CMST 22(2) 87-93 (2016) DOI:10.12921/cmst.2016.22.02.003

Highly Scalable Quantum Transfer Matrix Simulations of Molecule-Based Nanomagnets on a Parallel IBM BlueGene/P Architecture

Michał Antkowiak[1*], Łukasz Kucharski[1], Ryszard Matysiak[2], Grzegorz Kamieniarz[1]

[1] *Faculty of Physics, A. Mickiewicz University Umultowska 85, 61-614 Poznań, Poland*

[2] *Institute of Engineering and Computer Education, University of Zielona Góra, ul. Prof. Szafrana 4a, 65-516 Zielona Góra, Poland*

[*] *E-mail: antekm@amu.edu.pl*

Received: 15 June 2015; revised: 04 May 2016; accepted: 18 May 2016; published online: 31 May 2016

Abstract: In this work we present a very efficient scaling of our two applications based on the quantum transfer matrix method which we exploited to simulate the thermodynamic properties of Cr_9 and Mn_6 molecules as examples of the uniform and non-uniform molecular nanomagnets. The test runs were conducted on the IBM BlueGene/P supercomputer JUGENE of the Tier-0 performance class installed in the Jülich Supercomputing Centre.

Key words: parallelization of processing, MPI, numerical simulations, magnetic rings, Heisenberg model

I. INTRODUCTION AND PHYSICAL BACKGROUND

A lot of research on molecular nanomagnets has been done in the last years [1]. Such intensive scientific interest was to a large extent stimulated by expected applications in molecular spintronics, information storage [2], quantum information processing [3-6] and other. Yet even if these applications seem to be a matter of a rather distant future, molecular magnets are already very interesting materials where quantum effects can be observed in bulk samples. Thus, they are a very good test ground for various microscopic theories and computation tools. The interest in spin assemblies stems from the fact that they set the low-size limit for magnetic nanoparticles. They can display such behavior as magnetic quantum tunneling [7] and quantum-size effects in the thermodynamic properties [8].

One of the interesting phenomena that interfere with the quantum nature of the molecules is magnetic frustration. It appears e.g. in antiferromagnetic odd-membered molecular rings based on transition metal ions [9-13]. To observe the frustration effects, new odd nine-numbered homometallic rings Cr_9 have recently been synthesized. They are based on Cr^{3+} ions [14, 15] with spin $S = 3/2$ and their example is shown in Fig. 1. To simulate frustrated molecules it is impossible to use probabilistic methods like Monte Carlo.

An example of hetero-spin molecules with ferromagnetic couplings which we consider here [16, 17] is $[Mn(hfac)_2$ - NITPh$]_6$ compound (hfac, hexafluoroacetyl acetonate; NITPh, 2-phenyl-4,4,5,5-tetramethyl-4,5-dihydro-1H-imi-dazolyl-1-oxy -3-oxide) pictured in Fig. 2 referred to as Mn_6 molecule. It contains twelve paramagnetic centers (six Mn^{2+} ions with spin $S = 5/2$ and six organic radicals with non-paired $S = 1/2$ electrons) and belongs also to a class of nanocompounds actively investigated for their magnetic properties [18, 19].

The molecular magnetic rings are quantum objects modeled by the Heisenberg spin Hamiltonian. For a uniform

