

The AllScale API

1st Philipp Gschwandtner
Department of Computer Science
University of Innsbruck
Innsbruck, Austria
philipp.gschwandtner@uibk.ac.at

2nd Herbert Jordan
Department of Computer Science
University of Innsbruck
Innsbruck, Austria
herbert.jordan@uibk.ac.at

3rd Peter Thoman
Department of Computer Science
University of Innsbruck
Innsbruck, Austria
peter.thoman@uibk.ac.at

4th Thomas Fahringer
Department of Computer Science
University of Innsbruck
Innsbruck, Austria
thomas.fahringer@uibk.ac.at

Abstract—Effectively implementing scientific algorithms in distributed memory parallel applications is a difficult task for domain scientists, as evident by the large number of domain-specific languages and libraries available today attempting to facilitate the process. However, they usually provide a closed set of parallel patterns and are not open for extension without vast modifications to the underlying system.

In this work, we present the AllScale API, a programming interface for developing distributed memory parallel applications with the ease of shared memory programming models. The AllScale API is closed for modification but open for extension, allowing new, user-defined parallel patterns and data structures to be implemented based on existing core primitives and therefore fully supported in the AllScale framework. Focusing on high-level functionality directly offered to application developers, we present the design advantages of such an API design, detail some of its specifications and evaluate it using three real-world use cases. Our results show that AllScale decreases the complexity of implementing scientific applications for distributed memory while attaining comparable or higher performance compared to MPI reference implementations.

Index Terms—API, programming interface, parallel programming, shared memory, distributed memory, parallel operator, data structure

I. INTRODUCTION

Even with the recent trend of many-core processors, providing users with dozens of cores per chip in a single memory address space, distributed memory systems pose an essential aspect of HPC in order to achieve large-scale performance for scientific applications. Although there are certain system architectures that overcome the issue of distinct memory address spaces by hardware means (e.g. SGI’s UV [1] series using the NumaLink protocol), the conventional approach is still to handle distinct memory address spaces in the software stack by providing a global address space in software or by explicit message exchange.

This project has received funding from the European Unions Horizon 2020 research and innovation programme as part of the FETHPC AllScale project under grant agreement No 671603. The computational results presented have been achieved in part using the Vienna Scientific Cluster and the Regionales RechenZentrum Erlangen.

However, most of these ubiquitous software solutions entail several disadvantages that make them hard to use for domain scientists. Programming interfaces such as MPI are often too low-level for non-computer science experts and clutter up the application with non-domain-relevant source code. On the other hand, there are high-level domain-specific languages or libraries that lack extensibility in order to support new scientific problems [2]. In addition, many of these solutions often lack the composability required for building libraries and integrating them seamlessly into larger applications, they deny an incremental approach that allows parallelizing an application step by step, or are limited to shared memory only. Therefore, users often resort to combining several of these solutions (e.g. MPI + OpenMP), which presupposes knowledge in at least two different programming models and entails a lack of resource management coordination that is left to the user.

In contrast, the AllScale API aims at providing the application developer with a single, extensible programming interface to express their parallel algorithms on a high level of abstraction, with automatic support for distributed memory.

The specific contributions of this work are:

- a shared-memory-style API for high-level specifications of algorithms and data structures with implicit distributed memory support,
- the capability of expressing new algorithms by extending the API with full compatibility to the rest of the software stack, and
- an evaluation of its programmability and performance using three real-life use cases.

While documentation and tutorials introducing the novice to the AllScale API are available online¹, the remainder of this work focuses on the API specification and important properties.

The rest of the paper is structured as follows. Section II discusses API design motivation while Section III and Section IV detail API components. Implementation information

¹https://github.com/allscale/allscale_api/wiki

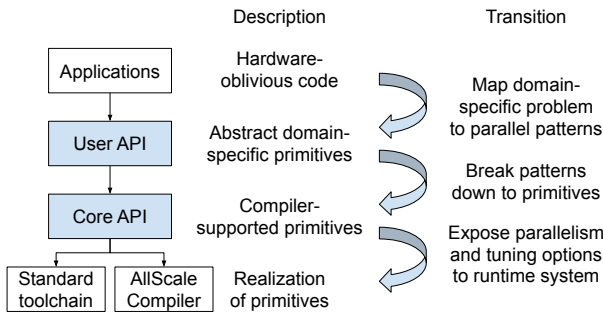


Fig. 1. AllScale API design and usage overview.

is given in Section V. Three real-world pilot applications and their respective API use are presented in Section VI followed by an evaluation in Section VII. Related work is discussed in Section VIII and Section IX provides the conclusion and future work.

II. API DESIGN

AllScale aims at providing domain scientists with the ability to write parallel applications for distributed memory using an API that is as easy to use as shared memory programming models such as OpenMP. While AllScale consists of many components including the API, a compiler, a distributed memory runtime system, and additional components for monitoring, resilience, etc., this work will present the API in detail.

The AllScale API is the façade of the AllScale Environment towards end-user applications. It provides the necessary primitives to express parallelism, data dependencies, and needed synchronization steps within application code. The API is subdivided into two layers: the *AllScale Core API* and the *AllScale User API*. Their relationship is illustrated in Fig. 1 and further discussed in the remainder this section.

The Core API provides a concise set of basic generic primitives, comprising parallel control flow, synchronization, and communication constructs. It furthermore offers a generic data item interface that enables automatic data management of user-defined data structures. The User API is harnessing the expressive power of the Core API to provide specialized primitives for particular use cases, including basic constructs like parallel loops or adaptive grids.

The purpose of the subdivision into a Core and User API is to enable the implementation of a variety of parallel primitives on top of a small, concise set of central constructs which can be utilized to provide portability among different implementations of the AllScale Core API. Currently there are two implementations available within AllScale:

- a shared memory, pure C++ implementation, also referred to as *standard toolchain*, which can be compiled by any C++14-compliant compiler with no further third-party library dependencies — this implementation serves as a development platform for AllScale applications and also represents a reference implementation; and
- the implementation utilizing the AllScale Compiler and Runtime System, also called the *AllScale toolchain*, which comprises a combination of static program analysis

(crucial for automatically deriving data dependencies required for distributed memory execution), code generation, scheduling, and resilience techniques to provide a highly scalable and portable implementation of the Core API on distributed memory systems.

Hence, applications developed within AllScale can be ported from shared to distributed memory simply by switching the toolchain, without any modifications required in the application. Additional parallel constructs may be introduced in the User API without the necessity of altering the underlying Core API implementation. Thus, the User API layer provides an effective way of extending the range of supported parallel patterns.

Furthermore, the User API shields application developers from the complexity of the Core API constructs. Due to the introduction of the User API efficient implementations of primitives native to the domain of the applications can be provided by parallelization experts. Therefore, AllScale provides a separation of concerns — with the overall task of providing efficient parallel codes — distributed among three contributors:

- the *domain expert*, aiming at obtaining the most effective algorithmic solution for the problem of interest;
- the *HPC expert*, able to develop efficient domain specific primitives to be used by the domain expert, focusing on e.g. communication and synchronization overheads or cache efficiency; and
- the *system-level expert* focusing on providing the most flexible and portable implementation of the Core API, hence handling load management, scheduling, resilience, and hardware management obligations.

