PERFORMANCE AND ENERGY-EFFICIENCY ASPECTS OF CLUSTERS OF SINGLE BOARD COMPUTERS

Christian Baun
Faculty of Computer Science and Engineering, Frankfurt University of Applied Sciences,
Nibelungenplatz 1, 60318 Frankfurt am Main, Germany

ABSTRACT

When a high performance cluster is demanded and the cost for purchase and operation of servers, workstations or personal computers as nodes is a challenge, single board computers may be an option to build inexpensive cluster systems. This paper describes the construction of such clusters and analyzes their performance and energy-efficiency with the High Performance Linpack (HPL) benchmark.

KEYWORDS

Single Board Computers, Performance evaluation, High Performance Linpack

1. INTRODUCTION

For research projects, dealing with parallel applications, options for implementing clusters of inexpensive nodes are useful. Compared with commodity hardware servers, such clusters require lesser purchase costs and operating costs.

This paper is organized as follows. In Section 2, options for providing cluster resources to research projects with limited financial resources are discussed. Section 3 contains a discussion of related work. Section 4 presents a list of components of the clusters, which were constructed for this work. In Section 5, the performance and speedup of the entire cluster of single board computers is analyzed, by using the High Performance Linpack (HPL). Section 6 contains a calculation of the energy costs and an analysis of the energy-efficiency of the clusters. Finally, Section 7 presents conclusions and directions for future work.

2. SINGLE BOARD COMPUTERS - AN OPTION FOR IMPLEMENTING CLUSTERS

Depending of the available funding resources, the purchase cost for physical servers, workstations or personal computers can be challenging for scientific projects. Decommissioned hardware can be bought for little money, but it requires much space and the maintenance is labour intensive. Another important characteristic are costs, which arise from running physical computer resources. These include electricity cost.

Building clusters of single board computers with ARM architectures like the Raspberry Pi, BeagleBone Black, PandaBoard, BananaPi or ODROID is a further option for implementing clusters. Table 1 contains the single board computers, which were used to build the clusters that are studied and compared in this paper

Table 1. The Single Board Computers, which were used to build the Clusters

	Raspberry Pi B	**BananaPi**	**Raspberry Pi 2 B**
CPU family	ARM 11	ARM Cortex A7	ARM Cortex A7
CPU cores	1	2	4
Clock rate	800 MHz[1]	900 MHz	900 MHz
Main memory	512 MB	1024 MB	1024 MB
Ethernet interface	100 Mbit	1000 Mbit	100 Mbit
Storage interfaces	SD	SD, SATA	microSD
[1]Increasing the clock rate from 700 to 800 MHz does not require overvolting the CPU and results in a noticeable increase of the processing power. For this reason, the Raspberry Pi nodes of the cluster were overclocked to 800 MHz			

3. RELATED WORK

In the literature, several works propose implementing high performance clusters of single board computers.

Cox et al. [1] assembled in 2012 at the University of Southampton for a total cost of around £3400 a cluster, called Iridis-pi, of 64 Raspberry Pi nodes with 256 MB main memory per node. This system performed 1.14 Gflops (Flops is an acronym for floating-point operations per second). The nodes were powered by using 64 individual 5 V power supplies. The power consumption of the cluster was not presented.

Balakrishnan [2] constructed in 2012 at the University of Edinburgh a cluster by using six PandaBoard single board computers and two Raspberry Pi nodes, which performed 6.484 Gflops using the six PandaBoard nodes. The work provides the power consumption of the cluster, which is around 170 W in idle state and around 200 W during peak load.

Kiepert [3] assembled in 2013 at the Boise State University a cluster of 32 Raspberry Pi nodes with 512 MB main memory per node, which performed 10.13 Gflops [4]. To power the nodes, he used two standard PC power supplies and attached them by using one of the 5 V pins of the I/O header, each Raspberry Pi provides. The maximum total power usage of the cluster is 167 W.

