

H2020 FETHPC-1-2014



**An Exascale Programming, Multi-objective Optimisation and Resilience
Management Environment Based on Nested Recursive Parallelism**

Project Number 671603

D3.2 – API Implementation for Recursive Parallelism (b)

*WP3: High-level parallel API and API-aware
optimizing source-to-source compiler*

Version: 1.2

Author(s): Herbert Jordan (UIBK), Philipp Gschwandtner (UIBK)

Date: 30/06/17



D3.2 – API Implementation for Recursive Parallelism (b)

Due date:	PM21
Submission date:	30/06/2017
Project start date:	01/10/2015
Project duration:	36 months
Deliverable lead organization	UIBK
Version:	1.2
Status	Final
Author(s):	Herbert Jordan (UIBK), Philipp Gschwandtner (UIBK)
Reviewer(s)	Emanuele Ragnoli (IBM), Roman Iakymchuk (KTH)

Dissemination level	
PU	<i>Public</i>

Disclaimer

This deliverable has been prepared by the responsible Work Package of the Project in accordance with the Consortium Agreement and the Grant Agreement Nr 671603. It solely reflects the opinion of the parties to such agreements on a collective basis in the context of the Project and to the extent foreseen in such agreements.

Acknowledgements

The work presented in this document has been conducted in the context of the EU Horizon 2020. AllScale is a 36-month project that started on October 1st, 2015 and is funded by the European Commission.

The partners in the project are UNIVERSITÄT INNSBRUCK (UBIK), FRIEDRICH-ALEXANDER-UNIVERSITÄT ERLANGEN NÜRNBERG (FAU), THE QUEEN'S UNIVERSITY OF BELFAST (QUB), KUNGLIGA TEKNISKA HÖGSKOLAN (KTH), NUMERICAL MECHANICS APPLICATIONS INTERNATIONAL SA (NUMECA), IBM IRELAND LIMITED (IBM).

The content of this document is the result of extensive discussions within the AllScale Consortium as a whole.

More information

Public AllScale reports and other information pertaining to the project are available through the AllScale public Web site under <http://www.allscale.eu>.

Version History

Version	Date	Comments, Changes, Status	Authors, contributors, reviewers
0.1	29/02/16	First draft	Herbert Jordan
0.2	24/03/16	Integration of review feedback	Herbert Jordan
1.0	24/03/16	Final version	Herbert Jordan
1.1	23/06/17	(b) update for M21	Philipp Gschwandtner
1.2	30/06/17	Integrated review feedback	Philipp Gschwandtner, Roman Iakymchuk

Table of Contents

Executive Summary	6
Changes with respect to D3.1 API Implementation for Recursive Parallelism (a):	6
1 Introduction	6
2 Overview.....	7
2.1 Requirements.....	7
2.2 Getting the Sources.....	7
2.3 Package Organization	7
2.4 Build and Run.....	8
3 API Overview	8
3.1 AllScale Core API.....	8
3.2 AllScale User API.....	9
4 Future Work	10

Executive Summary

This document describes the status and structure of the AllScale API implementation deliverable (D3.2). It outlines the organization of the development process, instructions on how to obtain the sources, build the project, and execute test cases, as well as the internal file structure of the software project. After highlighting the changes with respect to the first version of this deliverable, D3.1, it concludes with a summary of future work.

Changes with respect to D3.1 API Implementation for Recursive Parallelism (a)

Since the first version of this deliverable, D3.1 (a), a number of changes were required to support additional features and increase development productivity when working with the AllScale API. These changes, already incorporated in this version, D3.2 (b), are listed hereafter for clarity in their order of appearance:

- Minor changes in the build system and requirements (newer cmake version required, slightly modified directory structure)
- API project repository is no longer hosted by UIBK but available publicly on the AllScale GitHub repository
- Tutorial codes are now included, demonstrating the use of the AllScale User API
- The AllScale Core API now includes input/output facilities
- The AllScale User API now includes implementations of additional operators (async, vcycle) and data items (static grid, unstructured multi-grid mesh)

1 Introduction

The deliverable D3.2 is a source code deliverable, providing implementations of the AllScale API. This document provides an overview on the internal structure and organization of the sources, as well as instructions on how to obtain and build them.

The AllScale API is implemented in two phases:

- Early prototypes of API constructs are developed to explore the design space, clarify requirements, and to investigate interface design options and usability issues. Those are developed in the context of an isolated prototype software project (*allscale-api-prototype*).
- The results obtained by working on the prototypes are then migrated to the actual AllScale API development project, covering the full functionality required for the codes to be processed by the AllScale toolchain.

The prototype stage enabled us to start investigating the design of the User API and the structure and requirements of the pilot applications long before the interfaces towards underlying components, in particular the AllScale compiler

D3.2 – API Implementation for Recursive Parallelism (b)

and the AllScale runtime system, have been established as a first draft. Also, relaxed requirements on the implementations lead to faster development cycles.