The separation of responsibilities also effects the code base. By shielding the domain expert from all the underlying details (e.g. synchronization, communication, cache efficiency, scheduling, utilization of low-level parallel APIs), the resulting application code remains free of the otherwise necessary management code. This improves the maintainability of the resulting applications and thus the productivity of the domain expert.

III. CORE API

This section will detail the Core part of the AllScale API, specifically the primitives for parallel control flow and the concept of data items and their requirements. The User API, discussed in the section thereafter, builds on-top of these basic constructs to provide more high-level operations to domain experts. Note that while the Core API also offers additional tools such as a small performance profiling tool or means for managing distributed IO, discussing those exceeds the scope of this paper.

A. Parallel Control Flow

The AllScale Core API provides a single primitive for running concurrent tasks, resulting in feasible yet profound compiler and runtime system support for automatic distributed memory management of parallel applications. This single

parallelism primitive forms the basis for all higher-level operators of the User API such as parallel loops, allowing the User API to be open for extension with new higher-level operators without any modifications required in the Core API or underlying compiler and runtime system [3].

This primitive, the *prec* [4] operator, is a higher order function combining three given functions into a new, recursive function. The three combined input functions are:

- a function testing for the base case of a recursion,
- a function processing the base case of a recursion, and
- a function processing the recursive step case.

The result is a new recursive function which, for a given input parameter, conducts the specified computation accordingly. To support an arbitrary input type, the *prec* operator has the type

$$\left(\begin{array}{l} \alpha \rightarrow \text{bool}, \\ \alpha \rightarrow \beta, \\ (\alpha, \alpha \rightarrow \text{treature}\langle\beta\rangle) \rightarrow \text{treature}\langle\beta\rangle \end{array} \right) \rightarrow (\alpha \rightarrow \text{treature}\langle\beta\rangle)$$

where α is the parameter type of the resulting recursive function and $\text{treature}\langle\beta\rangle$ is a parameterized abstract data type (ADT) modeling a handle on parallel tasks. The three parameters of the *prec* operator are the input functions discussed above. The resulting value of type $\alpha \rightarrow \text{treature}\langle\beta\rangle$ is a function which, upon invocation, spawns a new task conducting the specified recursive operation in parallel. The resulting task handle can be utilized to orchestrate the parallel execution of additional tasks. A more in-depth discussion of ADTs can be found online [5].

B. Data Structure Primitives

While the parallel control flow primitive has been covered so far, it is not sufficient to compose parallel applications for distributed memory. In order to properly manage data dependencies for parallel tasks executed in distinct memory address spaces, a specification for user-defined data structures needs to be defined as well. The purpose of this specification is to provide a single generic interface for HPC experts to implement new user-defined data structures while offering management access to the underlying runtime system for data distribution.

To this end, the data structure primitives offered by the Core API are a mere specification of any potential data type's interfaces and behaviors. Any data type T to be managed by an AllScale API implementation must provide a fragment type F for managing data storage and a range type R for addressing and managing sub-ranges of the data structure. Table I lists the operators required to be defined by F and R . Proper implementation of these operators for any arbitrary data structure ensures its suitability for automatic distributed data management by the AllScale Compiler and Runtime System. Several examples that implement widely-used data structures such as grids are discussed in Section IV-B while their implementation, among others, can be found online².

²<https://git.io/fj4Xj>

```

1 #include <array>
2 #include <allscale/api/user/algorithm/pfor.h>
3 namespace alg = allscale::api::user::algorithm;
4 using ArrayType = std::array<int,N>;
5 const int N = 200;
6 void initAndIncrement(const ArrayType& data, ArrayType& output) {
7     auto ref = alg::pfor(0,N,&{(int i) {
8         output[i] = ...; // initialization
9     }});
10    alg::pfor(1,N-1,&{(int i) {
11        output[i] += data[i+1] + data[i] + data[i-1];
12    }}, alg::neighborhood_sync(ref));
13 }

```

Fig. 2. Two *pfor* operators initializing and incrementing data in a `std::array` with fine-grained synchronization. The second *pfor* will execute iteration i after the first has finished its iterations $i - 1$, i , and $i + 1$. Constructs specific to the AllScale API are shown in blue and underlined.

IV. USER API

The generic nature of the Core API exceeds the complexity which could be effectively handled by domain experts for implementing parallel algorithms. For this reason, the AllScale User API aims at providing a set of more user-friendly, higher-level constructs for the composition of parallel applications by domain experts. The implementation of these constructs is carried out by HPC experts utilizing the primitives offered by the Core API.

A. Parallel Control Flow Constructs

While the User API is open for extension with new parallel patterns, several frequently-encountered patterns such as parallel loops are already provided and discussed below.

1) *Parallel Loops*: A vast majority of algorithms expressing data parallelism rely on parallel loops. They provide the means to perform computational work in an iteration space in parallel at the cost of executing the individual iterations concurrently and in an arbitrary order. To that end, the User API offers a parallel loop construct for realizing data-parallel programming within the AllScale environment.

Let *iterator* be a random access iterator. Then the *pfor* operator provides a parallel loop execution with the parameters defined in Table II. Fig. 2 shows a sample usage of the *pfor* operator with fine-grained synchronization. Several of these synchronization patterns are available, such as *neighborhood_sync* or *one_on_one*. HPC experts are free to extend these by new patterns not yet covered.

2) *Recursive space/time decomposition*: A frequently utilized template for large-scale high-performance applications are stencils. In a stencil-based application, an update operation is iteratively applied to the elements of an n -dimensional array of cells. Thereby, for each update, the update operation is combining the previous values of cells within a locally confined area surrounding the targeted cell location to obtain the updated value for the targeted cell. Since these update operations within a single update step (also known as timestep) are independent, this application pattern provides a valuable source for parallelism within a correspondingly shaped application. The User API offers the *stencil* operator, the parameters of which are defined in Table III.

TABLE I
OPERATORS TO BE DEFINED BY FRAGMENT F AND RANGE R TYPES OF AN ALLSCALE DATA STRUCTURE.

Name	Type	Description
Fragment		
create	$R \rightarrow F$	creates a fragment covering (at least) the specified range
delete	$F \rightarrow \text{unit}$	deletes the given fragment
resize	$(F, R) \rightarrow \text{unit}$	alters the capacity of given fragment F to cover at least the range R
mask	$F \rightarrow T$	provides access to the data stored in fragment F via the interface defined by type T
extract	$(F, R) \rightarrow \text{Archive}$	extracts the data addressed by R from fragment F and packs it into an archive; <i>Archive</i> is a generic type of a utility provided by the API implementations to serialize data to be transferred between address spaces
insert	$(F, R, \text{Archive}) \rightarrow \text{unit}$	imports the data stored in the given archive into fragment F at the specified range R
Range		
union	$(R, R) \rightarrow R$	computes the union of two ranges
intersect	$(R, R) \rightarrow R$	computes the intersection of two ranges
difference	$(R, R) \rightarrow R$	computes the set difference of two ranges
empty	$(R) \rightarrow \text{bool}$	determines whether the given range is empty, thus addressing no elements
pack	$(R) \rightarrow \text{Archive}$	serializes instances
unpack	$(\text{Archive}) \rightarrow R$	deserializes instances

TABLE II
PARAMETERS OF THE *pfor* OPERATOR.

