D4.3 “Preliminary report about the choice and the first profiling and optimisation efforts on a subset of scientific applications”

Version 1.0

Document Information

Contract Number	288777
Project Website	www.montblanc-project.eu
Contractual Deadline	M24 (30 Sept 2013)
Dissemination Level	PU
Nature	Report
Author	Stéphane Requena (GENCI)
Contributors	B. Videau (IMAG), D. Brayford (LRZ), P. Lanucara (CINECA), X. Saez (BSC), R. Halver (JSC), D. Broemmel (JSC), S. Mohanty (JSC), N. Sanna (CINECA), C. Cavazzoni (CINECA), JH. Meincke (JSC), D. Komatitsch (Université de Marseille) and V. Moureau (CORIA)
Reviewer	Name : Petar Radojković (BSC)
Keywords	Exascale, scientific applications, porting, profiling, optimisation

Notices:
The research leading to these results has received funding from the European Community's Seventh Framework Programme [FP7/2007-2013] under grant agreement n° 288777.

Change Log

Version	Description of Change
V0.1	Initial draft released to the WP4 contributors
V0.2	Version released to internal reviewer
V0.3	Comments of the internal reviewer
V1.0	Final version sent to EU

Table of Contents

Executive Summary

The Mont-Blanc project aims to assess the potential of HPC clusters based on low-power embedded components to address future Exascale HPC needs.

The role of work package 4 (WP4, "Exascale applications") is to port, co-design and optimise up to 11 real exascale-class scientific applications to the different generations of Mont-Blanc hardware platforms available in order to assess the global programmability and the performance of such systems.

After the first report D4.1 "Preliminary report of progress about the porting of the full-scale scientific applications" [1] and the latest report D4.2 "Final report about the porting of the full-scale scientific applications" [2], this report aims to present the work of the last year activity of WP4 based on a selection of a subset of scientific applications suited for the Mont-Blanc architecture, and a specific work of optimisation and taskification using OmpSs/OpenCL.

The first results related to the optimisation performed on the selected set of applications are detailed in deliverable D4.2.

1 Introduction

The Mont-Blanc project aims to assess the potential of HPC clusters based on low-power embedded components to address future Exascale HPC needs.
As complement of the activities of work package 3 (WP3, "Optimized application kernels"), a part of the activities of Mont-Blanc will be to assess on the different generation of platforms made available by the project the behaviour of up to 11 "real" exacale-class scientific applications. The objective of work package 4 (WP4, "Exascale applications") will be to evaluate the global programmability and the performance (in term of time and energy to solution) of the architecture and to assess the efficiency of hybrid OmpSs/MPI programming model.

These eleven real scientific applications, used by academia and industry, running daily in production into existing European (PRACE Tier-0 systems) or national HPC facilities have been selected by the different partners in order to cover a wide range of scientific domains (geophysics, fusion, materials, particle physics, life sciences, combustion, weather forecast) as well as hardware and software needs.
Some of these applications are also part of the PRACE Benchmark (flagged with the P symbol after the name of the code):

Code	Sc. Domain	Contact	Institution
YALES2	Combustion	V. Mouveau	CNRS/CORIA
EUTERPE (P)	Fusion	X. Saez	BSC
SPECFEM3D (P)	Geophysics	D. Komatitsch	Univ. Marseille
MP2C	Multi-particle collision	G. Sutmann, A. Schiller	JSC
BigDFT	Elect. Structure	B. Videau	IMAG
Quantum Expresso (P)	Elect. Structure	C. Carvazonni and N. Sanna	CINECA
PEPC (P)	Coulomb + gravitational forces	P. Gibbon, L. Arnold	JSC
SMMP	Protein folding	J. Meinke	JSC
ProFASI	Protein folding	S. Mohanty	JSC
COSMO	Wheather forecast	P. Lancura	CINECA
BQCD (P)	Particle physics	D. Brayford	LRZ

Figure 1 - List of the 11 WP4 scientific applications

This report refers to the activities planned in WP4 under Task 4.2 and Task 4.3:

T4.2. Profiling, benchmarking and optimization (m6:m36)
Following the work performed in task 4.1, a subset of applications which finally offers the best potential for exploiting the hardware and software characteristics of these prototypes will be elaborated. This selection will include the results of WP3 and WP5 activities in term of kernel and software libraries availability/performance as well as all the others components of the software stack. On this subset of applications, dedicated optimisation efforts will be focused on effective usage of SIMD vector units or hybridisation with potential accelerators using portable programming models like OpenCL8 since some of the proposed codes have already some OpenCL versions.

