



Evaluating Performance Portability of Accelerator Programming Models using SPEC ACCEL 1.2 Benchmarks

Swen Boehm^(✉), Swaroop Pophale^(✉), Verónica G. Vergara Larrea,
and Oscar Hernandez

Oak Ridge National Laboratory, Oak Ridge, TN 37831, USA
{boehms,pophale,vergaravg,oscar}@ornl.gov

Abstract. As heterogeneous architectures are becoming mainstream for HPC systems, application programmers are looking for programming model implementations that offer both performance and portability across platforms. Two directive-based programming models for accelerator programming that aim at doing this are OpenMP 4/4.5 and OpenACC. Many users want to know the difference between these two programming models, the state of their implementations, how to use them, and evaluate how suitable they are for their applications.

The Standard Performance Evaluation Corporation (SPEC) ACCEL benchmarks, developed by the SPEC High Performance Group (HPG), recently released SPEC ACCEL 1.2 benchmark suite to help the evaluation of OpenCL, OpenMP 4.5 and OpenACC on different platforms. In this paper we present our preliminary results that evaluates OpenMP 4.5 and OpenACC on a variety of accelerator-based systems: POWER9 with NVIDIA V100 GPUs (Summit), Intel Xeon Phi 7230 (Percival), and AMD Bulldozer Opteron with NVIDIA K20x (Titan). Comparing these benchmarks on different systems gives us insight into the support for OpenMP and OpenACC and their execution times provide insights about their quality of implementations provided by different vendors. We also compare best of OpenMP and OpenACC to see if a particular programming model favors a particular type of benchmark kernel.

1 Introduction

Benchmarks have been the backbone of performance modeling in HPC since the invention of parallel processing. They highlight key metrics that are important to a specific audience. Focusing on best practices, reproducible results and

This manuscript has been co-authored by UT-Battelle, LLC, under contract DE-AC05-00OR22725 with the US Department of Energy (DOE). The US government retains and the publisher, by accepting the article for publication, acknowledges that the US government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this manuscript, or allow others to do so, for US government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<http://energy.gov/downloads/doe-public-access-plan>).

© Springer Nature Switzerland AG 2018

R. Yokota et al. (Eds.): ISC 2018 Workshops, LNCS 11203, pp. 711–723, 2018.

https://doi.org/10.1007/978-3-030-02465-9_51

behavior are the main attributes that help in wide acceptance of a benchmark/benchmark suite. Though well maintained scientific applications would be excellent candidates, applications are usually not easily available, require a lot of domain specific information (e.g. input parameters), or are not ported to a variety of platforms. The SPEC [5] benchmark suites are a widely accepted set of benchmarks. They are constantly improved and optimized by vendors and the community. SPEC ACCEL benchmark suite provides a comparative measure of parallel compute performance between accelerator platforms. The SPEC ACCEL benchmarks are written using the offloading model in OpenCL, OpenMP and OpenACC to program accelerators. As per guidelines set by SPEC, the use of the benchmarks must conform to the set of rules [4] to ensure comparability and reproducibility of results.

One of the benefit on using benchmarks is that they help us understand how a programming model can be used and optimized to achieve good performance on a given code, compiler, and platform combination. The SPEC ACCEL benchmark has a set of codes that have been ported to OpenCL, OpenMP 4.5 (offload model) and OpenACC. The members of SPEC/HPG have discussed and agreed on the best practices to write the SPEC ACCEL benchmarks with performance portability in mind.

We find SPEC ACCEL useful to provide answers on the quality of implementations of OpenMP and OpenACC on a variety of platforms available to us. It can also be used to understand how programming models work and how they are used, which is important, since there is no single programming model that is best suited for all application domains.

The benchmarks are helpful to compare the quality of implementations across compiler implementations on a given system and across systems. Different vendors have support for one or the other programming model. For example, PGI has OpenACC implementations for different platforms but no OpenMP support. So we look at the support for an implementation (OpenMP or OpenACC) for the SPEC ACCEL 1.2 Benchmarks to see which implementations have a more complete support for the different kernels exercised in the benchmarks. In this paper, we focus on evaluating the OpenMP 4.5 (offload) implementations on IBM XL, clang/LLVM, GCC, and Cray CCE, and OpenACC support in PGI and GCC compilers. We also evaluate production closed source implementations vs open source compiler implementations, and see how they match up in terms of benchmark execution times.