Name	Type	Description
begin	<i>iterator</i>	inclusive beginning of the iterator range
end	<i>iterator</i>	exclusive end of the iterator range
body	$(\text{iterator}) \rightarrow \beta$	the function applied to each element
dependency	<i>dep(iterator)</i>	optional dependency for fine-grained synchronization

```

1 #include <allscale/api/user/data/grid.h>
2 // create a two-dimensional grid of integers of size 10x20
3 allscale::api::user::data::Grid<int, 2> grid({10,20});
4 // initialize all elements with 1.0
5 grid.pforEach([](int& element) { element = 1.0; });
6 // set element at position [7,9] to 5.0
7 grid[{7,9}] = 5.0;

```

Fig. 3. Example usage of the *Grid* data structure.

3) *Additional operations*: Beyond the *pfor* and *stencil* operators presented thus far, the User API offers additional parallel operations that are frequently encountered in parallel applications. These include e.g. the *map-reduce* operator for data aggregation, the *async* operator for single tasks, or the *vcycle* operator for multi-grid methods. However, a more detailed presentation is omitted for brevity.

B. Data Structures

1) *Grid*: A frequently-encountered data structure in high-performance codes is formed by n -dimensional arrays of values. While many programming languages support such structures for arbitrary dimensions, C/C++ only supports one-dimensional, dynamically sized arrays natively. However, this leaves creation and management of these structures to the user, forming a major obstacle for the usability of C++ on distributed memory systems.

To ease the use of C++ for use cases depending on such structures, the AllScale User API provides a uniform *Grid* data structure providing the following features:

- regular n -dimensional array of runtime-defined size

- efficient read/write random access operators
- efficient scan operation (processing all elements)
- type-parameterized in element type and no. of dimensions
- enforces the serializability of its element types
- implements data item concept for automated distribution

Let $\text{Grid}\langle\alpha, n\rangle$ be the abstract data type family implemented by the AllScale User API to represent n -dimensional grids, where α is a type variable specifying the element type. Furthermore, let $\text{type}\langle\alpha\rangle$ be the meta type of type α . Then Table IV lists the operators defined on *Grid* data structures.

2) *Adaptive Grid*: The *Adaptive Grid* is an advanced variant of the *Grid* structure also frequently encountered within simulation code. In addition to the properties of *Grid*, the *Adaptive Grid* provides means to nest grids within grid cells. For a given instance, each top-level grid cell contains an identically structured fixed-length sequence of grids. The first of those contains a single cell. Every consecutive grid contains a multiple number of cells per dimension of its predecessor. Each top-level grid cell comprising the sequence of its nested grids is referred to as an *Adaptive Grid Cell*.

Let $\text{AGrid}\langle\alpha, n, [r_1, \dots, r_l]\rangle$ be the ADT family implemented by the AllScale User API to represent n -dimensional *Adaptive Grids*, where α is a type variable specifying the element type, r_1, \dots, r_l the refinement factors, and l the number of refinement levels. Thus, the size of the grid at level i is defined by

$$s(i) = \begin{cases} [1, \dots, 1] \in \text{int}^n & \text{if } i = 0, \\ s(i-1) * r_i & \text{otherwise.} \end{cases}$$

To address elements within an *Adaptive Grid* an extension of *Grid* coordinates is required. While elements within a *Grid* can be addressed using a single coordinate of type int^n , the *Adaptive Grid* requires information regarding the location of the addressed element in the nested grid structure. Thus, additional coordinates to navigate through these refinement layers are required. Hence, to address an element within an *Adaptive Grid*, a hierarchical coordinate of type $(\text{int}^n)^+$ is

TABLE III
PARAMETERS OF THE *stencil* OPERATOR.

Name	Type	Description
timesteps	int	the number of time steps to be computed
size	int^n	the spatial size of the n -dimensional data to be processed
kernel function	$(int, int^n, \alpha^{s_1 \times \dots \times s_n}) \rightarrow \alpha$	the update function accepting the current time, location, and grid, computing the resulting value
kernel shape	$(int^n)^*$	a compile-time-constant list of offsets to cells accessed by the kernel, determining its shape
boundary function	$(int, int^n, \alpha^{s_1 \times \dots \times s_n}) \rightarrow \alpha$	the update function for boundary cases, where some elements are outside the grid
initialization function	$(int^n) \rightarrow \alpha$	computes the initial value for a cell at a given coordinate
finalization function	$(int^n, \alpha) \rightarrow unit$	a function consuming the value of a cell at the end of a computation
observers	$\left(\begin{array}{l} (int, int^n) \rightarrow bool \\ (int, int^n, \alpha) \rightarrow unit \end{array} \right)^*$	a list of pairs, each describing an observer with a time/location filtering function and the actual trigger function to be applied

TABLE IV
OPERATORS DEFINED ON *Grid* DATA STRUCTURES.

Name	Type	Description
create	$\left(\begin{array}{l} type(\alpha), \\ int^n \end{array} \right) \rightarrow Grid(\alpha, n)$	creates a new n -dimensional grid with element type α of the given size
destroy	$Grid(\alpha, n) \rightarrow unit$	deletes the given grid
read	$\left(\begin{array}{l} Grid(\alpha, n), \\ int^n \end{array} \right) \rightarrow \alpha$	reads the element from the given grid at the specified coordinates
write	$\left(\begin{array}{l} Grid(\alpha, n), \\ int^n, \\ \alpha \end{array} \right) \rightarrow unit$	updates the element within the given grid at the specified coordinates
scan	$\left(\begin{array}{l} int^n, \\ int^n, \\ (int^n) \rightarrow \beta \end{array} \right) \rightarrow treature(unit)$	applies the given function (in parallel) to all elements of the given interval in an arbitrary order

required. For instance, the coordinate $[[7, 3], [2, 4], [8, 2]]$ addresses the element located within the cell that can be reached by navigating first to the top-level cell $[7, 3]$, continuing to cell $[2, 4]$ of its first refinement layer, and ending up within cell $[8, 2]$ of the second refinement layer. Let $seq\langle r_1, \dots, r_l \rangle$ be the static meta-type of a sequence of integers r_1, \dots, r_l , then Table V lists the operators defined on *Adaptive Grid* data structures.

3) *Unstructured Mesh*: The *Mesh* data structure is designed to represent a graph structure of multiple node types that are connected through various types of edges. Furthermore, a *Mesh* may consist of several layers, which describe the same graph in different levels of detail. Hierarchical edges may connect the same nodes of different layers.

Besides the topological information maintained by *Mesh* instances, means to maintain attributes associated to nodes, edges, and hierarchical edges within a *Mesh* need to be included. For instance, node IDs, coordinates, volumes, temperatures, and other domain space specific properties may be incorporated through this facility.

