrix features and enforcing partition policies with each partition assigned one of the nine available formats.

Auto-tuners Auto-tuners such as OSKI [18], FFTW [19] and ATLAS [20] are few of the successful approaches in domain-specific performance-critical libraries. In the field of sparse linear algebra, the focus of auto-tuners is mainly on tuning performance through selecting the storage format, the kernel implementation or parameters. For example, SMAT [7], SMATER [21], clSpMV [8] and Zhao et al. [3] enable runtime selection of the best format and SpMV kernel implementation across architectures for a given matrix. In a similar manner, Xie et al. [22] implement an auto-tunner that automatically determines the best format and Sparse Matrix-Matrix Multiplication (SpGEMM) algorithm for arbitrary sparse matrices.

Libraries PETSc [23], Eigen [24], AMGX [25] and CUSP [26] are only few of the examples of software libraries developed for solving linear systems. Although these solutions constitute the current state-of-the-art, they are either specific to a single target hardware or support only a single format internally and it is often a significant undertaking to provide full support for a new format. Given the heterogeneous nature of modern High Performance Computing (HPC) systems, providing support for multiple target architectures and provisions for supporting future architectures is a requirement.

VII. Conclusions and Further Work

As there is no single sparse matrix storage format that will perform optimally across different operations, sparsity patterns and target architectures, the ability to dynamically change the underlying data-structure to better match the computational pattern and hardware characteristics enables a range of optimization opportunities. Extending this concept into a distributed environment allows us to create a dynamic distributed matrix capable of changing its storage format on each process and achieve the best possible performance irrespective of the number or type of resources that are available.

Providing end-users with a mechanism to optimize their applications without any code changes (beyond the initial porting effort) and allowing them to target new architectures effectively increases the life-time of their software as support for more formats and target architectures is added.

As a next step, we will add further storage formats to *Morpheus* and study the performance characteristics over a wider set of matrices with different sparsity patterns. In addition, updating the auto-tuner to automatically select the optimum format online and not from profiling runs remains an avenue for further research.

References

[1] P. Colella, "Defining Software Requirements for Scientific Computing," http://view.eecs.berkeley.edu/w/images/temp/6/6e/20061003235551!DARPAHPCS.ppt, 2006.

[2] S. Filippone, V. Cardellini, D. Barbieri, and A. Fanfarillo, "Sparse Matrix-Vector Multiplication on GPGPUs," *ACM Trans. Math. Softw.*, vol. 43, no. 4, Jan. 2017. [Online]. Available: https://doi.org/10.1145/3017994

[3] Y. Zhao, J. Li, C. Liao, and X. Shen, "Bridging the gap between deep learning and sparse matrix format selection," in *Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, ser. PPoPP '18. New York, NY, USA: Association for Computing Machinery, 2018, p. 94–108. [Online]. Available: https://doi.org/10.1145/3178487.3178495

[4] W. Liu and B. Vinter, "Csr5: An efficient storage format for cross-platform sparse matrix-vector multiplication," in *Proceedings of the 29th ACM on International Conference on Supercomputing*, ser. ICS '15. New York, NY, USA: Association for Computing Machinery, 2015, p. 339–350. [Online]. Available: https://doi.org/10.1145/2751205.2751209

[5] E. Coronado-Barrientos, M. Antonioletti, and A. Garcia-Loureiro, "A new axt format for an efficient spmv product using avx-512 instructions and cuda," *Advances in Engineering Software*, vol. 156, p. 102997, 2021. [Online]. Available: https://www.sciencedirect.com/science/article/pii/S0965997821000260

[6] M. Kreutzer, G. Hager, G. Wellein, H. Fehske, and A. Bishop, "A unified sparse matrix data format for efficient general sparse matrix-vector multiplication on modern processors with wide simd units," *SIAM Journal on Scientific Computing*, vol. 36, 07 2013.

[7] J. Li, G. Tan, M. Chen, and N. Sun, "Smat: An input adaptive auto-tuner for sparse matrix-vector multiplication," in *Proceedings of the 34th ACM SIGPLAN Conference on Programming Language Design and Implementation*, ser. PLDI '13. New York, NY, USA: Association for Computing Machinery, 2013, p. 117–126. [Online]. Available: https://doi.org/10.1145/2491956.2462181

[8] B.-Y. Su and K. Keutzer, "Clspmv: A cross-platform opencl spmv framework on gpus," in *Proceedings of the 26th ACM International Conference on Supercomputing*, ser. ICS '12. New York, NY, USA: Association for Computing Machinery, 2012, p. 353–364. [Online]. Available: https://doi.org/10.1145/2304576.2304624

[9] C. Stylianou, "Morpheus: a library for efficient runtime switching of sparse matrix storage formats," 2022. [Online]. Available: https://github.com/morpheus-org/morpheus

[10] J. Dongarra, M. A. Heroux, and P. Luszczek, "High-performance conjugate-gradient benchmark: A new metric for ranking high-performance computing systems," *The International Journal of High Performance Computing Applications*, vol. 30, no. 1, pp. 3–10, 2016. [Online]. Available: https://doi.org/10.1177/1094342015593158

[11] N. Bell and M. Garl, "Efficient sparse matrix-vector multiplication on cuda," NVIDIA, Tech. Rep., 2008.