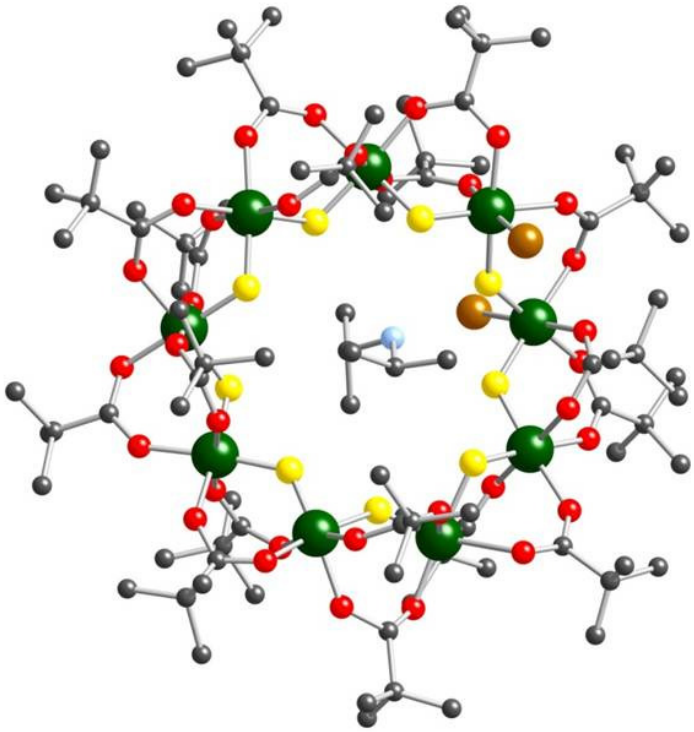

Fig. 1. The crystallographical structure of the Cr_9 molecule which consists of nine chromium spins $S = 3/2$ surrounded by organic radicals and is represented by the chemical formula $[^{i}Pr_2NH_2][Cr_9F_9(O_2C^{t}Bu)_{17}Cl_2]$

ring composed of N spins S, the dimension of the corresponding matrix representation of the Hamiltonian amounts to $(2S+1)^N$. This implies that within the exact diagonalization (ED) method, a solution of a physical problem needs the memory size of the order of $(2S+1)^{2N}$, whereas the computational complexity raises even more to $(2S+1)^{3N}$. Such an exponential complexity imposes strong bounds on the physical systems which can be considered within ED so that there are many methods devised to overcome the ED constraints. Among them the quantum transfer-matrix method (QTM) deserves special attention [20, 21] due to its high accuracy and direct addressing the partition function which is a key quantity needed to calculate the free energy and its derivatives, in particular the magnetic susceptibility, specific heat or magnetization. These quantities are measured experimentally and need theoretical interpretation. We note that the convergence of the QTM results depends on temperature and is deteriorated approaching the limit $T = 0$ and then the ED technique is more convenient. [15, 22, 23]

Within the QTM method, an exponential of the Hamiltonian which is called transfer matrix is addressed and can be found via the Suzuki-Trotter transformation [20, 21]. The transfer matrix has the same size as that in the ED method (i.e. $(2S+1)^N$), but it is expressed in a basis of a simple product of the Ising states. Its structure is complicated and contains m-fold replicas of the initial Hamiltonian, but can be reduced into a product composed of m sparse matrices which are easy to implement. Because the partition function is related to a trace of the transfer matrix, the QTM computational complexity is determined by the number of the diagonal elements $(2S+1)^N$, the number of the basis vectors $(2S+1)^N$ and size m of the products of the sparse matrices and amounts to $m(2S+1)^{2N}$. The complexity remains exponential but it is lower than that of the ED method due to the relation $m \ll (2S+1)^N$. Moreover, the diagonal matrix elements are determined by the scalar products of the same basis vectors so that they can be calculated independently and the trace of the transfer matrix can be efficiently executed in the parallel computing architecture. We also note the small QTM memory size needed which is only of the order of $(2S+1)^N$ and is easily accessible on supercomputing platforms.

The foregoing description refers to the QTM simulations of the non-uniform spin rings consisting of two types of species with the spin values S_1 and S_2. In this case it is enough to map $(2S+1) \rightarrow (2S_1+1)(2S_2+1)$ and consider N as a number of *pairs* of spins S_1 and S_2.

The computations needed to demonstrate the high scalability of our applications intended to simulate the homo- and hetero-spin molecular rings (referred to as 1 and 2, respectively) were performed on JSC JUGENE supercomputer