Abrahamsson et al. [5] presented in 2013 a cluster, called MegaRPi, which was assembled at the Free University of Bozen-Bolzano. The cluster consists of 300 Raspberry Pi nodes with 512 MB main memory per node. To power the cluster, standard PC power supplies were used. The work identified several challenges and a number of opportunities. Additionally, the work compared the power consumption of a single node with other computers when executing a HTTP server benchmark. Unfortunately, no further power measurements or Gflops results were presented.

Sukaridhoto et al. [6] presented in 2013 a cluster of 16 Pandaboard nodes which was assembled at the Electronics Engineering Polytechnics Institute Surabaya. The cluster used a single 200 W, 5 V, 40 A power supply to power the nodes. Only the performance of a single node was measured, which is 0.6 Gflops. The power consumption of the entire cluster was not presented.

Ou et al. [7] compared in 2012 the performance, energy-efficiency and cost-efficiency of a single PandaBoard computer with an Intel X86 workstation for the three applications web server throughput, in-memory database and video transcoding. The work examines how many nodes a cluster of single board computers are required to compete with the workstation.

Tso et al. [8] presented in 2013 the Raspberry Pi Cloud, which was assembled at the University of Glasgow. This cluster is a scale model of a data center, composed of 56 Raspberry Pi Model B nodes, that emulates the layers of a cloud stack. The work compares the acquisition cost, electricity costs (196 W) and cooling requirements of the cluster of single board computers with a

testbed of 56 commodity hardware servers. This work does not provide any performance measurements.

Pfalzgraf and Driscoll [9] assembled in 2014 at the Bradley University a cluster of 25 Raspberry Pi nodes and used a single 600 W PC power supply to power the cluster nodes. This work does not provide any power measurements or Gflops results.

These works show the potential of clusters of single board computers, but none of them compares the performance and energy-efficiency of different single board computer clusters with each other.

4. COMPONENTS OF THE CLUSTERS OF SINGLE BOARD COMPUTERS

Three clusters (see Table 2) have been constructed for this work. The acquisition cost for each cluster was less than 500 €. The prices of all components were checked in March 2015 in Germany and they may vary on the free market.

Table 2. Components of the Clusters

Quantity	Component	Cluster A	Cluster B	Cluster C
8	Raspberry Pi Model B	240 €		
8	Banana Pi		280 €	
8	Raspberry Pi 2 Model B			320 €
8	SD card (16 GB)	40 €	40 €	
8	microSD card (16 GB)			50 €
1	16-Port switch (100 Mbit/s)	40 €		40 €
1	16-Port switch (1 Gbit/s)		70 €	
8	Network cable CAT 5e	10 €	10 €	10 €
1	10-Port USB power supply	40 €	40 €	40 €
8	USB cable USB-A/Micro-USB	10 €	10 €	10 €
1	Power strip	10 €	10 €	10 €
some	Screws, cable ties, spacers, etc.	10 €	10 €	10 €
	Price for the entire cluster	400 €	470 €	490 €

The most expensive components are the nodes used. The price for one Raspberry Pi Model B was around 30 €. The price for one Banana Pi was around 35 € and the price for one Raspberry Pi 2 Model B was around 40 €. A 100 Mbit Ethernet switch is sufficient for the clusters A and C. Because the Banana Pi nodes provide a faster network interface, cluster B was equipped with a Gigabit Ethernet switch. The Raspberry Pi 2 nodes of cluster C require microSD flash memory cards, which are slightly more expensive compared with SD cards. All other components of the clusters are equal with each other.