Since compiler support for OpenMP 4.5 is work-in-progress, we use SPEC ACCEL benchmarks (1.2) as a reference to measure where we are in terms OpenMP and OpenACC functionality and performance support. We use the SPEC/ACCEL benchmarks to report the execution times. None of these results are *reportable runs* because the majority of compilers do not pass all the benchmarks. SPEC/HPG also documents official results and provides a wide range of benchmark results for a variety of systems and compilers. Results can be found at <http://www.spec.org/accel>.

2 Motivation

The OpenMP 4.0 specification [10] released in July 2013 introduced the OpenMP Accelerator Model. In the specification, the OpenMP API was extended to support accelerator and SIMD programming, allowing the user to specify regions of code that can be offloaded to one or more target devices. More recently, with the release of the OpenMP 4.5 specification [11], support for the accelerator model was further improved. Several major changes in OpenMP 4.5 affect the accelerator model including: new default data-mapping attributes, unstructured data mapping support and asynchronous execution, as well as runtime routines for memory management and extended attributes for SIMD loops, among others.

One of the goals of this paper is to evaluate the status of OpenMP 4.5 support on multiple architecture including NVIDIA GPUs and Intel Xeon Phi. We want to understand the maturity of OpenMP 4.5 offload support and how it compares to OpenACC across architectures.

Given the prescriptive nature of OpenMP, the SPEC/HPG organization designed the SPEC/ACCEL 1.2 OpenMP 4.5 benchmark implementation to rely on default behaviors to give as much freedom to the compiler and runtimes to optimize for a given architecture, without making optimization assumptions that favors one architecture over another.

For example, when a *teams* construct is executed, a league of threads is created, where the total number of teams is implementation defined but must be less than or equal to the number of teams specified by the *num_teams* clause. If the user does not specify the *num_teams* clause, then the number of teams is left completely to the implementation.

There are many concepts in the OpenMP and OpenACC specification that are implementation defined. All the implementation defined behaviors may have a significant effect on the performance of an implementation. Although deep analysis of the different OpenMP and OpenACC constructs is outside the scope of this paper, we provide the foundation by analyzing the execution times of the SPEC ACCEL 1.2 benchmarks.

3 The SPEC ACCEL Benchmark Suite

SPEC ACCEL 1.2 is a collection of the following benchmarks:

1. Stencil

The stencil code represents an iterative Jacobi solver. The heat equation is represented by a 3-D structured grid.

2. Lattice Boltzmann Method (LBM)

This program implements the LBM to simulate behavior of incompressible fluids in 3D. This simulation represents the most critical and computationally important part of calculations used in material science.

3. MRI-Q

This benchmark attempts to reconstruct a MRI. This is a compute bound problem where a large set of input is processed. The result is the Q Matrix

where each element of the Q matrix is computed by the summation of contributions from all trajectory sample points.

4. Molecular Dynamics (MD)
This benchmark performs molecular dynamics simulations of stellar objects. The simulations are applicable to all dense nuclear matter.
5. PALM
The benchmark simulates large-eddy simulation (LES) used for atmospheric and oceanic flows.
6. Clover Leaf
This benchmark uses domain decomposition to solve the compressible Euler Equations.
7. Conjugate Gradient (CG)
The CG benchmark performs irregular long distance communication to solve an unstructured sparse linear system.
8. Seismic Wave Modeling
The benchmark solves two and three dimensional isotropic or anisotropic elastic, viscoelastic or poroelastic wave equation using a finite-difference method.
9. Scalar Penta-diagonal solver(SP)
The benchmark solves scalar, pentadiagonal equations to simulate synthetic CFD problem solver. Both Fortran and C versions of this benchmark are included in the suite.
10. Mini Ghost
This benchmark supports Bulk Synchronous Parallel (BSP) model for finite difference stencil computation.
11. ILBDC
The benchmark kernel is a component of a flow solver that simulates the collision-propagation routine of an advanced 3-D lattice Boltzmann flow solver.
12. Swim
The benchmark is a weather prediction benchmark that use finite-difference approximation.
13. Block Tridiagonal (BT) Solver
The benchmark solves a 3D discretization of Navier-Stokes equation.