Let n_1, \dots, n_m be a list of node types, $e_1, \dots, e_k \in \{n_1, \dots, n_m\}^2$ a list of edge types, and $h_1, \dots, h_o \in \{n_1, \dots, n_m\}^2$ a list of hierarchical edge types. Then the type $Mesh\langle [n_1, \dots, n_m], [e_1, \dots, e_k], [h_1, \dots, h_o], l \rangle$ represents the type of a *Mesh* structure including the given node, edge, and hierarchical edge types on l layers. Furthermore, let $id\langle \alpha, l \rangle$ be an identifier for an element of type α on layer l within a *Mesh* — thus the type of ID used for addressing nodes, edges, or hierarchical edges within meshes. Also, let $MData\langle n, l, \alpha \rangle$ be the type of an attribute collection associating values of type α to

nodes of type n located on layer l of some *Mesh* instance. Finally, let $MBuilder\langle [n_1, \dots, n_m], [e_1, \dots, e_k], [h_1, \dots, h_o], l \rangle$ be the type of construction utility for creating meshes. Then Table VI lists the operations defined on these types.

V. IMPLEMENTATION

The AllScale API is based on C++, which allows the reuse of existing tools such as debuggers, and makes heavy use of template-based meta-programming. This built-in language feature of C++ enables the scripted generation of code during compilation. Widely utilized examples include the generation of data structures like vectors, sets, or maps specialized to specific type parameters. However, the capabilities of this feature reach much further. It also enables the generic implementation of primitives, where a single primitive may cover a wide range of use cases, without the introduction of any abstraction overhead. All primitives of the AllScale Core API are generic primitives, making heavy use of C++ meta-programming features for the automated synthetization of program code. The same applies for all AllScale User API constructs, to improve their (re-)usability and flexibility.

In addition, the standard toolchain implementation of the API only requires a C++14-compliant compiler and standard library (e.g. recent versions of GCC, Clang, Apple-Clang, and Visual Studio), and hence supports application development on at least three different operating systems (Linux, OS X, Windows). In order to mitigate the initial adoption barrier of porting applications to AllScale, an SDK comprising a build system infrastructure and setup scripts is provided³.

³https://github.com/allscale/allscale_sdk

TABLE V
OPERATORS DEFINED ON *Adaptive Grid* DATA STRUCTURES.

Name	Type	Description
create	$\left(\begin{array}{c} \text{type}\langle\alpha\rangle, \\ \text{int}^n, \\ \text{seq}\langle r_1, \dots, r_l \rangle \end{array} \right) \rightarrow AGrid\langle\alpha, n, [r_1, \dots, r_l]\rangle$	creates a new n-dimensional adaptive grid with element type α of the given size and grid cell structure
destroy	$(AGrid\langle\alpha, n, [r_1, \dots, r_l]\rangle) \rightarrow \text{unit}$	deletes the given adaptive grid
read	$\left(\begin{array}{c} AGrid\langle\alpha, n, [r_1, \dots, r_l]\rangle, \\ (\text{int}^n)^+ \end{array} \right) \rightarrow \alpha$	reads the element from the given grid at the specified hierarchical coordinates
write	$\left(\begin{array}{c} AGrid\langle\alpha, n, [r_1, \dots, r_l]\rangle, \\ (\text{int}^n)^+, \\ \alpha \end{array} \right) \rightarrow \text{unit}$	updates the element within the given grid at the specified hierarchical coordinates
refine	$\left(\begin{array}{c} AGrid\langle\alpha, n, [r_1, \dots, r_l]\rangle, \\ (\text{int}^n)^+, \\ Grid\langle\alpha, n \rangle \end{array} \right) \rightarrow \text{unit}$	refines the resolution of a cell addressed by the given hierarchical coordinate by inserting the given grid data as refinement information
coarsen	$\left(\begin{array}{c} AGrid\langle\alpha, n, [r_1, \dots, r_l]\rangle, \\ (\text{int}^n)^+, \\ \alpha \end{array} \right) \rightarrow \text{unit}$	coarsens the resolution of a cell addressed by the given hierarchical coordinate and inserting the given value data as coarsened information
getLevel	$\left(\begin{array}{c} AGrid\langle\alpha, n, [r_1, \dots, r_l]\rangle, \\ (\text{int}^n)^+ \end{array} \right) \rightarrow \text{int}$	get the currently active resolution level at the specified hierarchical grid position
scan	$\left(\begin{array}{c} \text{int}^n, \\ \text{int}^n, \\ AGrid\langle\alpha, n, [r_1, \dots, r_l]\rangle, \\ ((\text{int}^n)^+) \rightarrow \beta \end{array} \right) \rightarrow \text{treeture}\langle\text{unit}\rangle$	applies the given function (in parallel) to all active hierarchical coordinates of the given interval in an arbitrary order

VI. USE CASES

This section presents our real-world pilot applications that build on the AllScale API. The first, iPIC3D [6], is a particle-in-cell simulation code developed together with KTH Stockholm and employs multiple *pfor* operators and 3-dimensional *Grid* data structures. The second, AMDADOS [7], is an advection-diffusion code developed together with IBM Research Ireland and uses a 2-dimensional *stencil* operator and *adaptive grid* structure. While the full implementation of these applications is available online, we only present code excerpts of the main computation here for brevity.

A. iPIC3D

The iPIC3D pilot application is an iterative particle-in-cell space weather simulation code and its main computation loop is shown in Fig. 4. Its underlying data structure is a 3-dimensional regular equidistant *Grid* (line 13) where each element is a cell representing a cuboid and maintaining a dynamically-sized list of particles (line 8) located in this cuboid. Furthermore, each particle stores physical properties such as location, velocity, charge, and mass.

In each iteration of the simulation, the physical effects of the particles are aggregated to compute a set of induced force fields (lines 21–26). These force fields are also represented by 3-dimensional *Grid* structures (lines 9 and 14). In a next step, electromagnetic field equations are solved (lines 27–29), the forces affecting each particle’s position and velocity are computed and the particles are updated accordingly (lines 35–37). Particles moving beyond the boundary of a cell need to be migrated (lines 33–35) to the respective target cell, which can be any of 26 neighbor cells. Once the migration of particles is completed, the next iteration can be computed.