5. ANALYZING THE CLUSTERS' PERFORMANCE WITH THE HPL

The High Performance Linpack (HPL) benchmark is a method to investigate the performance of cluster systems. As described by Dunlop et al. [10] and Luszczek et al. [11], the benchmark solves a linear system of equations of order n.

$$A \times x = b; \quad A \in \mathbb{R}^{n \times n}; \quad x, b \in \mathbb{R}^n \tag{1}$$

That is divided into blocks of size $P \times Q$, by using double-precision (8 Bytes) floating-point arithmetic (Gaussian elimination with partial pivoting) on computer systems with distributed memory. $P \times Q$ is equal to the number of processor cores used. The developers of the HPL recommend in [12] that P (the number of process rows) and Q (the number of process columns) should be approximately equal, with Q slightly larger than P.

Parameter N specifies the problem size. To find the largest problem size that fits into the main memory of a specific system, the main memory capacity for storing double precision (8 Bytes) numbers is calculated. Utilizing the entire main memory for the benchmark is impossible because the operating system and running processes still consume memory. Thus, it is promising to set N to a value 80-90% of the available main memory [13].

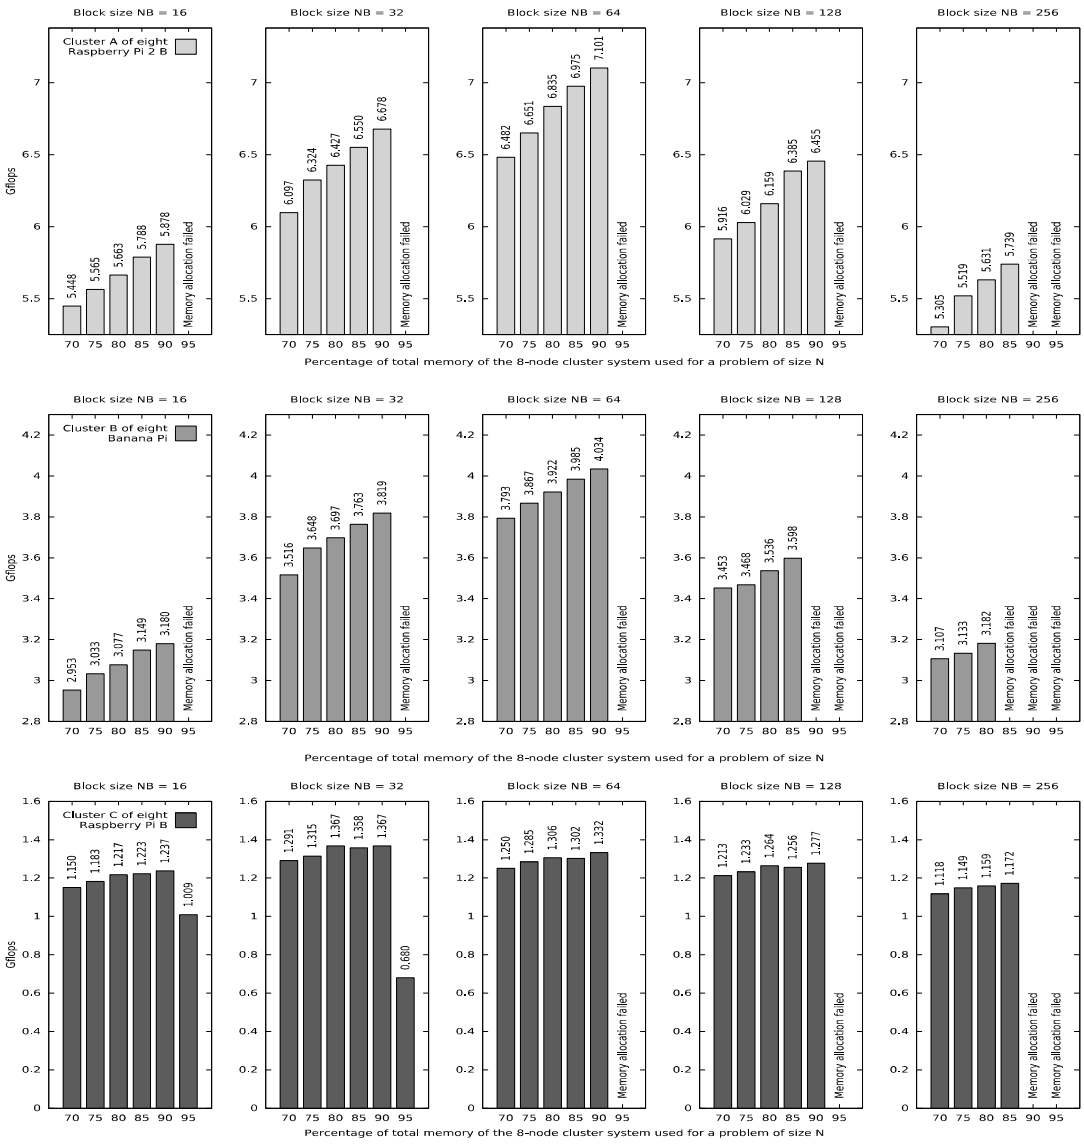

Figure 1. Gflops performance of all three clusters, when using all eight nodes, by using the HPL benchmark. The concrete values for problem size *N* provides Table 3

N can be calculated with equation 2. It depends on the number of nodes X in the system, the reduction coefficient R which specifies how much percent of the entire main memory of the cluster shall be utilized by the benchmark and the main memory capacity M of a single node. Table 1 shows how much main memory the single nodes have. A part of the main memory must be assigned as video memory to the GPU, which lacks own dedicated memory. Because the GPUs are not used at all in the clusters, the minimal GPU memory was set, which is for all nodes 16 MB. This results in 496 MB main memory left on each node of cluster A and 1008 MB on each node of the clusters B and C. After the Raspbian operating system and all daemons are started, approximately 400-430 MB main memory remains available on each node of cluster A. Each node of the clusters B and C has around 890-920 MB free main memory.