[12] Y. Saad, *Iterative Methods for Sparse Linear Systems*, 2nd ed., ser. Other Titles in Applied Mathematics. Society for Industrial and Applied Mathematics, 2003. [Online]. Available: https://www-users.cs.umn.edu/~saad/IterMethBook_2ndEd.pdf

[13] C. R. Trott, D. Lebrun-Grandié, D. Arndt, J. Ciesko, V. Dang, N. Ellingwood, R. Gayatri, E. Harvey, D. S. Hollman, D. Ibanez, N. Liber, J. Madsen, J. Miles, D. Poliakoff, A. Powell, S. Rajamanickam, M. Simberg, D. Sunderland, B. Turcksin, and J. Wilke, "Kokkos 3: Programming model extensions for the exascale era," *IEEE Transactions on Parallel and Distributed Systems*, vol. 33, no. 4, pp. 805–817, 2022.

[14] E. Gamma, R. Helm, R. Johnson, and J. Vlissides, *Design Patterns: Elements of Reusable Object-Oriented Software*. USA: Addison-Wesley Longman Publishing Co., Inc., 1995.

[15] A. Monakov, A. Lokhmotov, and A. Avetisyan, "Automatically tuning sparse matrix-vector multiplication for gpu architectures," in *Proceedings of the 5th International Conference on High Performance Embedded Architectures and Compilers*, ser. HiPEAC'10. Berlin, Heidelberg: Springer-Verlag, 2010, p. 111–125. [Online]. Available: https://doi.org/10.1007/978-3-642-11515-8_10

[16] W. Yang, K. Li, Y. Liu, L. Shi, and L. Wan, "Optimization of quasi-diagonal matrix–vector multiplication on gpu," *The International Journal of High Performance Computing Applications*, vol. 28, no. 2, pp. 183–195, 2014. [Online]. Available: https://doi.org/10.1177/1094342013501126

[17] H. Liu, S. Yu, Z. Chen, B. Hsieh, and L. Shao, "Sparse matrix-vector multiplication on nvidia gpu," 2012.

[18] R. Vuduc, J. W. Demmel, and K. A. Yelick, "OSKI: A library of automatically tuned sparse matrix kernels," *Journal of Physics: Conference Series*, vol. 16, pp. 521–530, jan 2005. [Online]. Available: https://doi.org/10.1088/1742-6596/16/1/071

[19] M. Frigo and S. Johnson, "The design and implementation of fftw3," *Proceedings of the IEEE*, vol. 93, no. 2, pp. 216–231, 2005.

[20] R. C. Whaley and J. J. Dongarra, "Automatically tuned linear algebra software," in *Proceedings of the 1998 ACM/IEEE Conference on Supercomputing*, ser. SC '98. USA: IEEE Computer Society, 1998, p. 1–27.

[21] G. Tan, J. Liu, and J. Li, "Design and implementation of adaptive spmv library for multicore and many-core architecture," *ACM Trans. Math. Softw.*, vol. 44, no. 4, Aug. 2018. [Online]. Available: https://doi.org/10.1145/3218823

[22] Z. Xie, G. Tan, W. Liu, and N. Sun, "Ia-spgemm: An input-aware auto-tuning framework for parallel sparse matrix-matrix multiplication," in *Proceedings of the ACM International Conference on Supercomputing*, ser. ICS '19. New York, NY, USA: Association for Computing Machinery, 2019, p. 94–105. [Online]. Available: https://doi.org/10.1145/3330345.3330354

[23] S. Balay, S. Abhyankar, M. F. Adams, S. Benson, J. Brown, P. Brune, K. Buschelman, E. M. Constantinescu, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, V. Hapla, T. Isaac, P. Jolivet, D. Karpeev, D. Kaushik, M. G. Knepley, F. Kong, S. Kruger, D. A. May, L. C. McInnes, R. T. Mills, L. Mitchell, T. Munson, J. E. Roman, K. Rupp, P. Sanan, J. Sarich, B. F. Smith, S. Zampini, H. Zhang, H. Zhang, and J. Zhang, "PETSc Web page," https://petsc.org/, 2022. [Online]. Available: https://petsc.org/

[24] G. Guennebaud, B. Jacob *et al.*, "Eigen v3," http://eigen.tuxfamily.org, 2010.

[25] M. Naumov, M. Arsaev, P. Castonguay, J. Cohen, J. Demouth, J. Eaton, S. Layton, N. Markovskiy, I. Reguly, N. Sakharnykh, V. Sellappan, and R. Strzodka, "Amgx: A library for gpu accelerated algebraic multigrid and preconditioned iterative methods," *SIAM Journal on Scientific Computing*, vol. 37, no. 5, pp. S602–S626, 2015. [Online]. Available: https://doi.org/10.1137/140980260

[26] S. Dalton, N. Bell, L. Olson, and M. Garland, "Cusp: Generic parallel algorithms for sparse matrix and graph computations," 2014, version 0.5.0. [Online]. Available: http://cusplibrary.github.io/