It is important to note that we use these benchmarks under the Academic license and as such the results presented here are for scientific curiosity only and may not be included as official results for the different architectures we are exploring.

4 SPEC ACCEL 1.2 Results

In this section we present our results. First we describe the systems where we run, then we report our findings on the different systems from the point of view of functionality. Then, we present the timings we found for OpenMP 4.5 offload and OpenACC. At the end we compare production closed-source implementations of OpenMP 4.5 and OpenACC versus open-source implementations.

4.1 Experimental Systems

For this study, we present results obtained on the OLCF’s flagship systems:

Summit [6] is the OLCF’s next generation high performance supercomputer. Summit compute nodes are IBM Power9 AC922 servers with two 22-core 3.45 GHz (turbo) IBM POWER9 processors. Each core supports 4-way SMT. Summit compute nodes have 512 GB of DDR4 memory and 96 GB of HBM2 memory, as well as 1.6 TB non-volatile memory that can be used as a burst buffer. Each compute node has 6 NVIDIA Tesla V100 Volta GPUs connected via NVLINK2 which provides up to 25 GB/s unidirectional and 50 GB/s bidirectional bandwidth between GPUs and from the CPU to the GPU. Summit compute nodes are interconnected via Mellanox EDR InfiniBand in a non-blocking fat-tree topology.

Titan [8] is the OLCF’s flagship supercomputer. Titan is a Cray XK7 system with 18,688 compute nodes each with a 16-core 2.2 GHz AMD Opteron 6274 Interlagos CPU, 32 GB of RAM, and one NVIDIA Kepler K20X GPU. Titan nodes are connected via Cray’s high speed Gemini interconnect.

Percival [3] is one of the supporting systems available at the OLCF and it was deployed to assist with performance portability efforts. Percival is a 168-node Cray XC40 supercomputer. Each Percival node is equipped with one 64-core 1.30GHz Intel Xeon Phi 7230 (KNL) processor and 110 GB of RAM. Percival nodes are connected via Cray’s Aries proprietary interconnect in a Dragonfly topology.

We execute reportable runs (three iterations of each benchmark, with the runtime averaged over those three runs). For more details on running the SpecACCEL benchmarks see [4].

4.2 Performance

In this section we present the timing information of the benchmarks for OpenMP 4.5 and OpenACC using different compilers and platforms.

4.3 Correctness and Functionality

Table 1 shows successes and failures (Verification, Runtime, Compile time errors) of SPEC ACCEL 1.2 benchmarks on Summit, Titan and Percival. We can see that in terms of functional implementations the PGI OpenACC supports majority of the benchmarks on all of the systems. Intel compiler on Xeon Phi has the most support for OpenMP 4.5. On Summit, the XL compiler passes the most benchmarks and the Clang/LLVM compiler passes the majority of the C benchmarks. The GCC compiler is still in the early stages of supporting OpenMP 4.5 offload on NVIDIA GPUs, hence it only compiles and (correctly) executes three of the fifteen benchmarks.

Experiments on Summit are performed using CUDA 9.2.64, PGI 18.3 and XL V16.1.0 (Beta 4), GCC 7.2 and CLANG (ykt branch) where PGI and GCC provide support for OpenACC and IBM XL and CLANG provide support for

Table 1. Successes and failures of running the SPEC ACCEL 1.2 benchmarks on different architectures with OpenMP 4.5 and OpenACC. The compiler versions used are: On Summit: PGI 18.3, XL V16.1.0, Clang/LLVM (ykt branch), GCC 7.2 (gomp branch), on Titan Cray CCE 8.7.0, PGI 18.4 and Percival Intel 18.0.0.128 and PGI 18.5