$$N = \sqrt{(\frac{M \times 1024 \times 1024 \times 1024 \times X}{8}) * R} \quad (2)$$

If for example N shall be big enough to fill around 80% of the memory capacity of four nodes ($X = 4$) of cluster system C, the calculation is as follows:

$$N = \sqrt{(\frac{1008\ \text{GB} \times 1024 \times 1024 \times 1024 \times 4}{8}) * 0.8} \approx 18{,}610$$

A further important parameter is the block size NB. As optimization, N should be NB aligned [13]. For this example, if we consider $NB = 64$, we calculate $\frac{18{,}610}{64} \times 64 = 290{,}78125 \approx 290$ and next $290 \times 64 = 18{,}560 = N$. For this work, the HPL benchmark was executed with different parameters in the three clusters. Figure 1 shows the Gflops when executing the benchmark with different values for the parameter NB when using all eight nodes and utilizing different proportions of the systems' total main memory. These tests were carried out to find for each cluster the most recommendable value for NB.

As BLAS (Basic Linear Algebra Subprograms) implementation was ATLAS (Automatically Tuned Linear Algebra Software) revision 3.8.4 used on all three clusters.

The results in Figure 1 show that increasing the problem size N results in a better Gflops performance. The performance drop for $NB = 16$ and $NB = 32$ in cluster A, when utilizing 95% of the main memory, is caused by the heavy use of swap memory. The peak performance of cluster A (1.367 Gflops) is achieved with $NB = 32$. The peak performance of cluster B (4.034 Gflops) and cluster C (7.101 Gflops) is achieved with $NB = 64$.

5.1. Analysis Of The Speedup

The Tables 3, 4 and 5 show the values of the parameters N, P, Q and NB, as well as the runtimes, required to solve the linear system and the resulting Gflops.

The benchmark was executed in the clusters with just a single node, two nodes, four nodes and eight nodes to investigate the speedup. The speedup S_P, that can be achieved when running a program on P processors is defined as

$$S_P = \frac{F_P}{F_1} \quad (3)$$

where F_1 is the Gflops on a single-processor system and F_P is the Gflops on a multiprocessor system. The theoretical maximum speedup is equal to the number of single-processor nodes. Therefore it is value 2 for two nodes, value 4 for four nodes, etc.