T4.3. Efficiency, performance and productivity evaluation (m24:m36)
Finally, on such optimised applications, additional experiments will be conducted in order to compare power consuming and computing performance on the Mont-Blanc platform and the Tier-0 platforms of the PRACE European infrastructure. A best practice document describing productivity evaluation in term of time to port and optimise selected codes, impact of the programming models into the rewriting of source codes, portability, assessment of programming models easiness for "regular programmers", choice of metrics to evaluate the efficiency of the system will be delivered.

2 Mont-Blanc prototype

In the Supercomputing Conference 2012 held in Salt Lake City (USA) in November 2012, the Mont-Blanc project announced that the Samsung Exynos 5250 will become the reference SoC for the first Mont-Blanc prototype.
This SoC is composed by a 32nm HKMG, dual-core ARM Cortex-A15 clocked at 1.7 Ghz, 2 32-bits channels of DDR3-1600 memory, a MALI T604 GPU and an USB3 to GbE bridge.
The ARM Cortex-A15 is capable of sustaining a 6.8 GFlops peak performance while the MALI T604 is sustaining 25.5 GFlops. The full SoC offers a 32 GFlops peak performance with a 10 W TDP giving a 3.2 GFlops/W peak power/performance ratio.

During the first half of 2013, the Mont-Blanc partners involved into WP7 (BSC, ARM, BULL, Gnodal) worked in designing a new Mont-Blanc low power compute blade integrated by BULL. This blade and the full chassis architecture were disclosed by Alex Ramirez, Mont-Blanc project leader during the ISC 2013 conference in Leipzig (Germany) in June 2013.
Fitted into a standard BullX Carrier Blade, each Mont-Blanc compute blade will carry up to 15

sockets Samsung Exynos 5250 and an integrated GbE switch. Each blade will deliver a 485 GFlops peak performance. The Mont-Blanc prototype architecture will be composed by stack of chassis, each 7U chassis containing 9 blades and delivering around 4.3 TFlops peak performance.

