To link to this article : DOI : 10.1002/cpe.4073
URL : http://dx.doi.org/10.1002/cpe.4073

To cite this version : Astsatryan, Hrachya and Narsisian, Wahi and Kocharyan, Aram and Da Costa, Georges and Hankel, Albert and Oleksiak, Ariel *Energy Optimization Methodology for e-Infrastructure Providers.* (2017) Concurrency and Computation: Practice and Experience, vol .29 (n° 1). pp. 1. ISSN 1532-0626

DOI: 10.1002/cpe.4073

Energy optimization methodology for e-infrastructure providers

Hrachya Astsatryan[1] | **Wahi Narsisian[1]** | **Aram Kocharyan[1]** | **Georges Da Costa[2]** | **Albert Hankel[3]** | **Ariel Oleksiak[4]**

[1]Institute for Informatics and Automation Problems of the National Academy of Sciences of the Republic of Armenia, 1, P. Sevak Str. Yerevan, 0014, Armenia
[2]Institut de Recherche en Informatique de Toulouse, 118 Route de Narbonne Toulouse, F-31062 Toulouse Cedex 9, France
[3]SURFNET, Moreelsepark 48 Utrecht, 3511 EP, Netherlands
[4]Poznan Supercomputing and Networking Center, ul. Noskowskiego 10 Poznan, 61-704, Poland

Correspondence
Hrachya Astsatryan, Institute for Informatics and Automation Problems of the National Academy of Sciences of the Republic of Armenia, 1, P. Sevak Str., Yerevan 0014, Armenia.
Email: hrach@sci.am

Funding Information
GN4 VI-SEEM EaPConnect Projects and the COST Action IC1305 Network for Sustainable Ultrascale Computing (NESUS) State Thematic Project

Summary

The environmental protection is a dominant concern for all types of industries, organizations, and governments. In this regard, the reduction of the energy consumption is substantial in bringing down the CO2 gas emission, which is considered as an important factor causing global warming. The e-infrastructure service providers, such as National Research and Education Networks or National Grid Initiatives have crucial role in the context of energy awareness because the energy consumption of the networking, data, and computational infrastructures keeps increasing exponentially over the time. In addition to this, scientific gateways and cloud services are becoming more significant to tackle scientific and societal challenges. Therefore, there is a need to provide robust and reliable services taking into account energy consumption aspect of e-infrastructures. The aim of the article is to introduce an energy optimization methodology for the beneficiaries of the e-infrastructures to explore, optimize, and report the energy consumption and CO2 emission of data, computing, and networking facilities. The suggested methodology has been implemented within the Armenian e-infrastructure aiming at the reduction of the energy consumption and thereby the CO2 emission.

KEYWORDS

e-infrastructure, energy optimisation, DVFS, green ICT, memory ballooning, virtualization

1 | INTRODUCTION

In the last decade, the e-Infrastructures became vital in enabling high-quality research and information and communications technology (ICT) developments by providing networking, data, and computational resources.[1] The e-Infrastructures consisting of data centers, networks, and computational facilities are a cost-effective solution for hosting computational and data intensive applications within the distributed infrastructures for research purpose.

However, e-infrastructures incur tremendous energy costs and CO2 emissions. In electricity consumption, ICT (including datacenters) is already going upwards to 8-10% of global consumption[2] and the total global footprint of ICT is 2% of global emissions.[3] For example, the CO2 emissions in tonnes of the backbone network and the data centre of the Greek Research and Educational Network in 2015 were about 807 and 2483, which is a part of the South East European Grid Infrastructure consisting of 55 resource centers with more than 6600 CPUs and 750 TB of disk storage distributed in 16 participating countries.[3]

The electricity consumption of such infrastructures is increasing exponentially because of the growth in amount of resources and digital data, which needs high-performance processing and big storage resources, as well as high-bandwidth network. Because of these facts, data centers can be considered as one of the biggest power consumption platforms contributing the highest CO2 emission among the ICT infrastructures.[4] Therefore, it is important to minimize the energy consumption of these infrastructures. This needs to be done on different levels, such as virtualization of computational resources in deploying cloud infrastructures by shifting the most of the dedicated servers to virtual machines or to executing optimization on software levels to decrease the computing time of all used applications. For instance, the energy efficiency of e-infrastructures has become a major issue with the growing demand of cloud computing, as more and more companies are investing in building large datacenters to host cloud services.[5]

Becoming more informed on CO2 emission, users and providers will have more insight on the impact of their infrastructure usage, and thus will be able to reduce this impact and their ecological footprint.

