

Seismic Processing Performance Analysis on Different Hardware Environment

*E. O. Tyutlyayeva*¹, *S. S. Konyukhov*¹, *I. O. Odintsov*¹, *A. A. Moskovsky*¹

© The Authors 2017. This paper is published with open access at SuperFri.org

In this research we have used computational-intensive software that implements 2D and 3D seismic migrations to study mini-application behavior for a set of the computational architectures. In addition to three architecture type comparative analyses, two CPU generation comparisons have been done.

The dynamic behavior of chosen mini-applications was studied using BSC performance analysis tools to identify their common features.

In summary, we observe the best performance of mini-applications on Intel Xeon E5-2698 CPU generation 4. Intel Xeon Phi 7250 peculiar architectural characteristics requires careful source code optimizations to help the compiler to effectively vectorize time-consuming loops and to improve the cache locality in order to achieve higher performance level. Elbrus-4S CPU is theoretically suitable for such kind of applications, but the currently observed performance is an order of magnitude smaller than on Xeon E5 family; we believe that the frequency and RAM bandwidth increase, as well as source code optimization work could improve its performance.

Keywords: Performance analysis, architecture comparison, seismic processing profiling, power-usage analysis, application behavior analysis.

Introduction

In this research, the suitable seismic processing mini-applications were selected with active collaborations with practitioners in seismic data analysis. These mini-applications can serve as a basis for detailed performance study of reverse time migration algorithms, which are actively used in reconstruction of under-surface Earth structure from the seismic sensor readings.

- Dynamic behavior of the chosen mini-apps is studied using the performance analysis tools to identify their common features;
- Analysis was performed for a set computational architectures, including both common architectures, such as x86, and non-standard one, VLIW; and
- Performed analysis demonstrates a scalability potential for the chosen mini-applications, and we expect more performance /speedup for these mini-applications if run on computational cluster in multi-threaded/multi-MPI processes way. That is planned for the future work.

The main emphasis was put on the computations, while mini-applications' I/O requirements, which play important role during data processing and affect total processing time, need to be investigated further.

1. Tested Applications

For performance analysis we have used the most typical seismic mini-applications, that implemented 2D and 3D seismic migrations, based on the algorithms used in practice. These applications have been chosen in cooperation with practicing researchers in this field. The source codes were provided by "GEOLAB" Company, [3].