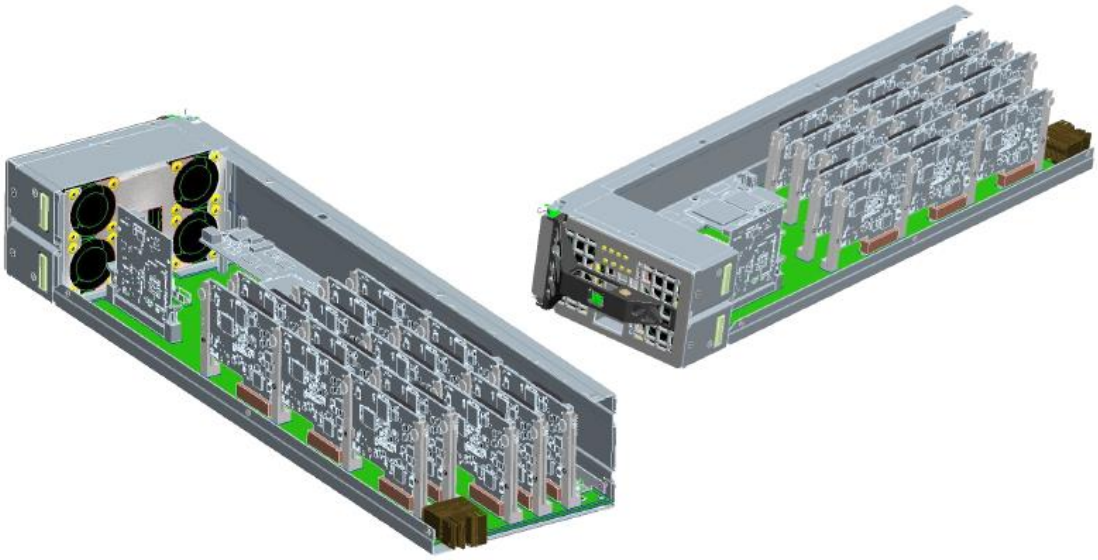

Figure 2 - Picture of a Mont-Blanc blade fitted into BULL standard format

From the software point of view, the Exynos 5250 come with an optimised OpenCL stack from ARM and BSC is working on a full optimised implementation of the Mercurium compiler in order to generate OmpSs/OpenCL codes.

3 Selection of the WP4 applications subset

After the first 2 years all the 11 applications have been ported and scaled out. Some of the applications are optimised using different platforms available: from low power single boards to full low-power clusters like Tibidabo, or large scale PRACE Tier0 systems like CURIE, SuperMUC or JUQUEEN.
In parallel, WP4 worked with WP3 to extract pertinent numerical kernels from some of the 11 HPC applications under study, optimise them and perform a taskification using OmpSs. Such kernels are planned to be integrated back into full applications during the last year of Mont–Blanc project.

The choice of the Mont-Blanc SoC, the Samsung Exynos 5250 with a dual-core ARM Cortex-A15 and a MALI T604 GPU has a strong, but expected impact into the choice of the applications to select for performing specific optimisation and port them on the Mont-Blanc prototype.
The ARM Cortex-A15 is capable of sustaining a 6.8 GFlops peak performance while the MALI T604 is sustaining 25.5 GFlops.

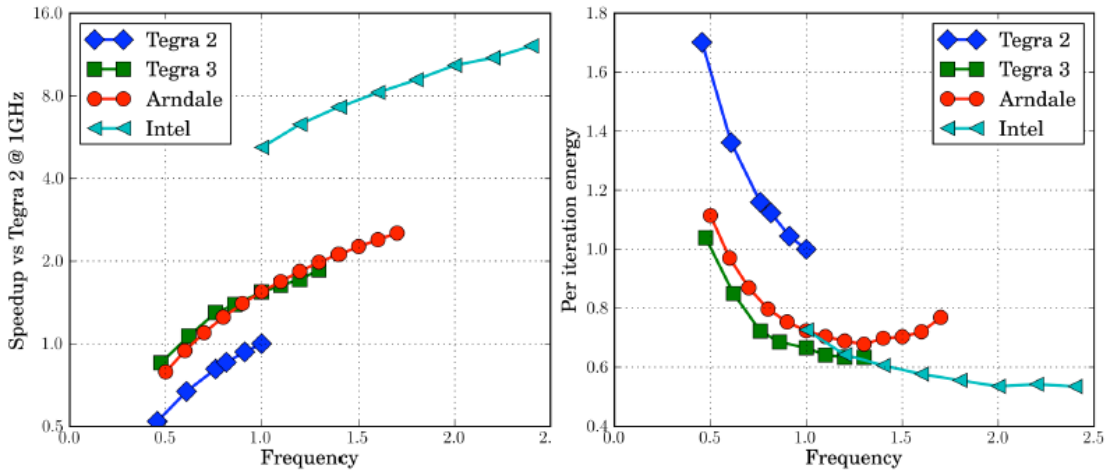

Figure 3 - Multicore performance and energy of Arndale boards (dual-core ARM Cortex- A15) against Tegra2 (dual-core ARM Cortex-A9), Tegra 3 (quad-core ARM Cortex-A9) or Intel x86 Core i7 (quad-core Sandybridge at 2.4 GHz)

As shown in Figure 3, even with some major improvements in term of power/performance of the ARM Cortex-A15 core with respect to the ARM Cortex-A9 (which was the reference processor for the Tibidabo cluster) or with respect to regular Intel core i7 x86 processors, it appears mandatory in order to exploit the full performance of the SoC to use the MALI T604 GPU or both the ARM Cortex-A15 and the MALI GPU together.
As the MALI T604 is programmable using an ARM optimised OpenCL 1.1 software stack, it become mandatory to rely on hybrid codes written in OpenCL or in programming models able to generate an OpenCL back end.