These major simulation steps are all parallel update operations using higher-dimensional variations of the *pfor* operator.

```

1  unsigned numSteps = ...; // number of time steps
2  auto zero = utils::Coordinate<3>(0); // point of origin
3  auto size = ...; // size of domain
4
5  namespace alg = allscale::api::user::algorithm;
6  namespace data = allscale::api::user::data;
7
8  struct Cell { std::vector<Particle> particles; };
9  struct FieldNode { ... // electric and magnetic field components };
10 struct DensityNode { Vector3<double> J; // current density };
11
12 // 3D grids for cells, electromagnetic field and current density
13 data::Grid<Cell, 3> cells = ...;
14 data::Grid<FieldNode, 3> field = ...;
15 data::Grid<DensityNode, 3> density = ...;
16 // create a grid of buffers for density projection from particles to grid nodes
17 data::Grid<DensityNode> densityContributions(size * 2);
18
19 // run time loop for the simulation
20 for(unsigned i = 0; i < numSteps; ++i) {
21     alg::pfor(zero, size, [&](const utils::Coordinate<3>& pos) {
22         //STEP 1a: collect particle density contributions and store in buffers
23         particleToFieldProjector(cells[pos], pos, densityContributions); });
24     alg::pfor(zero, density.size(), [&](const utils::Coordinate<3>& pos) {
25         // STEP 1b: aggregate density in buffers to density nodes
26         aggregateDensityContributions(densityContributions, pos, density[pos]); });
27     alg::pfor(fieldStart, fieldEnd, [&](const utils::Coordinate<3>& pos){
28         // STEP 2: solve electromagnetic field equations
29         fieldSolver(pos, density, field); });
30     alg::pfor(zero, size, [&](const utils::Coordinate<3>& pos){
31         // STEP 3: project forces to particles and move particles
32         particleMover(cells[pos], pos, field, particleTransfers); });
33     alg::pfor(zero, size, [&](const utils::Coordinate<3>& pos){
34         // STEP 4: transfer particles into destination cells
35         transferParticles(cells[pos], pos, particleTransfers); });
36 }

```

Fig. 4. Code excerpt of main data structures and simulation loop of iPIC3D. The full code is available online at https://github.com/allscale/allscale_ipic3d.

Thus, the resulting simulation code is structured like a list of update loops, enclosed within a single time step loop.

B. AMDADOS

AMDADOS is a numerical simulation of an oil spill, with an excerpt of the main computation code is shown in Fig. 5. It is based on a 2-dimensional stencil (lines 20–44) and

TABLE VI
OPERATORS DEFINED ON *Mesh* BUILDER, *Mesh* STRUCTURE AND *Mesh* DATA.

Name	Type	Description
Mesh Builder		
create	$(type\langle Mesh\langle n, e, h, l \rangle \rangle) \rightarrow MBuilder\langle n, e, h, l \rangle$	creates a builder for the given mesh type; initially, the mesh under construction is left empty
destroy	$(MBuilder\langle n, e, h, l \rangle) \rightarrow unit$	destroys a builder instance
addNode	$\left(\begin{array}{c} MBuilder\langle [n_1, \dots, n_m], e, h, l \rangle, \\ type\langle n \rangle, \\ type\langle i \rangle \end{array} \right) \rightarrow id\langle n, i \rangle$	creates a new node of the given type n on the given level i within the mesh under construction
link	$\left(\begin{array}{c} MBuilder\langle n, [(n_{1a}, n_{1b}), \dots, (n_{ka}, n_{kb})], h, l \rangle, \\ id\langle n_{ia}, j \rangle, \\ id\langle n_{ib}, j \rangle \end{array} \right) \rightarrow unit$	ads an edge to the mesh under construction
link	$\left(\begin{array}{c} MBuilder\langle n, e, [(n_{1a}, n_{1b}), \dots, (n_{oa}, n_{ob})], l \rangle, \\ id\langle n_{ia}, j + 1 \rangle, \\ id\langle n_{ib}, j \rangle \end{array} \right) \rightarrow unit$	adds a hierarchical edge to the mesh under construction
toMesh	$(MBuilder\langle n, e, h, l \rangle) \rightarrow Mesh\langle n, e, h, l \rangle$	obtains a copy of the mesh under construction
Mesh Structure		
store	$(Mesh\langle n, e, h, l \rangle) \rightarrow byte^*$	serializes a mesh into a byte array
load	$\left(\begin{array}{c} byte^*, \\ type\langle Mesh\langle n, e, h, l \rangle \rangle \end{array} \right) \rightarrow Mesh\langle n, e, h, l \rangle$	deserializes a given byte array into a mesh
destroy	$(Mesh\langle n, e, h, l \rangle) \rightarrow unit$	destroys the given mesh
getNeighbors	$\left(\begin{array}{c} Mesh\langle n, [(n_{1a}, n_{1b}), \dots, (n_{ka}, n_{kb})], h, l \rangle, \\ type\langle n_{ia}, n_{ib} \rangle, \\ id\langle n_{ia}, j \rangle \end{array} \right) \rightarrow id\langle n_{ib}, j \rangle^*$	obtains a list of neighbors of a given node following a given kind of edge
getParents	$\left(\begin{array}{c} Mesh\langle n, e, [(h_{1a}, h_{1b}), \dots, (h_{oa}, h_{ob})], l \rangle, \\ type\langle h_{ia}, h_{ib} \rangle, \\ id\langle h_{ia}, j \rangle \end{array} \right) \rightarrow id\langle h_{ib}, j + 1 \rangle^*$	obtains a list of parents of a given node following a given kind of hierarchical edge
getChildren	$\left(\begin{array}{c} Mesh\langle n, e, [(h_{1a}, h_{1b}), \dots, (h_{oa}, h_{ob})], l \rangle, \\ type\langle h_{ia}, h_{ib} \rangle, \\ id\langle h_{ib}, j \rangle \end{array} \right) \rightarrow id\langle h_{ia}, j - 1 \rangle^*$	obtains a list of children of a given node following a given kind of hierarchical edge
scan	$\left(\begin{array}{c} Mesh\langle [n_1, \dots, n_m], e, h, l \rangle, \\ type\langle n_i \rangle, \\ type\langle j \rangle, \\ (id\langle n_i, j \rangle) \rightarrow \beta \end{array} \right) \rightarrow treeture\langle unit \rangle$	applies a given operation to every instance of a selected node type on a selected level within the given mesh
scan	$\left(\begin{array}{c} Mesh\langle n, [(n_{1a}, n_{1b}), \dots, (n_{ka}, n_{kb})], h, l \rangle, \\ type\langle (n_{ia}, n_{ib}) \rangle, \\ type\langle j \rangle, \\ (id\langle n_{ia}, j \rangle, id\langle n_{ib}, j \rangle) \rightarrow \beta \end{array} \right) \rightarrow treeture\langle unit \rangle$	applies a given operation to every instance of a selected edge type on a selected level within the given mesh
scan	$\left(\begin{array}{c} Mesh\langle n, e, [(n_{1a}, n_{1b}), \dots, (n_{oa}, n_{ob})], l \rangle, \\ type\langle (n_{ia}, n_{ib}) \rangle, \\ type\langle j \rangle, \\ (id\langle n_{ia}, j + 1 \rangle, id\langle n_{ib}, j \rangle) \rightarrow \beta \end{array} \right) \rightarrow treeture\langle unit \rangle$	applies a given operation to every instance of a selected hierarchical edge type on a selected pair of adjacent levels within the given mesh
Mesh Data		
create	$\left(\begin{array}{c} Mesh\langle [n_1, \dots, n_m], e, h, l \rangle, \\ type\langle n_i \rangle, \\ type\langle j \rangle, \\ type\langle \alpha \rangle \end{array} \right) \rightarrow MData\langle n_i, j, \alpha \rangle$	creates an attribute storage associating a value of type α to each node instance of type n_i on layer j present in the given mesh
destroy	$(MData\langle n, l, \alpha \rangle) \rightarrow unit$	deletes the given attribute storage
read	$\left(\begin{array}{c} MData\langle n, l, \alpha \rangle, \\ id\langle n, l \rangle \end{array} \right) \rightarrow \alpha$	retrieves the value of an attribute associated to the given node from the given attribute store
write	$\left(\begin{array}{c} MData\langle n, l, \alpha \rangle, \\ id\langle n, l \rangle, \\ \alpha \end{array} \right) \rightarrow unit$	updates the value of an attribute associated to the given node in the given attribute store

incorporates data assimilation events (line 28) using external sensor data (line 16) in order to mitigate simulation errors. The basic data structure of this application is a regular, *Adaptive Grid* (line 18). The number of refinement levels is a compile-time constant and can be hard coded within the application (lines 4–8). However, coarsening and refinement steps are applied dynamically during execution based on the state of the simulation as well as data assimilation events (inside the functions called in lines 26 and 28, not shown in detail here).