¹ZAO RSC Technologies, Moscow, Russian Federation

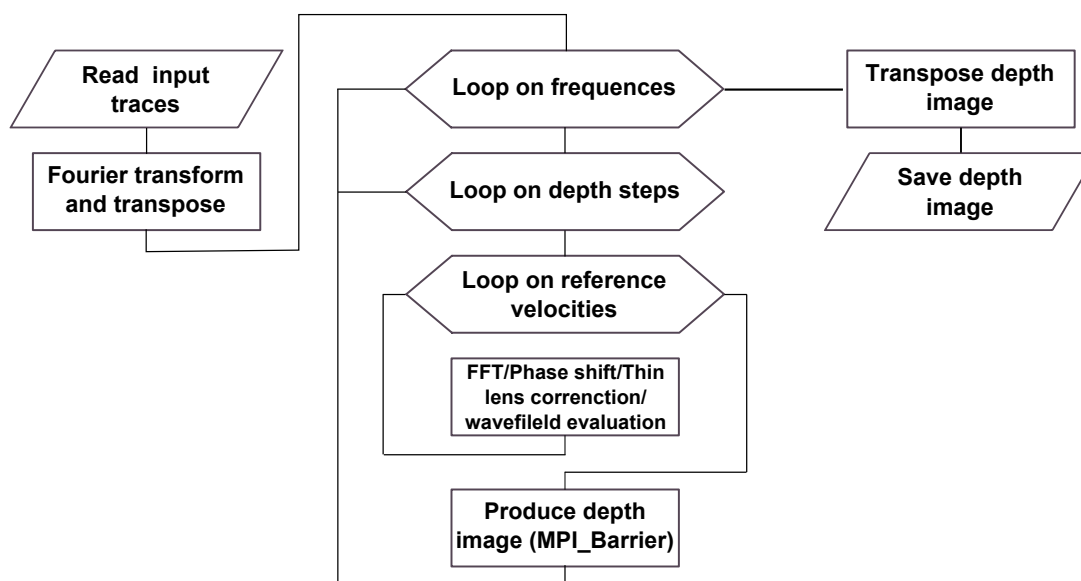


Figure 1. Seismic migration basic flowgraph

At the same time, the applications are characterized by acceptable level of computation complexity, which allows to use different techniques for testing different computational platforms. The basic flowgraph of the seismic migration is represented on Fig. 1

2D-seismic migration application (*Wemig*, [7]) uses reverse-time wavefield continuation in frequency/space domains and depth imaging. The MPI parallel programming model has been implemented with basic auto-vectorization for Intel architectures. Input data amount for test: 206 MB.

3D seismic migration application (*Cazmig*, [2]) implements the Cazdag migration algorithm, based on 3D data migration. In this method all computations are performed in the frequency domain where the source and the receiver positions are aligned with the phase shift by the rotation operation of Fourier coefficients. The hybrid parallel programming model has been used (MPI+OMP) with basic auto-vectorization for Intel architectures. Input data amount for test: 13.4 GB.

2. Simulation Stage

At the first stage, the algorithm structure analysis has been studied using BSC Performance Analysis Tools [6], with early efforts focused on simulation of multi-node cluster.

The execution trace-files have been generated using Extrae [4], a tool for post-mortem analysis. Then a simulation tool Dimemas [1] has been used for first approximation of geological processing software and hardware interaction.

During the simulation stage, the workload intensity has been evaluated, as well as scalability limits and DRAM impact on computation performance.

As an example, the MIPS (Million Instructions Per Second) distribution during the 2D seismic migration execution on the simulated 1-node and 4-node configurations of Intel Xeon E5-2697 v3, 64 GB RAM is shown on Fig. 2. Each horizontal line represents the timelined view of each MPI rank, and the color intensity reflects the workload (darker color means more intensive computation workload, then lighter).




Figure 2. Tracing of MIPS during 2D seismic migration Mini-App simulation

According to the simulation result, the workload evenly distributed between all used MPI ranks during execution time. The simulated results for four similar nodes also show balanced workload and good application scalability, while computation intensity is decreased.

Further simulation and analysis with different amounts of cores and DRAM shows good scalability for 2D and 3D seismic migration and results and strong correlation between amount of DRAM and workload. Additional practical testing using different amount of RAM has been conducted, based on the simulation results.

3. Testbeds

The real cluster nodes prototypes have been chosen, based on the simulation results and the main trends in geological computations. The basic specifications of studied testbeds are listed in the Table 1.

4. Efficiency

Experiments with different numbers of used computation cores have shown the efficiency of the 2D seismic migration mini-app depending on the hardware (see Fig. 3).

The application efficiency depending of the amount of RAM is presented on Fig. 4. The dotted vertical line divides physical cores efficiency (left side of the line graph) from the hyperthreading technology efficiency, [5] (right side of the line graph, where number of threads exceed the number of physical cores).

It is worth noting, that the doubling of memory capacity leads to the significant performance increase on the Broadwell testbed, especially in the hyperthreading range, while on the Haswell testbed productivity gains are not substantial for the tested mini-apps. It seems the Broadwell cores with 64GB RAM configuration were stalled due to memory demands; the Haswell 64GB results are more balanced.

The absolute values of the execution time confirm these observations for 2D and 3D seismic migration cases (see Table 2). The provided results are the best from the tested range (for 2D seismic migration we have tested a number of MPI from range 1 .. N_cores x Hyperthreading; for 3D migration, implemented using hybrid parallelization scheme, we have tested all suitable MPI + OMP configurations in this range).

