H2020 FETHPC-1-2014

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

Project Number 671603

D6.1 – AllScale Computing Infrastructure

WP6: Integration, testing and pilot applications

Version: 1.0
Author(s): Roman Iakymchuk (KTH), Kiril Dichev (QUB), Khalid Hasanov (IBM), Arne Hendricks (FAU), Thomas Heller (FAU), Erwin Laure (KTH)
Date: 27/03/17

Due date:	PM18
Submission date:	31/03/2017
Project start date:	01/10/2015
Project duration:	36 months
Deliverable lead organization	KTH
Version:	1.0
Status	Final
Author(s):	Roman Iakymchuk (KTH), Kiril Dichev (QUB), Khalid Hasanov (IBM), Arne Hendricks (FAU), Thomas Heller (FAU), Erwin Laure (KTH)
Reviewer(s)	Peter Zangerl (UIBK) and Emanuele Ragnoli (IBM)

Dissemination level	
PU	

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 (NUMEXA), 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	31/01/17	Document structure	Roman Iakymchuk
0.2	20/02/17	Initial information from partners regarding their computing infrastructure	Kiril Dichev, Thomas Heller, Philip Gschwandtner, Roman Iakymchuk
0.3	26/02/17	Introduction and executive summary	Roman Iakymchuk
0.4	27/02/17	FAU local infrastructure	Arne Hendricks
0.5	02/03/17	IBM local infrastructure	Khalid Hasanov
0.6	02/03/17	Draft for internal review	Roman Iakymchuk
0.6	13/03/17	Review	Emanuele Ragnoli
0.7	24/03/17	Addressed reviewers' comments	Roman Iakymchuk
1.0	*27/03/17*	*Final version*	*Roman Iakymchuk*

Table of Contents

Executive Summary

This document establishes and characterizes the AllScale computing infrastructure comprising a range of computing systems – from various small-scale prototype systems to Tier-1 supercomputers – provided by different partners. Furthermore, the document outlines our plans to prepare PRACE Tier-0 and XSEDE applications for obtain access different architectures (e.g. the Xeon-PHI-based “Stampede” system), which will also allow large validation runs. The consortium members are enabled, upon request, to access the target partners’ architectures remotely and are provided with support in the deployment and testing of their software.

1 Introduction

AllScale's focus on nested parallelism will provide the necessary flexibility to map and adapt applications to a wide range of parallel architectures including future Exascale systems. AllScale aims at seamless scaling of parallel applications across all of those architectures by scaling both up, on any number of any type of hardware unit (vector, core, socket, node, cabinet) of each level of the hierarchy, and out, across levels of the hierarchy. This will be particularly important for Exascale architectures, which will likely comprise several levels of hierarchical parallelism, including vector units, multi-threaded multi-core CPUs and many-core accelerators. To validate the AllScale approach, access to large-scale computing infrastructures as well as innovative hardware prototypes is required. AllScale follows a two-tiered approach to ensure the availability of appropriate hardware resources: partners' and external systems. Hence, the partner computing systems are complemented by PRACE Tier-0 and XSEDE applications for access to external supercomputers.

This document provides in Section 2 characteristics of partners' computing infrastructure and describes their access procedures; and outlines in Section 3 the desired external computing infrastructure.

2 AllScale Partners' Computing Infrastructure

This section is focused on presenting partners' computing infrastructure. After completing the corresponding procedures, in most cases, the access to partners' systems can be granted within, at most, few days.