This became the main criteria of selecting a subset of applications that are to be ported, optimised and scaled-out during the third year of project into the final Mont-Blanc prototype, when available.
In order to select the sublist it has been asked to all the WP4 application owners to give feedback about the potential portability of their application on the final Mont-Blanc prototype.
Portability means in that case, either having an existing native OpenCL implementation, having in plan to achieve the development of a native OpenCL version before mid 2014 (end of the Mont-Blanc project) or having an existing CUDA, OpenACC version to be ported to OpenCL or OmpSs/OpenCL during the final period.

The following table illustrates the answers of the 11 code owners of WP4 :

Code	CUDA version	OpenACC version	OpenCL version	Planned OmpSs/OpenCL version ?
BigDFT	X		X	
BQCD	X			X
COSMO		X		X
EUTERPE				X
MP2C		X		
PEPC		X		X
ProFASI				X
Quantum Espresso	X			X
SMMP	X		X	
SPECFEM3D	X			X
YALES2				

Regarding the EUTERPE code for modelling fusion, the code owner gave the information that this software is not any more OpenSource, as a consequence it could not be possible to follow activities in the field of Mont- Blanc.

Regarding the YALES2 application for combustion, the code is not planned to be hybridized on GPUs during the period. With the developer it has been discussed the idea of porting a new kernel related to the modelling of dynamic chemistry models on OpenCL with the support of WP3. So, the activity of the YALES2 team could move from WP4 to WP3 during the last year of Mont-Blanc project.

4 Conclusion and next steps

During the first 2 years of the project, a total of 11 real large-scale applications, daily used in HPC national and European centers by wide scientific and industrial communities, have been ported to the low-power architectures made available by the Mont-Blanc partners.
End of 2012 the Mont-Blanc project announced the choice of the Samsung Exynos 5250 as reference SoC for the Mont-Blanc final prototype. As the main performance of this SoC will rely on the use of the MALI T604 GPU or both MALI and ARM Cortex-A15 processor, WP4 worked internally on selecting applications that could be suitable for this upcoming architecture, by running in native OpenCL mode or in OmpSs with an OpenCL backend.
Out of the 11 applications, during the lifetime of the project, up to 9 have or will have a version which will be suitable for Mont-Blanc architecture. EUTERPE is not an open source application anymore, thus it will be excluded from the Mont-Blanc applications. Porting of YALES2 application, will move from WP4 to WP3 during the last year of the project.

This choice will allow all the participants to keep an active role in WP4 and will strengthen the relation between WP3 and WP4, as during the last year of the project, WP3 will provide the

optimised kernels that will be re-inserted into WP4 applications for a final demonstration on the Mont-Blanc prototype, when available.

List of figures

Acronyms and Abbreviations

- DEISA Distributed European Infrastructure for Supercomputing Application
- GbE Gigabit Ethernet
- GPL General Public Licence
- GPU Graphics Processing Unit
- HKMG High-k Metal Gate
- HPC High Performance Computing
- I/O Input (read), Output (write) operations on memory or on disks/tapes
- MD Molecular Dynamics
- PRACE Partnership for Advanced Computing in Europe (http://www.prace-ri.eu)
- SoC System On Chip
- TDP Thermal Dissipation Power
- WP Work Package
- WP2 Work Package 2 ("Dissemination and Exploitation")
- WP3 Work Package 3 ("Optimized application kernels »)
- WP4 Work Package 4 ("Exascale applications")
- WP5 Work Package 5 ("System software")
- WP6 Work Package 6 ("Next-generation system architecture")
- WP7 Work Package 7 ("Prototype system architecture")
- WPL Work Package Leader

List of references

[1] Mont-Blanc D4.1 "Preliminary report of progress about the porting of the full-scale scientific applications"
[2] Mont-Blanc D4.2 "Final report of progress about the porting of the full-scale scientific applications"