Table 1. Testbeds Specifications

Codename	CPU	# Cores	Memory	GB/Core
Haswell_64GB	Intel Xeon E5-2697 v3	2x 14	8x DRAM Micron 8GB DDR4/2133MHz	2.28
Haswell_128GB	Intel Xeon E5-2697 v3	2x 14	8x DRAM Samsung 16GB DDR4/2133MHz	4.57
Broadwell_64GB	Intel Xeon E5-2698 v4	2x 20	8x DRAM Micron 8GB DDR4/2133MHz	1.6
Broadwell_128GB	Intel Xeon E5-2698 v4	2x 20	8x DRAM Samsung 16GB DDR4/2133MHz	3.2
KNC	Intel Xeon Phi 7120D	61	SDRAM Intel 16GB GDDR5/2750MHz	0.26
KNL	Intel Xeon Phi 7250	68	MCDRAM Intel 16GB + 6x DRAM Micron 32GB DDR4/2133MHz	2.8
Elbrus	Elbrus-4C	4x 4	12x DRAM Micron 4GB DDR3/1600MHz	3

Table 2. Impact of the Amount of DRAM on the Execution Times of 2D and 3D Seismic Migration Mini-Apps

Testbed	2D Seismic Migration	3D Seismic Migration
Haswell_64GB	56 sec (56 MPI)	57 min 35 sec (14 MPI, 14 OMP)
Haswell_128GB	50 sec (56 MPI)	55 min 39 sec (14 MPI, 14 OMP)
Broadwell_64GB	55 sec (80 MPI)	56 min 13 sec (4 MPI, 16 OMP)
Broadwell_128GB	36 sec (80 MPI)	43 min 2 sec (16 MPI, 16 OMP)