2.1 KTH: The PDC Center for High-Performance Computing

For access to Beskow (Cray XC40):

1. Register on the SUPR (https://supr.snic.se/)
    a. Click the "Register New Person" button
    b. If you have SWAMID, then register with it. Otherwise, follow the "Register without SWAMID" procedure
2. Become a member of the corresponding project on the SUPR: *SNIC 2016/1-257*
3. Once you have your project code, apply for an account at PDC following the "Apply using SUPR" procedure (https://www.pdc.kth.se/support/accounts/user)

For access to Tegner (a pre/post processing machine equipped with NVIDIA GPUs): Please contact Roman Iakymchuk (riakymch@kth.se)

Details regarding the Beskow supercomputer configuration can be found on the PDC web page as follows

1. Hardware: https://www.pdc.kth.se/resources/computers/beskow/hardware
2. Software: https://www.pdc.kth.se/resources/computers/beskow/software

2.2 UIBK

For access to Ortler (a shared memory system equipped with 4x Intel Xeon E5-4650 and 256 GB of RAM): Please contact Philipp Gschwandtner (philipp@dps.uibk.ac.at).

2.3 IBM

IBM is striving to make it possible for the AllScale runtime and/or applications to be tested and validated on Power8/+ systems.

2.4 QUB

Queen's University Belfast has so far provided a test server platform consisting of

- Four test nodes (run as VMs)
- A Gitlab server (repository) and a Jenkins server (build server) (both run as VMs)

In addition, Queen's can offer access to the QUB Kelvin Cluster to the consortium. It is as follows

- HP hardware
- Running CentOS
- Compute nodes: 16 dedicated to Medical School and 45 general compute nodes
- High memory nodes ranging from 128GB to 1TB of RAM
- 500TB of lustre parallel file system for scratch
- Each node has 20 cores -- Up to 900 Cores total for jobs.
- Visualisation node for OpenGL applications

For access to the QUB computing infrastructure: Please contact Kiril Dichev (K.Dichev@qub.ac.uk)

2.5 FAU

For access to FAU's computing infrastructure:

1. Register an account: Email cs3-admin@lists.fau.de with the following information:
    a. Desired email address
    b. Desired account name
    c. Desired password combination
    d. Full name and affiliation
2. Wait for account approval/confirmation email.
3. Access infrastructure by logging in to the front node as
    a. ssh yourusername@i3login.informatik.uni-erlangen.de

FAU's computing infrastructure is as follows

- Siegfried partition: 4 nodes equipped with 16-cores Intel Xeon E5-2650; 4x NVIDIA K40M plus 3x NVIDIA K20C; 1x Intel KNC
- Accel partition: 3 nodes equipped with 24-cores Intel Xeon E5-2650; 5x NVIDIA K20C
- Whistler partition: 21 nodes equipped with 2-cores AMD Opteron

2216HE
- KNC Partition: 2x Intel KNC co-processors

3 AllScale External Computing Infrastructure

This section describes the state-of-the-art large-scale external computing infrastructure. Access to these facilities can be obtained via application procedures led by corresponding partners, see below. Such procedures often require a justification of the resource usage with performance results, e.g. from small partners clusters. Given an estimated time of few months between applications submissions and their approval, we foresee to apply for these resources starting from August-September 2017. Hence, we aim to attain access to these resources and use them for our needs during the third year of the project.

3.1 QUB

Queen's can explore the access to the ARCHER supercomputer (www.archer.ac.uk), which is a Cray XC30 supercomputer. The costs should be included at the non-partner rate, and are detailed here:
http://www.archer.ac.uk/access/cost.

3.2 FAU

Through FAU, the AllScale partners also have guaranteed access to several Tier-0 systems including the “SuperMuc”, a three petaflop system at the Leibniz Supercomputing Center in Munich and the JUQUEEN supercomputer at Juelich Supercomputing Centre, which is a five petaflop system. Furthermore, AllScale will also have access to the BlueGene systems “Mira” and “Vesta” located at the Argonne Leadership Computing Facility.

3.3 PRACE and XSEDE Resources

The above-listed systems will be complemented through PRACE Tier-0 and XSEDE applications for access to different architectures (e.g. the Xeon-PHI-based “Stampede” system or the Cray XC40 based “Hornet”), which will also allow large validation runs. We will also seek collaborations with Riken, Japan, for access to the next-generation K computer as well as NUDT, China, for access to Tianhe-2, the currently most powerful system in the world.