	Summit (NV100 GPU)					Titan (K20X GPU)		Percival (Xeon Phi)	
	XL	PGI	GCC		Clang	PGI	CCE	PGI	Intel
	OMP	ACC	ACC	OMP	OMP	ACC	OMP	ACC	OMP
Stencil	✓	✓	✓	✓	✓	✓	✓	✓	✓
LBM	✓	✓	✓	✓	✓	✓	✓	✓	✓
MRI-Q	✓	✓	✓	x ^{RE}	✓	✓	✓	✓	✓
MD	x ^{VE}	✓	✓	x ^{RE}		✓	x ^{RE}	✓	x ^{CE}
PALM	x ^{RE}	✓	✓	x ^{RE}		x ^{CE}	x ^{CE}	✓	✓
EP	✓	✓	✓	x ^{VE}	✓	✓	✓	✓	✓
CLVRLEAF	✓	✓	✓	x ^{RE}		✓	x ^{VE}	✓	✓
CG	✓	✓	✓	x ^{VE}	x ^{RE}	✓	✓	✓	✓
SEISMIC	✓	✓	✓	x ^{RE}		✓	x ^{RE}	✓	✓
SP	F	x ^{RE}	✓	✓	x ^{RE}		✓	x ^{RE}	✓
	C	x ^{VE}	✓	✓	x ^{RE}	✓	✓	✓	✓
MiniGhost	x ^{RE}	✓	✓	x ^{RE}		x ^{CE}	x ^{RE}	x ^{CE}	x ^{CE}
LBDC	✓	✓	✓	✓		✓	x ^{RE}	✓	✓
Swim	x ^{NR}	✓	✓	x ^{RE}		✓	x ^{RE}	✓	✓
BT	x ^{NR}	✓	✓	x ^{RE}	✓	✓	✓	✓	✓
Passed	8	15	15*	3	6	13	7	14	13

*GCC/OpenACC only offloads 4 out of the 15 benchmarks, the remaining 11 benchmarks utilize the CPU.

VE: Verification error

RE: Runtime error

CE: compile error

NR: benchmark excluded from run

OpenMP. Figure 1 compares the timings for SPEC ACCEL 1.2 OpenACC benchmarks when using the PGI 18.3 and GCC 7.2 (GOMP branch) compilers, both using CUDA 9.2.64. We can see that of the fifteen benchmarks PGI has better execution times for all except two (LBM and LBDC). From the timings we observe that PGI has a more optimized implementation for OpenACC. Another reason for the differences in the timing is that GCC only offloads STENCIL, LBM, MRI-Q, and LBDC to the device as GCC has only functional support for ACC kernels (without offloading) used in the remaining benchmarks. Both PGI and GCC execute all the benchmarks successfully showing that their OpenACC functionality support is complete.

Figure 2 shows the timings for SPEC ACCEL 1.2 OpenMP 4.5 (offload) benchmarks when using the XL V16.1.0 (Beta 4), the CLANG (ykt branch) and GCC 7.2 (GOMP branch). The CLANG specific version used is an IBM version loosely based on the latest CLANG version available in trunk. Clang




Fig. 1. Execution times of SPEC ACCEL 1.2 OpenACC using the PGI 18.3 and GCC 7.2 (GOMP branch) with CUDA 9.2.64 on Summit

compiles most of the C benchmarks (Stencil, LBM, MRI-Q, EP, CG, SP(C), BT). It produces a runtime error for CG. CLVRLEAF and Minighost are mixed C and Fortran codes, and are not supported using CLANG. The XL compiler always shows better timing results than CLANG on the sub-set of benchmarks that both implementations compile and execute correctly. However, XL does not have execution results for BT (excluded from run) and SP-C (verification error). GCC has minimal support for OpenMP 4.5 offload, as only three of the fifteen benchmarks compile and execute correctly. However we observe that for LBM it outperforms the CLANG compiler.

Figure 3 compares the execution times of SPEC ACCEL 1.2 with production implementations from PGI for OpenACC and XL for OpenMP on Summit. Both compilers use the CUDA 9.2.64 tool chains. From Fig. 3 we see that OpenACC (PGI) has better performance for four of the benchmarks (MRI-Q, EP, Clover-Leaf, CG) when compared to the OpenMP 4.5 (XL). Even when XL compiler only compiles and executes seven of the fifteen OpenACC benchmarks, three of the OpenMP 4.5 benchmarks (Stencil, LBM, Seismic) show better performance than OpenACC. This indicates that OpenMP 4.5, though newer, is promising and can have better performance compared to OpenACC. Maturity of the compiler is a relevant factor, as OpenACC implementations have been available for a longer time.