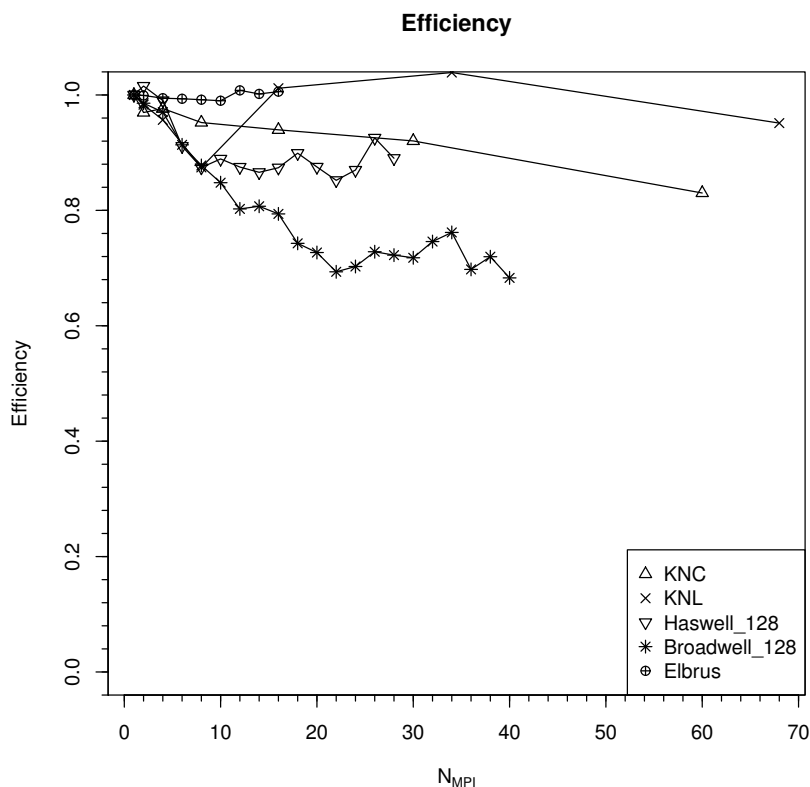


Figure 3. Efficiency of the 2D seismic migration depending on the hardware

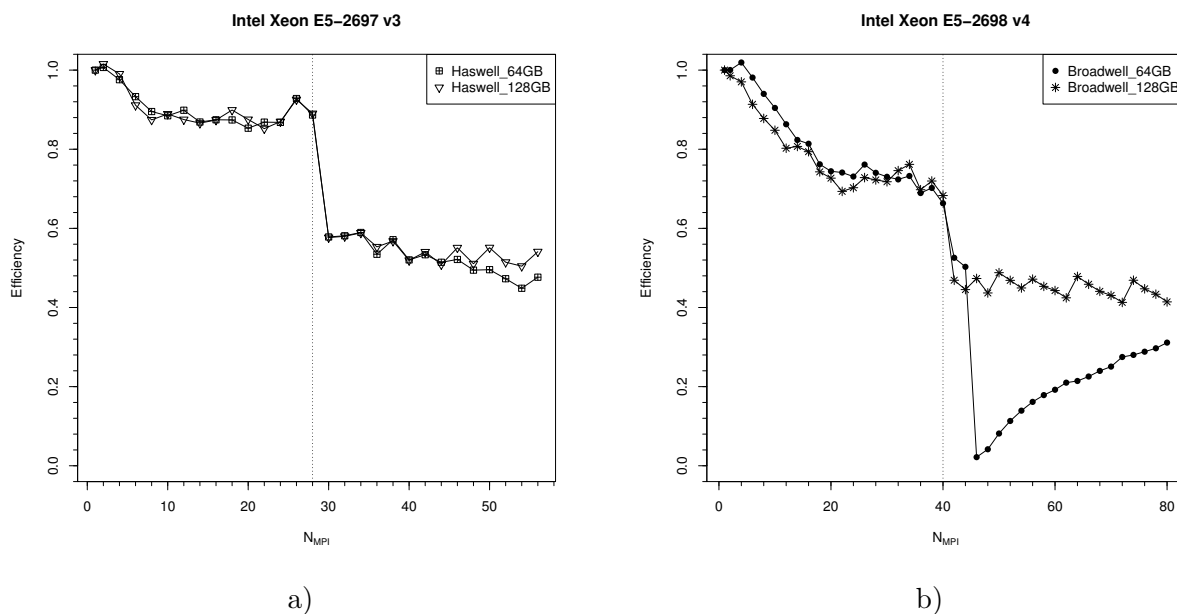


Figure 4. Efficiency of the 2D seismic migration depending on the amount of DRAM for a) Haswell, b) Broadwell testbeds.

However, the increase in number of MPI processes results in declining efficiency rate for all architectures despite enabled hyperthreading technology, that provides some performance increase in absolute values.

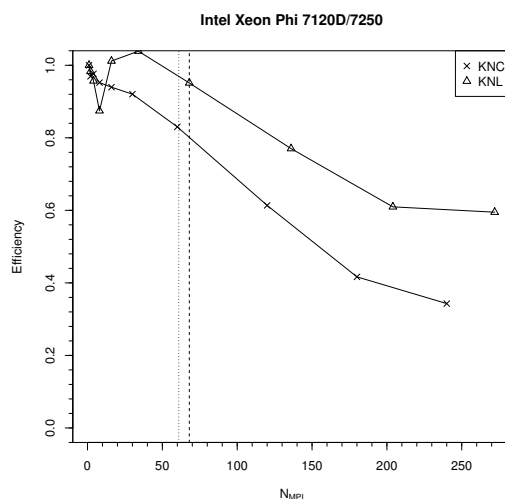


Figure 5. Efficiency of the 2D seismic migration on two generations of Intel Xeon Phi

The parallel computation efficiency for two generations of Intel Xeon PHI is shown on Fig. 5. Intel Xeon Phi architectures support hyperthreading technology for up to 4 virtual threads/core.