The best benchmark results were obtained when N is set to a value 80-90% of the available main memory. These results contain Table 3, 4 and 5 and they show that increasing the number of nodes also increases the speedup significantly.

Table 3. Analysis of the performance and speedup of Cluster A (Raspberry Pi B) by using the HPL benchmark

Memory utilized	N	Nodes used	NB	P	Q	Time [s]	Gflops	Speedup
≈ 80%	6496	1	32	1	1	722.29	0.212	1.00
	9216	2	32	1	2	1503.04	0.347	≈ 1.63
	13024	4	32	1	4	2328.03	0.632	≈ 2.98
	18432	8	32	2	4	3055.37	1.367	≈ 6.44
≈ 85%	6912	1	32	1	1	1037.84	0.212	1.00
	9792	2	32	1	2	1705.12	0.367	≈ 1.73
	13856	4	32	1	4	2782.23	0.637	≈ 3.00
	19584	8	32	2	4	3688.51	1.358	≈ 6.40
≈ 90%	7328	1	32	1	1	1246.73	0.210	1.00
	10368	2	32	1	2	1993.88	0.372	≈ 1.77
	14656	4	32	1	4	3274.85	0.641	≈ 3.05
	20768	8	32	2	4	4370.31	1.367	≈ 6.50

Table 4. Analysis of the performance and speedup of Cluster B (Banana Pi) by using the HPL benchmark

Memory utilized	N	Nodes used	NB	P	Q	Time [s]	Gflops	Speedup
≈ 80%	9280	1	64	1	2	928.09	0.574	1.00
	13120	2	64	1	4	1456.33	1.034	≈ 1.80
	18560	4	64	2	4	2120.46	2.010	≈ 3.50
	26304	8	64	2	8	3094.01	3.922	≈ 6.83
≈ 85%	9856	1	64	1	2	1116.91	0.571	1.00
	13952	2	64	1	4	1714.48	1.056	≈ 1.84
	19712	4	64	2	4	2510.06	2.035	≈ 3.56
	27904	8	64	2	8	3635.13	3.985	≈ 6.97
≈ 90%	10432	1	64	1	2	1313.42	0.576	1.00
	14784	2	64	1	4	2018.12	1.068	≈ 1.85
	20928	4	64	2	4	2988.12	2.045	≈ 3.55
	29568	8	64	2	8	4272.77	4.034	≈ 7.00

Table 5. Analysis of the performance and speedup of Cluster C (Raspberry Pi 2 B) by using the HPL benchmark

Memory utilized	N	Nodes used	NB	P	Q	Time [s]	Gflops	Speedup
≈ 80%	9280	1	64	1	4	493.13	1.081	1.00
	13120	2	64	2	4	753.07	2.000	≈ 1.85
	18560	4	64	2	8	1126.66	3.784	≈ 3.50
	26304	8	64	4	8	1775.34	6.835	≈ 6.32
≈ 85%	9856	1	64	1	4	589.52	1.083	1.00
	13952	2	64	2	4	892.98	2.028	≈ 1.87
	19712	4	64	2	8	1324.80	3.855	≈ 3.55
	27904	8	64	4	8	2076.80	6.975	≈ 6.44
≈ 90%	10432	1	64	1	4	684.53	1.106	1.00
	14784	2	64	2	4	1056.70	2.039	≈ 1.84
	20928	4	64	2	8	1594.88	3.832	≈ 3.46
	29568	8	64	4	8	2427.23	7.101	≈ 6.42

5.2. ANALYSIS OF THE EFFICIENCY

Especially for the Top500 list of the most powerful computer systems, two performance indicators are considered important. These are:

- the theoretical peak performance $Rpeak$ of the system. It is determined by counting the number of floating-point additions and multiplications (in double precision), that can be completed during a period of time, usually the cycle time of the machine. The $Rpeak$ of a system can be calculated with equation 4. The ARM 11 and ARM Cortex A7 processors, used by the nodes, can process one floating-point addition in one cycle and require two cycles for a floating-point multiplication. Table 4 contains the $Rpeak$ of the three clusters.

$$Rpeak\ [\text{Gflops}] = \text{clock speed}\ [\text{GHz}] \times \text{cores} \times \text{operations per cycle} \quad (4)$$

- The maximal performance $Rmax$ in Gflops that was achieved with the HPL (see Figure 1 and Table 3, 4 and 5).