The resolution refinement follows a hierarchical pattern. On the top level, a fixed size, regular 2D grid defines the domain

of the overall simulated area. Each of these top-level cells (also called sub-domains) may then be itself recursively subdivided into small regular grids, up to a statically fixed maximum resolution. The simulation algorithm updates each sub-domain independently for a single time step at the currently active level of resolution. This update operation may take several iterations, yet does not necessitate the exchange of any information with neighboring sub-domains. Once complete, boundary information is exchanged between adjacent sub-domains. Thus, sub-domains being n top-level cell-widths apart may be n time steps apart in their simulation time.

```

1 namespace alg = algorithm;
2 namespace data = data;
3
4 typedef data::CellConfig<2, data::layers<
5     data::layer<1,1>, // layer 2, size 1 x 1
6     data::layer<8,8>, // layer 1, size 8 x 8
7     data::layer<2,2> // layer 0, size 16 x 16
8 >> subdomain_config_t;
9
10 using subdomain_t = data::AdaptiveGridCell<double, subdomain_config_t>;
11 using domain_t = data::Grid<subdomain_t, 2>;
12
13 struct Sensor { ... };
14
15 // 2D grid of sensor data
16 const data::Grid<Sensor, 2> sensors = ...;
17 const size_t Nt = ...; // number of time steps
18 domain_t state_field = ...; // 2D grid of sub-domains constitutes entire domain
19
20 alg::stencil(state_field, Nt,
21     [&,Nt](time_t t, const point2d_t& idx, const domain_t& state)
22     -> const subdomain_t
23     { // Computation of subdomains
24         subdomain_t temp_field;
25         if(contexts[idx].sensors.empty()) // subdomain without sensors
26             SubdomNoSensors(t, state, temp_field, contexts[idx], idx);
27         else // subdomain with sensors, assimilation occurs
28             SubdomKalman(sensors[idx], t, state, temp_field, contexts[idx], idx);
29         return temp_field;
30     },
31     alg::observer( // Monitoring: periodically write full subdomain to file
32         [=](time_t t) { return (t % output_interval == 0); // time filter },
33         [(const point2d_t& idx) { return true; // Space filter: no constraints },
34         // Append a full field to the file of simulation results.
35         [&](time_t t, const point2d_t& idx, const subdomain_t& cell) {
36             cell.forAllActiveNodes([&](const point2d_t& loc, double val) {
37                 point2d_t glo = computeGlobalIndex(loc, idx);
38                 out_stream.atomic([&](auto& file) {
39                     file << ... << "\n"; // write output data
40                 });
41             });
42         })
43     );
44 );

```

Fig. 5. Code excerpt of the main stencil computation of AMDADOS. The full code is available online at https://github.com/allscale/allscale_amdados.

While this application could be implemented using the *pfor* operator, this would lead to a flat parallelism structure with synchronization enforced at the end of each time step. For this reason, it utilizes the *stencil* operator instead, which exposes recursive space-time decomposition and allows multiple time steps on spatially sufficiently separated sub-domains to be computed in parallel. In addition, it shows the observer functionality of the *stencil* operator, which allows for time- and space-controlled output of the simulation state.

C. FINE/Open

The FINE/Open application is a computational fluid dynamics (CFD) solver developed by NUMECA [8]. The underlying data structure is a static, unstructured *Mesh* comprising objects such as cells, faces, edges, nodes, or boundary faces. The computation is based on a *vcycle*, which is further detailed below.

While it is not possible to show the actual implementation of FINE/Open due to non-disclosure concerns, Fig. 6 shows a basic code example which is built on-top of the same *Mesh* data structure and also performs computations using the *vcycle* operator. The geometric information is covered by a list of relations connecting cells, faces and vertices (line 3 and lines

```

1 using namespace allscale::api::user;
2 // -- elements --
3 struct Cell {}; struct Face {}; struct Vertex {};
4 // -- edges --
5 struct Cell2Vertex : public data::edge<Cell, Vertex> {};
6 struct Cell2Face_In : public data::edge<Cell, Face> {};
7 struct Cell2Face_Out : public data::edge<Cell, Face> {};
8 struct Face2Cell_In : public data::edge<Face, Cell> {};
9 struct Face2Cell_Out : public data::edge<Face, Cell> {};
10 // -- hierarchical edges --
11 struct Parent2Child : public data::hierarchy<Cell,Cell> {};
12 // -- property data --
13 struct CellTemperature : public data::mesh_property<Cell,double> {};
14 struct FaceArea : public data::mesh_property<Face,value_t> {};
15 struct FaceVolRatio : public data::mesh_property<Face,value_t> {};
16 struct FaceConductivity : public data::mesh_property<Face,double> {};
17 // -- define the mesh and builder --
18 template<unsigned levels = 1>
19 using MeshBuilder = data::MeshBuilder<
20     data::nodes<Cell, Face, Vertex>,
21     data::edges<Cell2Vertex, Cell2Face_In, Cell2Face_Out, Face2Cell_In,
22         Face2Cell_Out>,
23     data::hierarchies<Parent2Child>,
24     levels
25 >;
26 // -- type of a mesh --
27 template<unsigned levels = 1, unsigned PDepth = P_DEPTH>
28 using Mesh = typename MeshBuilder<levels>::template mesh_type<PDepth>;
29 // -- type of the properties of a mesh --
30 template<typename Mesh>
31 using MeshProperties = data::MeshProperties<Mesh::levels,
32     typename Mesh::partition_tree_type,
33     CellTemperature, FaceConductivity, VertexPosition
34 >;
35 // V-Cycle stage
36 template<typename Mesh, unsigned Lvl>
37 struct TemperatureStage {
38     const Mesh& mesh; // mesh structure, properties and cell data
39     MeshProperties<Mesh>& properties;
40     attribute<Cell, value_t> temperature;
41     attribute<Face, value_t> fluxes;
42 };
43 void jacobiSolver() {
44     auto& fConductivity = properties.template get<FaceConductivity, Lvl>();
45     auto& fArea = properties.template get<FaceArea, Lvl>();
46     auto& fVolumeRatio = properties.template get<FaceVolRatio, Lvl>();
47
48     mesh.template pforAll<Face, Lvl>([&](auto f) { // compute per-face flux
49         auto in = mesh.template getNeighbor<Face2Cell_In>(f);
50         auto out = mesh.template getNeighbor<Face2Cell_Out>(f);
51         value_t gradTemp = temperature[in] - temperature[out];
52         fluxes[f] = fVolumeRatio[f] * fConductivity[f] * fArea[f] * gradTemp;
53     });
54
55     mesh.template pforAll<Cell, Lvl>([&](auto c) { // update per-cell solution
56         auto subtractingFaces = mesh.template getNeighbors<Face2Cell_In>(c);
57         for(auto sf : subtractingFaces) { temperature[c] -= fluxes[sf]; }
58         auto addingFaces = mesh.template getNeighbors<Face2Cell_Out>(c);
59         for(auto af : addingFaces) { temperature[c] += fluxes[af]; }
60     });
61 }
62
63 void computeFineToCoarse() { ... }
64 void computeCoarseToFine() { ... }
65 void restrictFrom(TemperatureStage<Mesh,Lvl-1>& childStage) { ... }
66 void prolongateTo(TemperatureStage<Mesh,Lvl-1>& childStage) { ... }
67 };