The aim of the article is to introduce an energy optimization methodology for the beneficiaries of the e-infrastructures to explore, optimize, and report the energy consumption and CO2 emission of data, computing, and networking resources. This method has been applied for the Armenian heterogeneous networking (Academic Scientific Research Computer Network of Armenia) and computational (Armenian National Grid Initiative) resources and in a case study to demonstrate it and to illustrate its effects in example experiments.

The content of this paper is organized as follows. The methodology is presented in Section 2. The report about energy consumption and CO2 emission in details can be found in Section 3. The energy minimization methods are given in Section 4. Finally, the conclusion and directives for future research are drawn in Section 5.

2 | METHODOLOGY

To optimize the energy consumption and reduce CO2 emission and based on the green policy terms, which are accepted by all beneficiaries, we find that it is important to take an effective and easy steps to achieve this goal. A green policy within the regulation agreement

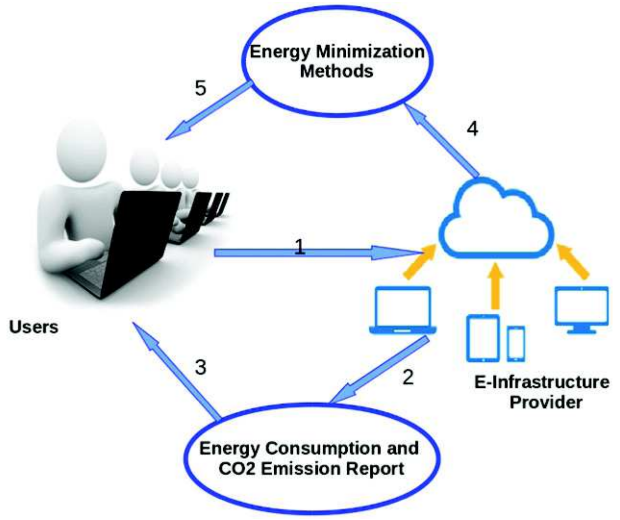

FIGURE 1 Methodology skeleton

between the (NREN) and the beneficiaries using the networking facilities has been adopted, and all the beneficiaries accepted their term, optimizing the energy consumption and reducing CO2 emission will be based on this policy steps.

The energy optimization methodology (see Figure 1) enables both infrastructure providers and beneficiaries to explore, optimize, and report the energy consumption of data, computing, and networking facilities.

In the first phase, to have a clear view of energy consumption of the applications or facilities used by the beneficiaries, a green report is prepared, which is introduced in Section 3. The report enables beneficiaries and stakeholders to calculate and analyze the energy consumption of their applications and resources, which needs to be tackled with the aim of using more energy-aware equipment and optimizing the energy consumption of their applications. For instance, in case of networking facilities to reduce the number of the most energy consuming equipments. Thus, there will be a suggestion to replace them by more energy-aware equipment, conduct this approach for future purchase plan, or reduce the number of network devices (routers, switches, etc) without affecting the overall network performance.

In the next stage, we try to minimize the energy consumption of applications using high-performance computational or data resources. The arrows 4 and 5 on the Figure 1 indicate the special methods, which are suggested and implemented by the provider to optimize the applications of beneficiaries. Special methods and energy aware computational resources are implemented to optimize applications. For instance NAMD,[6] GROMACS,[7] CHARMM,[8] and AMBER[9] packages are widely used by life science user community and it is important to optimize them.