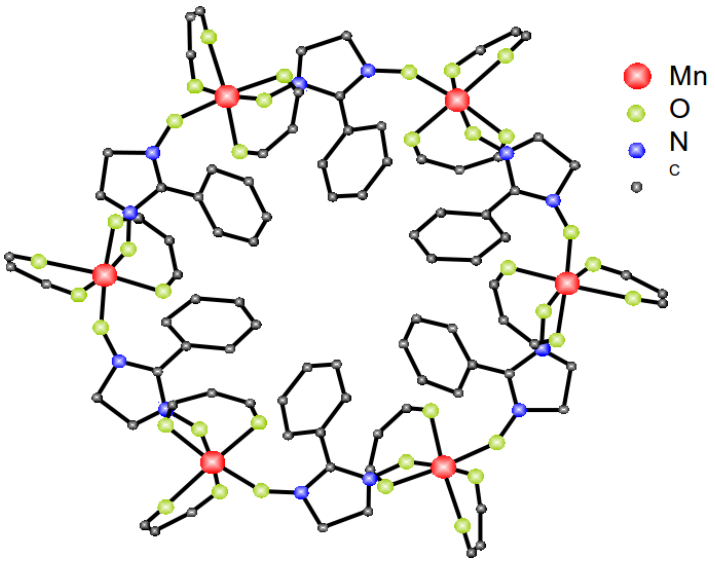

Fig. 2. A scheme of the Mn_6 molecule composed of six manganese ions carrying the spin $S = 5/2$ and six unpaired $S = 1/2$ electrons localized on the adjacent radicals

during the BlueGene/P Extreme Scaling Workshop 2011 organized in Jülich. At that time it was the biggest supercomputer in the world in terms of the number of computing cores.

II. DESCRIPTION OF THE CODES AND ITS PARALLELIZATION

Our codes have been designed to calculate thermodynamic properties of spin rings and chains which can be modeled by an anisotropic quantum Heisenberg model. In particular we use these codes to simulate molecular nanomagnets belonging to the Cr_8 family [13-15, 24, 25].

However, in our research the ED-based codes have also been executed, depending on the circumstances as the appropriate sizes of the molecules or the type of physical properties pursued [15, 23].

In our QTM applications each diagonal element of the final transfer matrix is calculated separately and requires approximately the same number of arithmetic operations. The number of elements calculated by a give core was defined by the ratio of the total number of basis vectors and the number of cores involved. The communication was based on the MPI directives and took place only at the beginning and end of the processes running in each core. The only limiting factors to the efficiency are the uniform distribution of the work among the computing cores and the size of the matrix. This size should not be too small, to avoid an impact of the latencies not directly related to the parallel implementation. Our application 1 was tested imposing $N = 11$, whereas it was designed to solve the problem of $N = 9$ interacting spins $S = 3/2$ in the Cr_9 molecules.

The size of the computational problem is determined by the number L of the diagonal matrix elements which is equal

Tab. 1. The hardware characteristics and the data showing the scalability of applications 1 and 2. The notation of columns is the following: N_r – number of racks, N_c – number of cores, n_i – the range of matrix elements assigned to a single core (the index $i = 1, 2$ and refers to a given application), t_i – the total running time in seconds, e_i – the corresponding efficiency.

N_r	N_c	n_1	n_2	t_1[s]	t_2[s]	e_1[%]	e_2[%]
8	32768	128	91-92	7428	804	100	99.0
16	65536	64		3713		100	
32	131072	32		1857		100	
64	262144	16		929		99.9	
72	294912	14-15	10-11	870	104	94.9	85.9

to the size of the transfer matrix. Thus the applications 1 and 2 have dealt with the significant computational problems of the sizes $L = 4^{11} = 4194304$ and $L = 12^6 = 2985984$, respectively. Of course such large matrices would not fit into memory limits for BlueGene/P systems. The memory footprint was reduced to the values of the order of L by exploiting the sparseness of the matrices. In fact only a few vectors of size L are needed to be remembered. Furthermore, using the same mechanisms the computational complexity for calculating a single element was also reduced from matrix-vector multiplications to vector-vector operations leading to the temporal complexity of the order of L^2 due to a small m value adopted. The magnitude of the integer m depends on the convergent rate. The slower convergence, the higher value of m is needed. However, for the sake of testing the code parallelization, the size of m is irrelevant.