```

Fig. 6. Code excerpt of a sample application using the *vcycle* operator on a multi-grid mesh. The full version is available online at <https://git.io/fjBTq>.

5–9) with each other (e.g. a relation of a cell to its faces). These relations can be easily navigated by templated member functions (lines 49–50, 56 and 58). Note that the type system of C++ ensures proper object navigation during compile-time (e.g. attempting to illegally access a vertex from a face would result in a compiler error). Furthermore, for each object, a set of properties influencing the simulation is maintained (lines

TABLE VII
PILOT APPLICATION CODE METRICS.

	AMDADOS		iPIC3D	
	AllScale	MPI	AllScale	MPI
P.SLOC	25	70	23	56
SLOC	1136	1420	1443	1717
TOT CY	154	181	204	264

13–16). These may comprise static information like e.g. the volume of a cell or dynamic information such as the heat flow through a face. The latter is the state of the conducted simulation and the result end users are interested in. Finally, to aid the effective computation of the desired solution, multiple meshes describing the same objects in different resolutions are combined into a hierarchy of meshes to enable the use of multi-grid solvers (lines 43–66). The hierarchy of meshes can be navigated via hierarchical edges (line 11).

In each simulation step, updates to the various properties associated to the mesh objects are conducted. Updates start in the mesh layer exhibiting the finest resolution. Thereby, physical effects are propagated through the connections between the various objects on this layer. After a fixed number of iterations, the current state of the simulated properties are aggregated and projected to the next coarser-grained level of the hierarchical mesh. There, the same propagation and aggregation operations are repeated. After completing updates on the coarsest layer, modifications are projected recursively down towards the finer layers and the simulation continues with the next time step.

VII. EVALUATION

A. Productivity

Table VII lists absolute values for code metrics collected for the AllScale and MPI versions of iPIC3D and AMDADOS, in order to get a grasp on the productivity of working with the AllScale API compared to MPI. *P.SLOC* denotes the lines of code spent on parallelizing the application, counting only the minimal set of lines containing explicit interface calls. Additional code required for e.g. preparing arguments is not accounted for, and hence these results present a best-case perspective for MPI. Nevertheless, as the numbers show, AllScale clearly outperforms MPI. Furthermore, we have measured the sum of the cyclomatic complexity [9] (*TOT CY*) as well as the total count of non-comment lines of code (*SLOC*) over all translation units. Compared to *P.SLOC*, these give an indication of how much of the overall pilot application complexity is related to their manual MPI parallelization rather than actual domain science content.

B. Performance

In order to ascertain the performance of the AllScale API and the underlying AllScale toolchain, we conducted weak scaling experiments for AMDADOS and iPIC3D on the Vienna Scientific Cluster (VSC) 3 and the Meggie cluster of the University of Erlangen-Nuremberg. Table VIII lists their hardware characteristics. The initial problem size for a single node was chosen such that application throughput did not noticeably improve when further increasing the problem size.

Fig. 7 compares the performance results of the AllScale implementations against MPI reference implementations, with performance measured as application throughput. The figure illustrates that for AMDADOS, both AllScale and MPI show degrading performance for increasing node counts, however the AllScale implementation clearly outperforms the MPI variant (up to 160% higher performance for 256 nodes). Note that the performance for MPI for 512 nodes is missing. This is due to the fact that we set the job time limit to 30 times the execution time of a single node run, and the MPI experiments for 512 nodes exceeded this time limit. For iPIC3D, AllScale is on-par with MPI, both showing stable, almost linearly scaling performance for higher node counts.

These results demonstrate the feasibility of using automatically managed user-defined data structures in large-scale high performance applications.

VIII. RELATED WORK

Conventional, low-level HPC infrastructures comprising combinations of MPI with some per-node parallelism APIs are still the default platforms for building HPC applications, but require programmers to manually implement workload decomposition. Systems providing a higher level of abstraction, such as the AllScale API, can be grouped into three broad categories: new general purpose languages, domain-specific frameworks, and general purpose libraries. Note that there are a large number of parallelism approaches constrained to single-node shared memory hardware. We omit these from our overview provided here as they do not address the same problem space as the AllScale API.

In terms of languages, X10 [10] and Chapel [11] have targeted (recursive) parallelism on large scale, distributed systems, but left locality and data management to the user. Charm++ [12], on the other hand, is a C++ extension aiming at isolating the user from low-level mapping activities, thus facilitating portability. Its design is based on message-exchanging entities exposed to the user and lacks automated data distribution management. Recently, the ANTAREX research project [13] proposed a DSL-based approach, facilitating the separation of concerns between functional and non-functional aspects of HPC applications. However, due to its DSL-focused design, users require additional tools and may not rely on the experience of an established developer community.

Several new frameworks such as Lift [14], Delite [15], or AnyDSL [16] provide environments for implementing DSLs. Internally, DSL constructs are encoded using functional IR constructs like `map`, `reduce`, or `zip`. However, the resulting programming interface for the domain experts remains a DSL, targeting very specific application domains and inheriting the difficulties of DSLs noted above.

Domain-specific, C++-based libraries such as PETSc [2] or TensorFlow [17] handle several of the challenges addressed by our framework successfully for their respective domains. However, they are tailored towards specific domains instead of supporting a wider range of applications.

TABLE VIII
EXPERIMENTAL PLATFORM DESCRIPTION. THE NUMBER OF NODES REFERS TO THE MAXIMUM USED IN THIS WORK.