The Armenian e-infrastructure is used for the studies and experiments, which is a complex national IT infrastructure consisting of both networking and distributed computing infrastructures.[10] Academic Scientific Research Computer Network of Armenia is the core-networking infrastructure providing a high-quality infrastructure

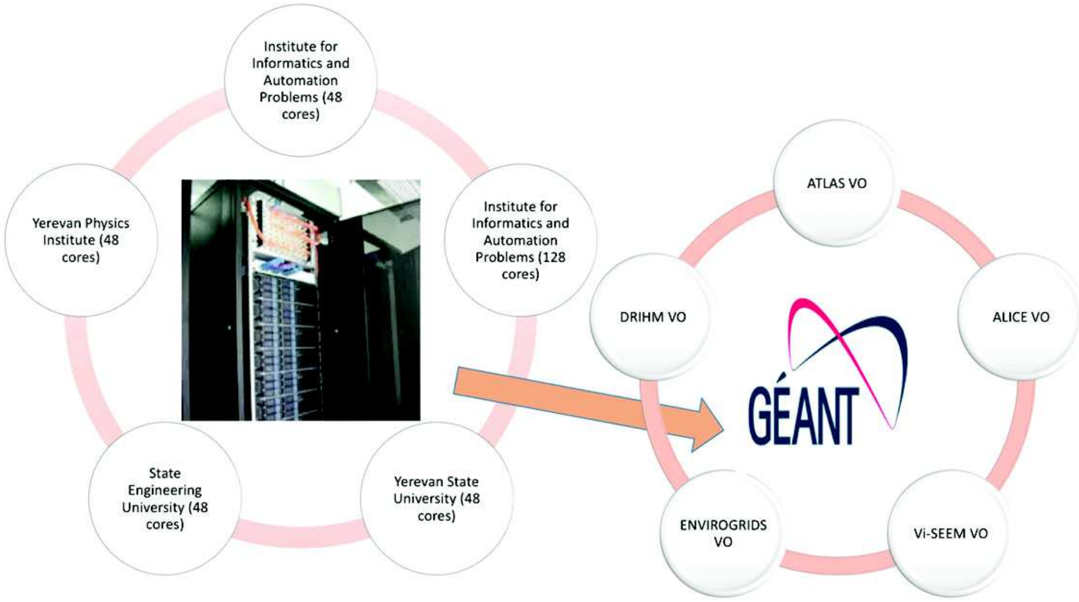

FIGURE 2 Topology of Armenian National Grid Initiative (ArmNGI) infrastructure

and services to the academic, research, and educational community ofArmenia by connecting more than 60 scientific research, academic, and cultural organizations. The backbone of Academic Scientific Research Computer Network of Armenia consists of network communication nodes in 6 cities in Armenia, which are interconnected by fiber optics and wireless links.

The core computational infrastructure is the Armenian National Grid Initiative, which is a national effort to establish a nationwide grid environment for computational science and research.[11,12] Now, the computational resources (about 500 cores) of Armenian Grid infrastructure distributed among our leading research (National Academy of Sciences, Yerevan Physics Institute) and academic (Yerevan State University, State Engineering University of Armenia) organizations are located in the cities of Yerevan and Ashtarak (see Figure 2).

3 | ENERGY CONSUMPTION AND CO2 EMISSION REPORT

As the energy consumption of computational, data, and networking infrastructures are growing very fast, e-infrastructure providers take a responsibility of optimizing of the energy consumption on different levels to gain economic and ecological sustainability. To overcome the challenge, it is important to direct the whole infrastructure and community towards low-carbon operations through different actions and implemented services, such as the following:

- Maintain and apply some standards, ISO standards-based audits of the greenhouse gas emissions. Particularly, the ISO-14064 standard is used in this regard for its carbon audits.
- Monitor and report all types of energy consumption and CO2 emission to get a vivid view about the situation through all infrastructures.
- Save energy by using energy efficient equipments (switches and routers) by replacing available high–energy-consuming equipment with the energy-aware equipment and to have the same approach for the planning equipment, reduce travels, and use online platforms instead, recycle when possible.
- Promote and push vendors and other suppliers to improve their environmental standards through procurement and service level agreements.
- Actively get involved in different research activities in this domain.
- Collaborate with HE institutions, governments, and industry partners to maximize overarching impacts and avoid local maxima.