The resulting comparative curves show that the second generation architecture (KNL) has significant efficiency and scalability advantage over the first generation (KNC) for the tested mini-apps, especially in the hyperthreading range. The absolute values of execution times presented in Table 3 also support this observation. Numbers of MPI processes and MPI threads were carefully selected for each test run to achieve maximal performance.

Table 3. Execution Times of 2D Seismic Migration Mini-App on Two Generations of Intel Xeon PHI

Testbed	2D Seismic Migration
KNC	13 min (240 MPI)
KNL	3 min (272 MPI)

Finally, Elbrus-4S architecture demonstrates the almost perfect efficiency of up to 16 processes (i.e. up to one MPI process per computational core), because the amount of computations is high for this architecture. (Fig. 6.)

5. Tracing

The application tracing results are presented on Fig. 7 in the same way, as simulation traces, where each horizontal line represents one MPI rank. The server clients communication model is implemented, root computation rank read data and generates packages for clients to process. Client ranks are waiting and synchronizing.

The processors are highly loaded during the main computation stage, workload intensity balances for KNL and Haswell testbeds, while there is some performance swings on the Broadwell testbed due to high performance.

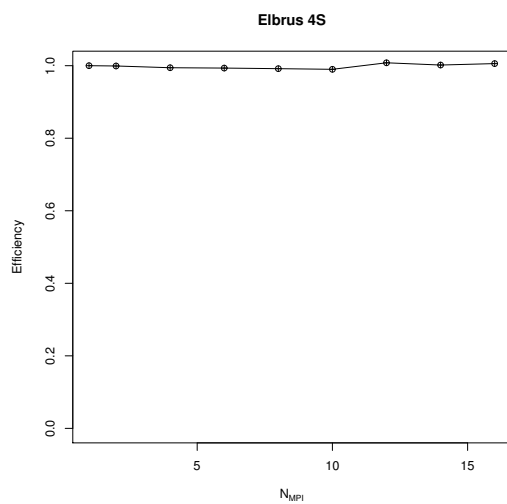
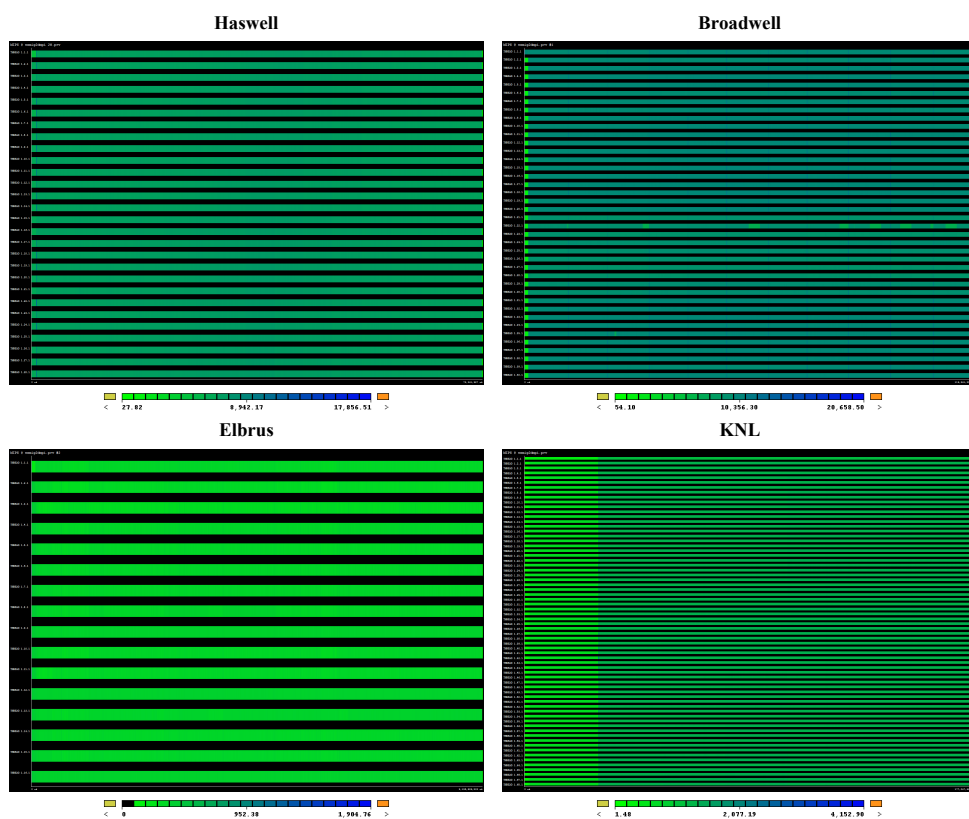