Name	#Nodes	CPU	RAM	Interconnect	Compiler	MPI Implementation
VSC-3	512	2x Intel Xeon E5-2650 v2	64 GB	Intel Infiniband QDR-80	GCC 7.2.0	OpenMPI 3.0.0
Meggie	256	2x Intel Xeon E5-2630 v4		Intel OmniPath 100 GBit/s	GCC 7.3.0	Intel MPI 2018.2

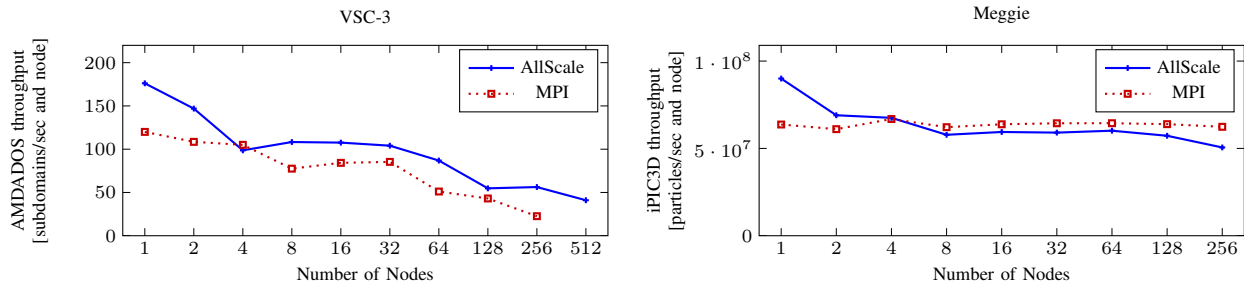


Fig. 7. Comparison of throughput per node for AMDADOS on VSC-3 and iPIC3D on Meggie for MPI and AllScale.

More general purpose parallel C++ library based frameworks like STAPL [18] and Kokkos [19] are exercising control over parallel algorithms and data structures similar to our architecture. STAPL envisions a separation of concerns strategy similar to ours. Kokkos, on the other hand, has a strong focus on multidimensional arrays and parallel loops, unlike the wider range of data structures and operations supported by our architecture. Due to a lack of compiler integration, these approaches require data dependencies of code regions to be expressed explicitly as part of the API, while this is covered implicitly in our approach.

IX. CONCLUSION

This work presented the AllScale API, a novel interface for implementing distributed memory parallel applications with the programmability of a shared memory API. We illustrated how the distinction into the User and Core components provides a separation of concerns for domain experts, HPC experts and system-level experts, and discussed several constructs of the AllScale API in detail. In addition, the three use cases presented show the suitability of our approach to real-world scientific problems, evaluated in both productiveness and parallel performance.

Future work includes better user feedback for programming errors, additional pre-provided operators in the User API along with new applications.

REFERENCES

- [1] G. Thorson and M. Woodacre, “SGI UV2: A fused computation and data analysis machine,” in *Proc. of the Int. Conf. on High Performance Computing, Networking, Storage and Analysis 2012*. Los Alamitos, CA, USA: IEEE Computer Society Press, 2012, pp. 105:1–105:9.
- [2] S. Balay, S. Abhyankar, M. Adams, P. Brune, K. Buschelman *et al.*, “PETSc Users Manual Revision 3.7,” Argonne National Lab.(ANL), Argonne, IL (United States), Tech. Rep., 2016.
- [3] H. Jordan, T. Heller, P. Gschwandtner *et al.*, “The allscale runtime application model,” in *2018 IEEE International Conference on Cluster Computing (CLUSTER)*, Sep. 2018, pp. 445–455.
- [4] H. Jordan, P. Thoman, P. Zangerl, T. Heller, and T. Fahringer, “A context-aware primitive for nested recursive parallelism,” in *Euro-Par 2016: Parallel Processing Workshops*. Cham: Springer International Publishing, 2017, pp. 149–161.
- [5] H. Jordan, R. Iakymchuk, T. Fahringer, P. Thoman, T. Heller *et al.*, “D2.4 - allscale system architecture,” 5 2018. [Online]. Available: http://www.insieme-compiler.org/pub/allscale_D2.4_system_architecture.pdf
- [6] S. Markidis, G. Lapenta, and Rizwan-uddin, “Multi-scale Simulations of Plasma With iPIC3D,” *Math. Comput. Simul.*, vol. 80, no. 7, Mar. 2010.
- [7] A. Akhriev, F. O’Donncha, P. Gschwandtner, and H. Jordan, “A localised data assimilation framework within the ‘AllScale’ parallel development environment,” in *OCEANS 2018 MTS/IEEE Charleston*, Oct 2018, pp. 1–7.
- [8] NUMECA International, “NUMECA International,” 2019. [Online]. Available: <https://www.numeca.com>
- [9] T. J. McCabe, “A Complexity Measure,” *IEEE Transactions on Software Engineering*, vol. SE-2, no. 4, pp. 308–320, 1976.
- [10] P. Charles, C. Grothoff, V. Saraswat *et al.*, “X10: An Object-oriented Approach to Non-uniform Cluster Computing,” in *Proc. of the 2005 ACM SIGPLAN Conf. on Object-oriented Programming, Systems, Languages, and Applications*. New York, NY, USA: ACM, 2005, pp. 519–538.
- [11] B. L. Chamberlain, D. Callahan, and H. P. Zima, “Parallel Programmability and the Chapel Language,” *The International Journal of High Performance Computing Applications*, vol. 21, no. 3, pp. 291–312, 2007.
- [12] L. V. Kale and S. Krishnan, “CHARM++: A Portable Concurrent Object Oriented System Based on C++,” in *Proc. of the Conf. on Object-oriented Programming Systems, Languages, and Applications*, ser. OOPSLA. New York, NY, USA: ACM, 1993, pp. 91–108.
- [13] C. Silvano, G. Agosta, S. Cherubin *et al.*, “The ANTAREX Approach to Autotuning and Adaptivity for Energy Efficient HPC Systems,” in *Proc. of the Int. Conf. on Computing Frontiers*. New York, NY, USA: ACM, 2016, pp. 288–293.
- [14] M. Steuwer, T. Rimmelg, and C. Dubach, “Lift: A Functional Data-parallel IR for High-performance GPU Code Generation,” in *Code Generation and Optimization, Int. Symp. on*. IEEE, 2017, pp. 74–85.
- [15] K. J. Brown, A. K. Sujeeth, H. J. Lee, T. Rompf, H. Chafi *et al.*, “A Heterogeneous Parallel Framework for Domain-specific Languages,” in *Parallel Architectures and Compilation Techniques (PACT), Int. Conf. on*. IEEE, 2011, pp. 89–100.
- [16] Saarland University, “AnyDSL,” 2018. [Online]. Available: <https://anydsl.github.io>
- [17] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro *et al.*, “Tensorflow: Large-scale Machine Learning on Heterogeneous Distributed Systems,” *arXiv preprint arXiv:1603.04467*, 2016.
- [18] P. An, A. Jula, S. Rus, S. Saunders, T. Smith *et al.*, “STAPL: An Adaptive, Generic Parallel C++ Library,” in *Int. Workshop on Languages and Compilers for Parallel Computing*. Springer, 2001, pp. 193–208.
- [19] H. C. Edwards, C. R. Trott, and D. Sunderland, “Kokkos: Enabling Manycore Performance Portability Through Polymorphic Memory Access Patterns,” *Journal of Parallel and Distributed Computing*, vol. 74, no. 12, pp. 3202–3216, 2014.