One of the most important steps toward green ICT is to have annual report about CO2 emission, explore what is going on, and depend on this report to maintain and take some actions. The annual report is the main or starting point to have a general overview of energy consumption allowing to perform several activities. The management will use these reports to yearly decide which actions to take to further reduce energy consumption and CO2 emissions. The eCO2meter web tool[13] is used for the report and collecting of data by calculating the yearly energy consumption and GHG emissions. The eCO2meter is able to support the automatic extraction of statistics, the automatic conversion to CO2 emissions, and comparisons among different sectors. There are 4 categories for CO2 emission in the report: the office, the data centers, the network, and the transportation.

Using the tool, only the network and computational infrastructures were taken into account, by collecting all necessary information about available network backbone, computational, and data devices; their energy consumption was measured and the CO2 emission was indicated from each device. In the data centers, the total energy consumption of the IT storage, computing, and networking equipment is measured along with the energy consumption caused by the supportive infrastructure. The use of eCO2meter tool will help to partially automate the report generation, based on some user inputs, such as CO2-eq per KWh, cost of 1 KWh (in Euro), etc. The report includes a review of all emission related impact, enables to find ways to reduce it by providing user-friendly guidelines to all beneficiaries, and develop a Maturity Model to help NRENs in assessing how sustainable their data centers really are in their operational environment.

4 | ENERGY MINIMIZATION METHODS

The power consumption of computational resources is becoming a crucial challenge in the context of increasing the performance regardless of energy consumption. Therefore, finding ways to improve energy efficiency becomes a main issue for HPC applications.[14] In addition, all the dedicated servers and services within the Armenian e-Infrastructures have been transferred to virtual environment. By doing this, the energy consumption will be reduced due to minimizing the power used for cooling purposes for decreased number of dedicated servers, also, several servers will be hold on sleep state when there will be no request and jobs on them, etc. The work can be divided into optimizing the mostly used HPC applications and to virtualize some dedicated servers.

4.1 | Virtualization

Almost 90% of all the dedicated servers and services within the Armenian e-infrastructures have been transferred to virtual environment; this action reduced the energy consumption because of 2 important reasons:

- Minimizing a power used for cooling because the number of dedicated servers are decreased.
- Optimizing the resource usage, by holding them on sleep state when they are not used which is more easy task on virtual machines.

4.2 | Improve energy efficiency of HPC applications

The improvement is done by 2 approaches: dynamic voltage and frequency scaling (DVFS) and memory ballooning. The methods we use enable power capping of CPU and memory subsystems for modern hardware facilities on the basis of the analyses and studies of well-used applications by scientific communities.

4.2.1 | Dynamic voltage and frequency scaling

The DVFS is a very popular and widely used technique for various types of processors to save power by reducing their operating frequency and voltage.[15] The DVFS technique is applied in different ways:

- Constraint: real-time when critical energy consumption is observed.
- Policy determination: when we apply it statically or dynamically by some automation scripts.

The main methodology of our work is based on conducting several simulations and making benchmarking based on the results of those simulations. Series of experiments have been performed for life science community to study the benefits of the DVFS method for both CPU and GPU resources.

Several systems consisting of 5, 24, and 134 K, and about 2.5 M atoms, have been used for the simulations with the following parameters using GPU K40 and Tesla 1060 devices: 2fs time step, the PME electrostatics, and van der Waals forces truncated at 0.9 nm with the corresponding temperature control (v-rescale algorithm).[16] The experiments have been performed taking the average power consumptions of 10 executions, and it has been done typically for 10 000 steps without any writing outputs (see Figure 3). We stated that the performance hardly depends on the size of system studied and in case of small system, the parallelization is reduced.

To avoid wasting computational power, we recommend GPU with low frequency or few cores combination when the size of the system studied does not exceed the so-called "critical amount" of atoms such as for 5 K atoms, which in its turn will reduce the energy consumption on the system.