Within this report, however, the structure and properties of the second stage, the actual AllScale API development project (D3.2), are covered.

2 Overview

The AllScale API package provides the interfaces, reference implementations, use case demonstrations, and unit tests for the AllScale Core API, and the AllScale User API.

2.1 Requirements

The sources of this package are using features specified by the C++14 standard. Thus, a compliant C++ compiler (e.g. GCC ≥ 4.9) is required. For the build and testing process, cmake (version ≥ 3.5) is needed. Furthermore, for the build process an internet connection to download the sources for the utilized unit-testing solution (googletest) is required. No additional dependencies to third-party components are instated.

2.2 Getting the Sources

The sources are provided publicly on the AllScale GitHub¹ repository and can be obtained utilizing git as follows

```
git clone git@github.com:allscale/allscale_api.git
```

This command will download the latest version of the sources.

2.3 Package Organization

The directory and file structure of the package are organized as follows:

- <project root>
 - code – all headers, sources, unit tests, and build scripts
 - api – all headers, sources and tests related to the API
 - include/allscale/api – core and user API headers
 - src – sources
 - test – unit tests for the core and user API
 - tutorials/src – example codes demonstrating API usage
 - utils – general utilities
 - include/allscale/utils – utilities headers
 - test – unit tests for utilities
 - CMakeLists.txt – cmake build script
 - README.md – short summary and usage hints
 - createDebug.sh / createRelease.sh – project initialization scripts

¹ https://github.com/allscale/allscale_api
Copyright © AllScale Consortium Partners 2017

D3.2 – API Implementation for Recursive Parallelism (b)

2.4 Build and Run

To build the project, change to your local <project root> directory and create a debug build as follows:

```
mkdir build_debug
cd build_debug
. ../createDebug.sh
make -j8
```

The script *createDebug.sh* uses *cmake* to create a build environment. It uses *g++* as its default compiler and may be customized for other toolchains.

To run the test cases and prototypes within the project, run

```
make test ARGS=-j8
```

within the build directory or execute individual unit tests or prototypes by running the corresponding *ut_** executable directly.

To build and run a release build, run the commands

```
mkdir build_release
cd build_release
. ../createRelease.sh
make -j8
make test ARGS=-j8
```

within the local project root directory.

3 API Overview

This section briefly outlines the AllScale API implementation covered by Deliverable D3.2. For a detailed specification of the API we refer to D2.5 – AllScale API Specification.

3.1 AllScale Core API

The AllScale Core API covers three essential components:

- a higher-order function creating parallel recursive operations (*prec*)
- futures for synchronizing parallel recursive tasks
- an abstract *Data Item* interface for data objects being distributed and managed by the AllScale runtime system
- Input/Output utilities for reading and writing persistent data

The following table lists those central elements of the AllScale Core API and the corresponding files relative to the directory <project_root>/code/api/include/allscale/api/core :

Construct	Header	Reference Implementation
-----------	--------	--------------------------

D3.2 – API Implementation for Recursive Parallelism (b)

prec	prec.h	prec.h
decomposable futures	future.h	impl/reference/runtime.h
data items	data.h	data.h
input/output	io.h	impl/reference/io.h

3.2 AllScale User API

The following table lists the essential elements of the AllScale User API and the corresponding files relative to the directory `<project_root>/code/api/include/allscale/api/user` :

Construct	Header
async	operator/async.h
parallel loop	operator/pfor.h
stencil	operator/stencil.h
adaptive grid stencil*	operator/adaptive_grid_stencil.h
Vcycle	operator/vcycle.h
distributed key/value map	data/map.h
n-D static grid	data/static_grid.h
n-D dynamic grid	data/grid.h
n-D adaptive grid*	data/adaptive_grid.h
unstructured multi-grid mesh	data/mesh.h

Note: prototype implementations of most User API constructs have been developed in the context of the API prototype project to be obtained from

```
git@goedis.dps.uibk.ac.at:herbert.jordan/allscale_api-prototype.git
```

Currently (mid-June 2017), the remaining constructs from the API prototype project – marked by a star (*) – are migrated to the AllScale API project covered by this deliverable.

4 Future Work

Currently, the development of the AllScale API is focusing on porting the remaining implementations of various constructs from the relaxed prototype stage to the fully-functional development project. The main target here is the porting of the adaptive grid as well as the associated operators. This phase will be concluded by its successful usage in the AMDADOS pilot application.

Future steps will target the provisioning of improved data structures for the pilot codes, including improved implementations of collections (e.g. particle list for iPIC3D), memory efficient implementations of the adaptive grid (AMDADOS) and efficiently distributable, cache friendly mesh data structures (Fine/Open).

At the same time, the results of the development of the runtime interface, in particular its requirements on data objects to be managed by the system, will be integrated into the code base.