The efficiency of a specific system in percent is calculated via $Rmax/Rpeak \times 100$ in case of our clusters the efficiency depends of the executed operations. For cluster A it is between $\approx 21\%$ and $\approx 42\%$, for cluster B it is between $\approx 28\%$ and $\approx 56\%$ and for cluster C it is between $\approx 24\%$ and $\approx 49\%$ (see Table 6). The exact reason for this low efficiency was not evaluated. But the HPC Challenge benchmark test suite stresses not only the CPUs, but the memory system and the network interconnect too [11]. Therefore, the interconnect and memory performance of the single board computers may have a negative impact here [14].

Table 6. The $Rpeak$ and $Rmax$ of the Clusters

	Cluster A	Cluster B	Cluster C
Nodes	8 × Raspberry Pi B	8 × Banana Pi	8 × Raspberry Pi 2 B
Clock rate	800 MHz	900 MHz	900 MHz
CPU cores	8	16	32
$Rpeak$	3.2-6.4 Gflops	7.2-14.4 Gflops	14.4-28.8 Gflops
$Rmax$	1.367 Gflops	4.034 Gflops	7.101 Gflops
Efficiency	$\approx 21\% - 42\%$	$\approx 28\% - 56\%$	$\approx 24\% - 49\%$

6. ANALYSIS OF THE ENERGY-EFFICIENCY

Three options for power supply and its impact to the overall power consumption were evaluated. In one scenario, each cluster node has its own USB power supply (Samsung ETA-U90EWE, 5 V, 2 A). In a second scenario, two 5-port USB power supplies (Anker Model 71AN7105 40 W, 5 V, 8 A) delivered the required electric energy for the nodes. A third scenario, a single 10-port USB power supply (Anker Model A2133311 60 W, 5 V, 12 A) was used for the nodes. Table 7 shows the power consumption of the clusters in idle operation mode and in two different stress mode scenarios. In the first scenario, just the CPUs of the nodes were put into stress mode by using the command-line tool `stress`. In the second scenario, the HPL benchmark was used.

The results in Table 7 show that the more power supplies are used, the more energy is consumed. The reason for this observation is that each power supply wastes a part of the electric energy as heat. Therefore, each additional power supply results in additional waste of electric energy.

Table 7. Power Consumption of the Clusters

Cluster	Power Supply	Idle Mode	Stress Mode (stress)	Stress Mode (HPL)
A	1 × 10-port USB power supply	≈ 23 W	≈ 25 W	≈ 25 W
	2 × 5-port USB power supplies	≈ 24 W	≈ 26 W	≈ 26 W
	8 × 1-port USB power supplies	≈ 30 W	≈ 32 W	≈ 32 W
B	1 × 10-port USB power supply	≈ 27 W	≈ 35 W	≈ 38 W
	2 × 5-port USB power supplies	≈ 28 W	≈ 36 W	≈ 40 W
	8 × 1-port USB power supplies	≈ 33 W	≈ 44 W	≈ 48 W
C	1 × 10-port USB power supply	≈ 15 W	≈ 22 W	≈ 25 W
	2 × 5-port USB power supplies	≈ 16 W	≈ 23 W	≈ 26 W
	8 × 1-port USB power supplies	≈ 20 W	≈ 29 W	≈ 33 W