The optimization of memory usage has been studied and implemented to reduce the power consumed by all virtual machines. As the most and well-known approaches, the memory ballooning and page sharing techniques have been implemented to address the challenge.[17]

4.2.2 | Memory optimization

Memory ballooning is a memory optimization method introduced by VMware. However, it became so popular that almost all the hypervisors have implemented and supported this functionality to their products. It uses pseudodriver called balloon driver. The balloon inflates when it needs to allocate free memory pages in VMs and deflates in case it deallocates. The memory ballooning is an efficient technique to redistribute the memory resources to the VMs that have a memory shortage. To check the efficiency of the memory ballooning, we used 2 VMs (2 cores and 2GB of RAM initially each of them) for the experiments (see Figure 4) in the domain of linear algebra calculations.[18]

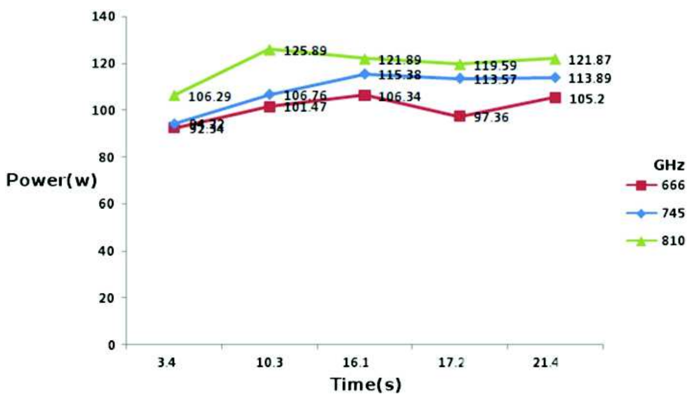

FIGURE 3 Gromacs execution of molecular system consists of 134 K atoms using several GPU frequencies on K40