Fig. 2. Execution times of SPEC ACCEL 1.2 OpenMP 4.5 (offload) using XL V16.1.0 (Beta 4), Clang (ykt branch) and GCC 7.2 (GOMP branch) with CUDA 9.2.64 on Summit

Figure 4 shows the SPEC ACCEL 1.2 results on Titan for OpenACC and OpenMP 4.5 (offload) using the PGI 18.4 and Cray CCE 8.7.0 compilers. OpenACC has better performance on four of the benchmarks (STENCIL, MRI-Q, SP_C and BT) and OpenMP 4.5 offload has better performance on three of the benchmarks (LBM, EP, and CG). PGI fails to compile two of the benchmarks and Cray fails to pass (because of compilation, runtime and verification errors) eight of the benchmarks. This is evidence that, on Titan, the OpenACC implementation PGI provides is more mature than the OpenMP 4.5 offload implementation Cray provides when compiling the SPEC 1.2 benchmarks.

Figure 5 shows the SPEC ACCEL 1.2 benchmarks timings on Percival. On this system, OpenACC has better performance than OpenMP 4.5 offload for four of the benchmarks (Stencil, LBM, MRI-Q, MD) when using the PGI 18.5 compiler for OpenACC and the Intel 18.0.0.128 compiler for OpenMP 4.5 (offload). OpenMP 4.5 has better performance for nine of the benchmarks (EP, CLVRLEAF, CG, Seismic, SP_F, SP_C, LBDC, Swim, BT). The PGI compiler fails to produce results for two of the OpenACC benchmarks (PALM and Minighost) and Intel fails to pass one of the OpenMP benchmarks (Minighost) due to compilation errors.




Fig. 3. Execution times of the SPEC ACCEL 1.2 OpenACC and OpenMP 4.5 (offload) benchmarks using PGI 18.3 and XL V16.1.0 (Beta 4) with CUDA 9.2.64 on Summit




Fig. 4. Execution times of the SPEC ACCEL 1.2 OpenACC and OpenMP 4.5 (offload) benchmarks using the Cray 8.7.0 and PGI 18.4 compilers on Titan




Fig. 5. Execution times of the SPEC ACCEL 1.2 OpenACC and OpenMP 4.5 (offload) benchmarks using Intel 18.0.0.128 and PGI 18.5 compilers on Percival.




Fig. 6. OpenMP’s Performance Improvement over OpenACC

4.4 OpenMP and OpenACC Performance Comparison

Since the benchmarks target the accelerator offload aspect for both OpenMP 4.5 and OpenACC, it is fair to compare the execution times as the algorithms are very similar. The performance improvement or relative speed up, shown in Fig. 6, is calculated by first choosing, for each platform, the best execution times for OpenACC and OpenMP and then dividing the OpenACC timing by the OpenMP for individual benchmarks.

For example, for Summit (Power 9 + NV100) we compare PGI's OpenACC with XL's OpenMP. For Percival we compare PGI's OpenACC and Intel's OpenMP implementations. While for Titan we use PGI's OpenACC and Cray's OpenMP implementations. From Fig. 6 it is clear that on most platforms OpenACC performs better than OpenMP 4.5 (offload). One of the reasons for it is that the OpenACC specification has been formalized for a longer time (OpenACC 1.0 was ratified in 2011), giving vendors time to optimize their implementation.

In contrast OpenMP first introduced accelerator support with OpenMP 4.0 [10] in 2013. That said, the figure also shows that for certain benchmarks like LBM OpenMP does better consistently across most platforms.

5 Related Work

EPCC has developed another suite of microbenchmarks that has both OpenMP and OpenACC variants. The OpenMP [2] versions focus on measuring the synchronization, loop scheduling and array operations overheads in the OpenMP runtime library. The benchmarks support OpenMP 3.1 API [12] and have not been modified to use `target` devices (introduced in OpenMP 4.0 and extended in OpenMP 4.5 Specification [11]). The OpenACC EPCC [1] microbenchmarks mimic kernels commonly seen in scientific codes. They focuses on low-level operations and are designed to test raw performance of compilers and hardware.