The energy costs per year C_Y for a 24/7 usage for a specific power consumption in kW during operation E can be calculated with equation 5. In the equation, energy costs of 0.30 € per kWh are assumed.

$$C_Y = \mathrm{E} \times 24\frac{\text{hours}}{\text{day}} \times 365.25\ \frac{\text{days}}{\text{year}} \times 0.30\ \frac{€}{\text{kWh}} \tag{5}$$

In a scenario where cluster C and one 10-port USB power supply runs all the time, the energy cost per year for 24/7 usage is between 39.45 € (if it runs in idle mode with 15 W all the time) and 65.75 € (if it runs in HPL stress mode with 25 W all the time) per year.

Knowing the clusters' electric energy consumption (see Table 7) and its performance when executing the HPL benchmark (see Section 5) is the precondition to analyze the clusters' energy-efficiency.

The Green500 list, which is a complement to the Top500 list, uses the flops per Watt metric [15] to rank the energy efficiency of supercomputers [16]. The metric is defined as

$$\text{flops per Watt} = \frac{Rmax\ [\text{flops}]}{P(Rmax)\ [\text{Watt}]} \tag{6}$$

$P(Rmax)$ is the average system power consumption while executing the HPL with a problem size that delivers $Rmax$. When executing the HPL benchmark, the power consumption of a cluster depends of the number of nodes used for the benchmark.

With $Rmax = 1.367$ Gflops, when using all eight nodes, and consuming 25 W, cluster A provides approximately 54.68 Mflops per Watt. With $Rmax = 4.034$ Gflops, when using all eight nodes, and consuming 38 W, cluster B provides approximately 106.15 Mflops per Watt. With $Rmax = 7.101$ Gflops, when using all eight nodes, and consuming 25 W, cluster C provides approximately 284.04 Mflops per Watt. These results demonstrate that more recent generations of single board computers operate more energy-efficient.

7. CONCLUSIONS AND FUTURE WORK

The performance of single board computer clusters cannot compete with higher-value systems. In the Top500 list from June 1993, the maximum observed performance $Rmax$ of cluster system A (1.367 Gflops) would be sufficient for 216th place, $Rmax = 4.034$ Gflops of cluster B would be sufficient for 45th place and $Rmax = 7.101$ Gflops of cluster B would be sufficient for 29th place. The most recent Top500 list in which our best performing cluster could achieve a place is the list of November 1996 where cluster C would be sufficient for 277th place.

Compared with recent cluster sites this performance is very low. In the most recent Top500 list from November 2015, the last entry (500th place) provides $Rmax = 206{,}400$ Gflops, which is more than factor $150{,}000$ better compared with cluster A, more than factor $51{,}000$ better compared with cluster B and more than factor $29{,}000$ better compared with cluster C.

Also the energy-efficiency cannot compete with higher-value systems. Cluster A provides approximately 54.68 Mflops per Watt, which would be sufficient for 166th place in the Green500 list from November 2007. Cluster B (106.15 Mflops per Watt) would be sufficient for 50th place and Cluster C (284.04 Mflops per Watt) would be sufficient for 6th place in the same list.

In the most recent list from November 2015, the best entry 1st place) provides 7,031.58 Mflops per Watt, which is more than factor 128 better compared with cluster A, more than factor 66 better compared with cluster B and more than factor 24 better compared with cluster C.

Regardless of the performance or energy-efficiency, clusters of single board computers like the evaluated ones are useful for several academic purposes and research projects because of the lesser purchase costs and operating costs compared with commodity hardware server resources.

Since March 2016, the Raspberry Pi 3 is available for purchase. This single board computer provides more computational power compared with the cluster nodes that were used in this project. Building a cluster of these computers is one of the next steps. Such a cluster should also have more than just eight nodes to further analyze the potential suitability for larger-scale parallel applications.

ACKNOWLEDGEMENTS

This work was funded by the Hessian Ministry of Higher Education, Research, and the Arts ('*Hessisches Ministerium für Wissenschaft und Kunst*') in the framework of research for practice ('*Forschung für die Praxis*').