Figure 6. Efficiency of the 2D seismic migration on the Elbrus 4S



* According to The Elbrus architecture specifics the information about VLIW instructions per cycle were gathered instead of the standard IPC metric (Instructions per Cycle).

Figure 7. Efficiency of the 2D seismic migration depending on the hardware

6. Energy Consumption

The energy consumption of 3D Seismic Migration was studied using Intel Running Average Power Limit (RAPL, [8]) counters. RAPL provides a way to measure power consumption on processor packages and DRAM. The power consumption tracing is shown on Fig. 8, where a curved line represents total power consumption, measured by RAPL*, and a dotted line represents measured average idle power consumption.

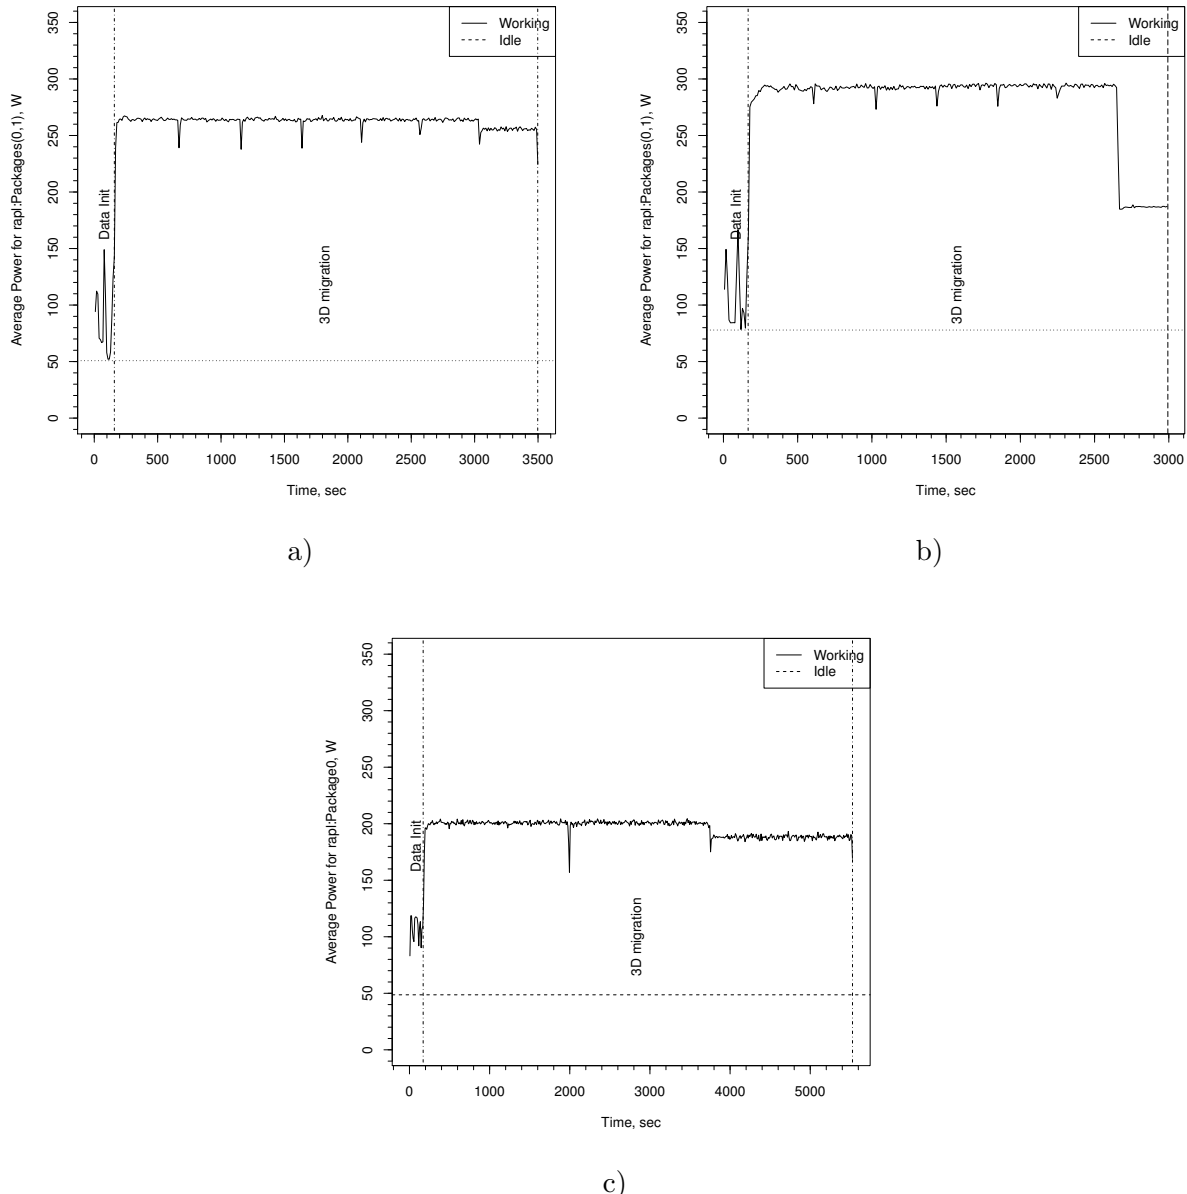


Figure 8. Energy consumption for a) Haswell_128GB, b) Broadwell_128GB c) KNL testbeds during 3D seismic migration

Although during 3D-seismic migration test power consumption rate for Broadwell was higher than for KNL (270 W vs. 188W), total power consumption for Broadwell was lower (0.2245 kW*h vs. 0.2993 kW*h) because of substantially lesser runtime (3080 sec vs. 5741 sec). The Haswell power consumption results show intermediate values (see Table 4).

Table 4. Energy Consumption during 3D Seismic Migration Execution

	KNL	Haswell_128GB		Broadwell_128GB	
		Package0	Package1	Package0	Package1
Processor, J	967255	418726	424767	350858	92335
DRAM, J	113129	22204	25260	80799	90557
Time, sec	5741.172		3510.175		3080.837
Total Energy, kWh	0,2993		0.2475		0.2245

It may be noted, that the energy consumption of Haswell sockets is similar (see Table 5), while at the Broadwell testbed the first socket (processor) consumed more energy then the second one more then three times. There can be correlation with workload disbalance detected on the tracing stage. It seems, that the Broadwell_128GB testbed still has room for code optimization to achieve maximum possible performance; larger amount of computation data and more sophisticated processing models could also be used successfully.