Currently there are no publicly accessible standardized OpenMP 4.5 benchmarks but an effort to build a Validation and Verification suite is underway [13] as part of a more elaborate Exascale Computing Project (ECP) - SOL-LVE. The OpenACC Standard [9] first introduced in 2011, focuses solely on the Application Program Interface (API) for compiler directives that offload compute kernels from a host CPU to an attached accelerator. Extensive work has been done in the field of benchmarks and validation and verification suites for OpenACC [7, 14].

6 Conclusion

This paper details our findings when trying to compile and execute SPEC ACCEL 1.2 benchmarks on our production systems: Titan, Summit, and Percival.

In the findings from these experiments we can see that the OpenACC implementations are more functionally mature than the OpenMP 4.5 implementations. However, performance of OpenACC and OpenMP 4.5 (offload) varies depending on the compiler and platform. This is an indication that OpenMP 4.5 is becoming more mature and it can be at par with OpenACC implementations. Since device offloading in OpenMP is a relatively new concept it comes as no surprise that there are more support for OpenACC than OpenMP. This is clear from the results across all systems, particularly Summit where we see fifteen OpenACC benchmarks in comparison to eight OpenMP benchmarks that have compiled and executed correctly.

For Summit we see that open source implementations of both OpenACC and OpenMP 4.5 (offload) are starting to come closer to production implementations for the case of GCC/OpenACC and CLANG/OpenMP 4.5. However, the GCC/OpenMP 4.5 (offload) is still in its early stages of supporting OpenMP 4.5 (offload) functionality and performance on GPUs.

Cray (Titan) also does not support OpenMP offloading as well as it does OpenACC, with thirteen OpenACC and only seven OpenMP benchmarks that have compiled and executed correctly.

For Intel we see that the performance gap between OpenMP and OpenACC narrows a little, with OpenACC still getting better performance for most benchmarks. But the overall high execution times with Intel and PGI compilers on Percival, when compared to other architectures, hints at the platform not being optimized to support accelerator models like OpenMP and OpenACC.

Acknowledgments. This material is based upon work supported by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under contract number DE-AC05-00OR22725. This research used resources of the Oak Ridge Leadership Computing Facility at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

References

1. EPCC OpenACC Microbenchmarks. <https://www.epcc.ed.ac.uk/research/computing/performance-characterisation-and-benchmarking/epcc-openacc-benchmark-suite>
2. EPCC OpenMP Microbenchmarks. <https://www.epcc.ed.ac.uk/research/computing/performance-characterisation-and-benchmarking/epcc-openmp-micro-benchmark-suite>
3. Percival quickstart guide. <https://www.olcf.ornl.gov/percival-quickstart-guide/>
4. SPEC ACCEL: Run and Reporting Rules. <https://www.spec.org/accel/docs/runrules.html>
5. Standard Performance Evaluation Corporation. <https://www.spec.org/>
6. Summit: Scale new heights. discover new solutions. <https://www.olcf.ornl.gov/summit/>

7. Friedline, K., Chandrasekaran, S., Lopez, M.G., Hernandez, O.: OpenACC 2.5 validation testsuite targeting multiple architectures. In: Kunkel, J.M., Yokota, R., Tauber, M., Shalf, J. (eds.) ISC High Performance 2017. LNCS, vol. 10524, pp. 557–575. Springer, Cham (2017). https://doi.org/10.1007/978-3-319-67630-2_39
8. Oak Ridge National Lab: Titan supercomputer. <https://www.olcf.ornl.gov/titan/>
9. OpenACC: OpenACC, Directives for Accelerators. <http://www.openacc.org/>
10. OpenMP: Openmp 4.0 specification. <http://www.openmp.org/wp-content/uploads/openmp-4.0.pdf>
11. OpenMP: Openmp 4.5 specification. <http://www.openmp.org/wp-content/uploads/openmp-4.5.pdf>
12. OpenMP Validation and Verification Suite: Openmp 3.1 Specification. <https://github.com/sunitachandra/omp-validation>
13. Pophale, S., Diaz, J.M., Hernandez, O., Bernholdt, D., Chandrasekaran, S.: OpenMP 4.5 Validation and Verification Suite for Device Offload. <https://crpl.cis.udel.edu/ompvvsolve/>
14. Wang, C., Xu, R., Chandrasekaran, S., Chapman, B., Hernandez, O.: A validation testsuite for OpenACC 1.0. In: 2014 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW), pp. 1407–1416. IEEE (2014)