REFERENCES

[1] Cox, S. & Cox, J. & Boardman, R. & Johnston, S. & Scott, M. & O'Brien, N (2009) "Iridis-pi: a low-cost, compact demonstration cluster", *Cluster Computing*, Vol. 17, No. 2, pp 349-358.

[2] Balakrishnan, N. (2012) "Building and benchmarking a low power ARM cluster". Master's thesis, University of Edinburgh.
[3] Kiepert, J. (2013) "Creating a Raspberry Pi-Based Beowulf Cluster", Technical report, Boise State University.
[4] Vaughan-Nichols, S (2013) "Build your own supercomputer out of Raspberry Pi boards", ZDNet.
[5] Abrahamsson, P. & Helmer, S. & Phaphoom, N. & Nicolodi, L. & Preda, N. & Miori, L. & Angriman, M. & Rikkila, J. & Wang, X. & Hamily, K. & Bugoloni, S. (2013) "Affordable and energy-efficient cloud computing clusters: The bolzano raspberry pi cloud cluster experiment", *Cloud Computing Technology and Science*, IEEE 5th International Conference on CloudCom 2013, Vol. 2, pp 170–175.
[6] Sukaridhoto, S. & KHalilullah, A. S. & Pramadihanto, D. (2013) "Further Investigation on Building and Benchmarking A Low Power Embedded Cluster for Education", 4th International Seminar on Applied Technology, Science, and Art (APTECS).
[7] Ou, Z. & Pang, B. & Deng, Y. & Nurminen, J. K. & Ylä-Jääski, A. & Hui, P. (2012) "Energy- and Cost-Efficiency Analysis of ARM-Based Clusters", 12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), pp 115–123.
[8] Tso, F.P. & White, D.R. & Jouet, S. & Singer, J. & Pezaros, D.P. (2013) "The Glasgow Raspberry Pi Cloud: A Scale Model for Cloud Computing infrastructures", 1st International Workshop on Resource Management of Cloud Computing (CCRM).
[9] Pfalzgraf, A. M. & Driscoll, J. A. (2014) "A Low-Cost Computer Cluster for High-Performance Computing Education" IEEE International Conference on Electro/Information Technology, Milwaukee, pp 362-366.
[10] Dunlop, D. & Varrette, S. & Bouvry, P. (2010) "Deskilling HPL" , Volume 6068 of *Lecture Notes in Computer Science*, Springer, Heidelberg, Berlin, pp 102–114.
[11] Luszczek, P. & Dongarra, J. & Koester, D. & Rabenseifner, R. & Lucas, B. & Kepner, J. & McCalpin, J. & Bailey, D. & Takahashi. D. (2005) "Introduction to the HPC Challenge Benchmark Suite", Technical Report, ICL, University of Tennessee at Knoxville.
[12] HPL Frequently Asked Questions, http://www.netlib.org/benchmark/hpl/faqs.html
[13] Sindi, M. (2009) "HowTo – High Performance Linpack (HPL)", Technical Report, Center for Research Computing, University of Notre Dame.
[14] Baun, C. (2016) "Mobile clusters of single board computers: an option for providing resources to student projects and researchers", SpringerPlus 5(360).
[15] Sharma, S. & Hsu, C.-H. & Feng W.-c. (2006) "Making a Case for a Green500 List", IEEE 20th International Parallel and Distributed Processing Symposium 2006.
[16] Ge, R. & Feng, X. & Pyla, H. & Cameron, K. & Feng, W. (2007) "Power Measurement Tutorial for the Green500 List", The Green500 List: Environmentally Responsible Super-computing.

Authors

Dr. Christian Baun is working as Professor at the Faculty of Computer Science and Engineering of the Frankfurt University of Applied Sciences in Frankfurt am Main, Germany. His research interest includes operating systems, distributed systems and computer networks.