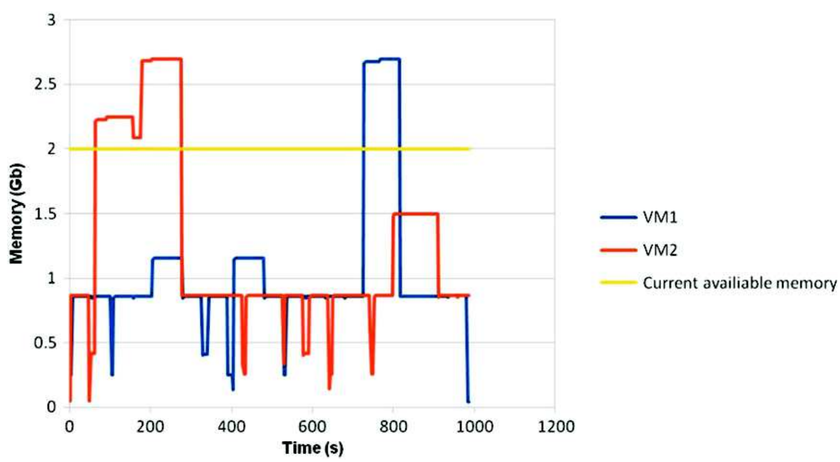

FIGURE 4 Memory usage of VMs before implementing memory ballooning

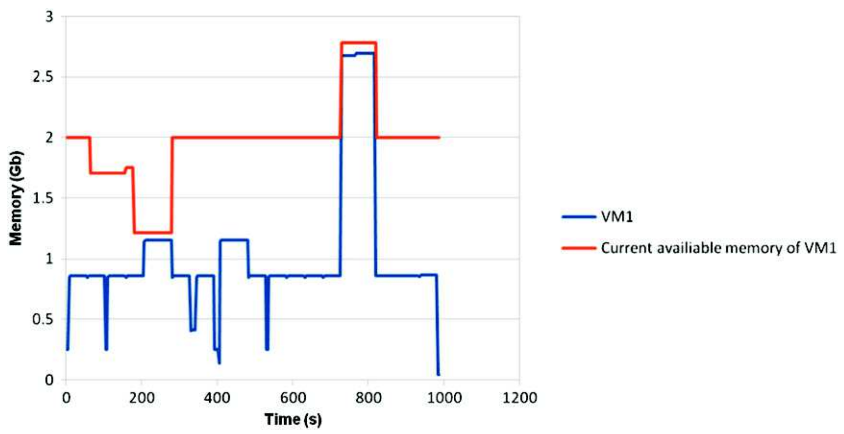

FIGURE 5 Memory usage of VM1 after implementing memory ballooning

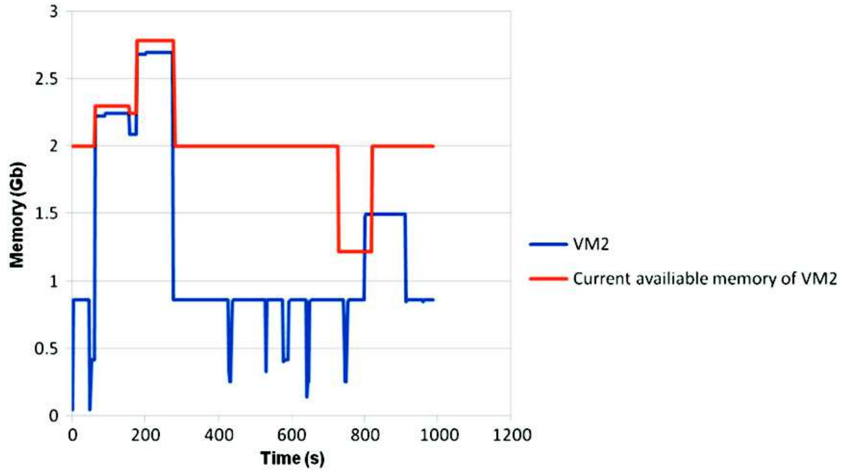

FIGURE 6 Memory usage of VM2 after implementing memory ballooning

According to the results, both VMs have memory shortage and overhead. However, the key point is that it happens at the different period of time, so the usage of memory ballooning technique optimizes the memory consumption as it is illustrated in the next 2 figures (see Figure 5 and 6).

According to the graphs, the swap activity has decreased for about 95%, which means that we have rebalanced the memory resources without affecting the efficiency of any VMs.

Another approach for energy minimization on memory level is the page-sharing method, which is a perfect solution for homogeneous systems introduced by Disco systems.[19] The idea behind the technique is to periodically scan the memory to find identical memory pages. In case of success, it merges them and redefine their addresses in for the processes that were using them. The experiments have been performed on 3 homogeneous dedicated servers. All VMs are running instances of the same guest operating system (Ubuntu server 14.04); use the same applications and components loaded to tackle linear algebra calculations by sharing common input data. In this case, the implementation of the method allows to raise consolidation of the VMs on servers by merging identical memory pages which in its turn allows to free up one of the servers and to put it on sleep mode, which saves about 13% of energy consumption.

5 | CONCLUSION AND FURTHER ACTIVITIES

The suggested energy optimization methodology allows e-infrastructure providers and users to explore, optimize, and report the energy consumption and CO2 emissions of data, computing, and networking facilities. For instance, it allows to reduce the energy consumption of Armenian e-infrastructure by about 12% without changing any HPC application source code. We plan to investigate all other packages, which are used widely by user communities but are not optimized yet to find out an effective way to optimize them on both hardware or software levels. A user-friendly interface will be developed for all beneficiaries to fill all required data to gain a green report later in an easier way.

ACKNOWLEDGEMENTS

The European Commission through the GN4, VI-SEEM, EaPConnect Projects and the COST Action IC1305, "Network for Sustainable Ultrascale Computing (NESUS)," and State Thematic Project entitled "Towards the Green e-Infrastructures" supported the work presented in this paper.