Table 5. Energy Consumption during 3D Seismic Migration Execution

Testbed	Data	Input	Processing
KNL	Time. sec	169.62	5571.552 sec
	DRAM. J	1051	112078
	CPU. J	14980.58	952274.42
Haswell_128	Time. sec	157.32	3352.855 sec
	DRAM. J	819.86	46644.14
	CPU. J	12189.47	831303.53
Broadwell_128	Time. sec	165.29	2913.873 sec
	DRAM. J	4626.6	166729.4
	CPU, J	10389.07	339303

Conclusions

In this research we have used computational-intensive software that implements 2D (Wemig) and 3D (Cazmig) seismic migrations to study the application behavior for a set of the computational architectures. In addition to three architecture type comparative analyses, two CPU generation comparisons have been done.

For Haswell/Broadwell testbeds with similar architecture there has been a substantial (about 2x times) performance growth between generations; for the KNC/KNL testbeds the performance increase amounted up to 4x times. Moreover, there is portability issues with KNC architecture

that are eliminated in KNL software stack. While the I/O overhead costs are non-essential (0.0% of overall runtime) for most studied architectures, for KNL it takes 0.73% of the runtime. KNC runtime results have worse scalability than the KNL due to lesser amount of RAM per core.

It is worth noting that the doubling of RAM memory capacity leads to the significant performance increase on the Broadwell testbed, while on the Haswell testbed productivity gains are not substantial. So the memory amount for seismic applications should be appropriate to avoid the CPU stalls. The Elbrus-4S CPUs show the best scalability while overall absolute values were lower than values for the Intel Xeon according to the theoretical performance value rates.

Average power consumption rate is the lowest for KNL and the largest for Broadwell; but total power consumption for 3D seismic migration run shows the best rates for Broadwell testbed.

In summary, it makes sense for seismic applications to use the Intel Xeon E5-2698 CPU (Broadwell) generation instead of E5-2697 (Haswell) only with large amount of RAM available; the Intel Xeon Phi (KNC/KNL) particular architectural characteristics requires careful source code optimizations to help the compiler to effectively vectorize time-consuming loops and to improve the cache locality for achieving higher performance level; The Elbrus-4S CPU is theoretically suitable for such kind of applications, but it requires the frequency and RAM bandwidth increasing, as well as sophisticated source code optimization work for achieving the best instruction-level parallelism.

Acknowledgments

This research was supported by the Russian Federal Science and Technology Program grant 14.607.21.0165 "Efficient co-design of massively parallel computer for multispectral night-time remote sensing". We thank our colleagues from GEOLAB, particularly Evgeny Kurin for the collaboration.

This paper is distributed under the terms of the Creative Commons Attribution-Non Commercial 3.0 License which permits non-commercial use, reproduction and distribution of the work without further permission provided the original work is properly cited.

References

1. Dimemas: predict parallel performance using a single cpu machine ; <https://tools.bsc.es/dimemas>, accessed: 04.07.2017
2. Gazdag J. Wave Equation migration with equation migration with Phase Shift method / J. Gazdag // Geophysics - 1978 - V. 43 P. 1342–1351
3. The GEOLAB-IT Company; <http://sk.ru/net/1120173/>, accessed: 04.07.2017
4. Home: Extrae, Trace-generation package; <https://tools.bsc.es/extrae>, accessed: 04.07.2017
5. Intel Hyper-Threading Technology, <http://www.intel.ru/content/www/ru/ru/architecture-and-technology/hyper-threading/hyper-threading-technology.html>, (accessed: 29.11.2016)
6. Performance Analysis Tools: Details and Intelligence; <https://tools.bsc.es/>, (accessed: 04.07.2017)

7. Popovici A.M. Prestack migration by split-step DSR / A.M. Popovici // Geophysics - 1996 - V. 61 - P. 1412-1416;
8. Intel 64 and IA-32 Architectures Software Developers Manual, Volume 3B: System Programming Guide, Part 2, <http://www.intel.com/content/dam/www/public/us/en/documents/manuals/64-ia-32-architectures-software-developer-vol-3b-part-2-manual.pdf>, accessed: ~04.07.2017