REFERENCES

1. Gentzsch W. E-infrastructures. *Proceedings of the 18th International Conference on Database and Expert Systems Applications (DEXA)*, Springer Lecture Notes in Computer Science, vol. 1997, Regensburg, Germany; September 2007:895–904.
2. Gelenbe E, FACM CaseauY. The impact of information technology on energy consumption and carbon emissions. Magazine Ubiquit, Volume 2015 Issue June, June; 2015.
3. Global e-Sustainability Initiative 2015. GeSI MARTer 2030: ICT solutions for 21st century challenges. Available: http://smarter2030.gesi.org, 2015.
4. Bala A, Prnjat O, Vudragovic D, Slavnic V, Liabotis I, Atanassov E, Jakimovski B, Savic M. Development of grid e-infrastructure in South-Eastern Europe. *Grid Comput*. 2011;9(2):135–154. doi:10.1007/s10723-011-9185-0.
5. Kliazovich D, Bouvry P, Khan S. A packet-level simulator of energy-aware cloud computing data centers. *J Supercomput*. 2012;62(3):1263–1283.
6. Garg SK, Yeo CS, Buyya R. Green cloud framework for improving carbon efficiency of clouds. *Proceedings of Euro-Par 2011 Parallel Processing*, Lecture Notes in Computer Science, vol. 6852. Berlin: Springer; 2011:491–502.
7. Kale L, Skeel R, Bhandarkar M, Brunner R, Gursoy A, Krawetz N, Phillips J, Shinozaki A, Varadarajan K, Schulten K. NAMD2: greater scalability for parallel molecular dynamics. *Comput Phys*. 1999;151(1):283–312. doi:10.1006/jcph.1999.6201.
8. van der Spoel D, Lindahl E, Hess B, Groenhof G, Mark AE, Berendsen HJS. GROMACS: fast, flexible, and free. *Comput Chem*. 2005;26(16):1701–1718.
9. Brooks BR, Bruccoleri RE, Olafson BD, States DJ, Swaminathan S, Karplus M. CHARMM: a pro-gram for macromolecular energy, minimization, and dynamics Calculations. *Comput Chem*. 1983;4(2):187–217. doi:10.1002/jcc.540040211.
10. Pearlman DA, Case DA, Caldwell JW, Ross WR, Cheatham TE, DeBolt S, Ferguson D, Seibel G, Kollman P. AMBER, a computer program for applying molecular mechanics, normal mode analysis, molecular dynamics and free energy calculations to elucidate the structures and energies of molecules. *Comput Phys Commun*. 1995;91(1):1–41. doi:10.1016/0010-4655(95)00041-D.
11. Astsatryan H, Sahakyan V, Shoukourian Yu, Cros PH, Dayde M, Dongarra J, Oster P. Strengthening compute and data intensive capacities of Armenia. *IEEE Proceedings of 14th RoEduNet International Conference - Networking in Education and Research (NER'2015)*, Craiova, Romania; September 2015:28–33.
12. Shoukourian Yu, Sahakyan V, Astsatryan H. E-infrastructures in Armenia: virtual research environments. *IEEE Proceedings of the International Conference on Computer Science and Information Technologies (CSIT'2013)*, Yerevan, Armenia; 2013:1–7.
13. Hankel A, Astsatryan H, Narsisian W. The GEANT green team: an example of how organizations can use a community approach to promote the use of ICT in sustainability efforts. *IEEE Proceedings of the International Conference on Computer Science and Information Technologies (CSIT'2015)*, Yerevan, Armenia; September 2015:125–127.
14. D'Hollander EH, Dongarra JJ, Foster I, Grandinetti L, Joubert GR. *Transition of HPC towards exascale computing*: IOS Press; 2013;232.
15. Guérout T, Monteil T, da Costa G, Calheiros RN, Buyya R, Alexandru M. Energy-aware simulation with DVFS. *Simul Modell Pract Theory*. 2013;39:76–91.
16. Poghosyan A, Arsenyan L, Astsatryan H. Dynamical features of complex systems: a molecular simulation study. *HP-SEE User Forum 2012*, Springer International Publishing Modelling and Optimization in Science and Technologies, vol. 2. Belgrade, Serbia; 2014:117–122, doi:10.1007/978-3-319-01520-0_14.
17. Carl A. Waldspurger, memory resource management in VMware ESX server. *Proceedings of the 5th symposium on Operating systems design and implementation*, Boston, Massachusetts; December 2002:181–194.
18. Astsatryan H, Sahakyan V, Shoukouryan Yu, Daydé M, Hurault A, Guivarch R, Terzyan A, Hovhannisyan L. Services enabling large-scale linear systems of equations and algorithms based on integrated P-grade portal. *Grid Comput*. 2013;11(2):39–248. doi:10.1007/s10723-013-9254-7.
19. Bugnion E, Devine S, Rosenblum M. Disco: running commodity operating systems on scalable multiprocessors. *Proceedings of the 16th ACM Symposium on Operating Systems Principles*; Saint-Malo, France, October 5-8, 1997:412–447.