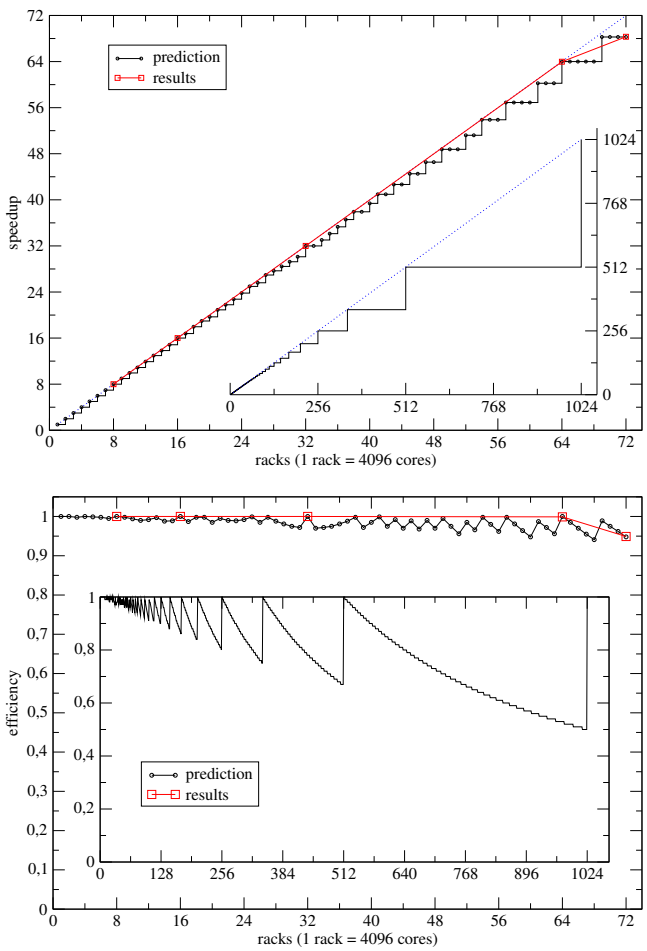

Fig. 3. Scaling plots for application 1: in the upper panel, speedup as a function of a number of racks; in the lower panel, the corresponding efficiency. Theoretical predictions are plotted by full red lines and match very well the computer-based data denoted by symbols. The insets depict the expected values calculated up to 1024 racks (4 194 304 cores) – the highest fully efficient number of cores for this problem

III. ANALYSIS OF CODE PERFORMANCE AND RESULTS

We have performed a number of successful large scale runs, proving an excellent scalability of both applications. We note that on JUGENE only full racks containing processors with 2^{12} cores could be assigned during execution. Application 1 was tested more thoroughly. Runs were performed from 8 racks up to the full number of 72 racks (see Table 1). In Table 1 we have assumed the following abbreviations: N_r means the number of racks assigned during a given run; N_c denotes the total number of cores engaged; n_i shows the range of the number of elements processed by a single core in the case of the ith application, where $i = 1, 2$; t_i describes the total processing time and e_i yields the efficiency of a given parallelized code. Some fields in Table 1 are missing, as the application 2 was not granted so much computing time as the first one.

As one element of the matrix was calculated during a fixed amount of time, the more cores computed the whole numerical task, the smaller part of the task each slave core had to calculate on its own. Knowing that, it is possible to predict theoretical speed up and efficiency as a function of the number of computing racks or cores. This was done for Application 1 and is presented in Fig. 3. This application would maintain perfect linear scalability when the work distribution was optimal, i.e. the ratio of the total number of matrix elements L and the number of cores N_c was an integer.

The performance of application 1 is illustrated in Fig. 3. The full red straight lines in both panels illustrate the optimal theoretical speed up, whereas the black curves in the form of steps show the values obtained from the real computing runs. As expected, for the optimal distribution of slave processors, the ideal efficiency is accomplished. These cases are marked by the full red dots. In the insets we demonstrate predictions of the speedup and efficiency for a computer with as many as 1024 racks.

The corresponding predictions for application 2 are visualized in Fig. 4 by the blue curves. The only two experimental points available are denoted by the red symbols. That obtained from the 8-rack run fits the theoretical predictions quite well, but that from the full-machine run clearly deviates. Although at this size the efficiency should be around 92%, we got 86% only. The inspection of worklogs showed that in both cases the time required to compute one element did not vary much, so the total time should be very close to the predicted one. We believe that the loss of efficiency might be caused by some intrinsic problems with final MPI reduction operation performed on the whole machine, as this is the only place between the single-step and total run-time measurements which could affect the performance.

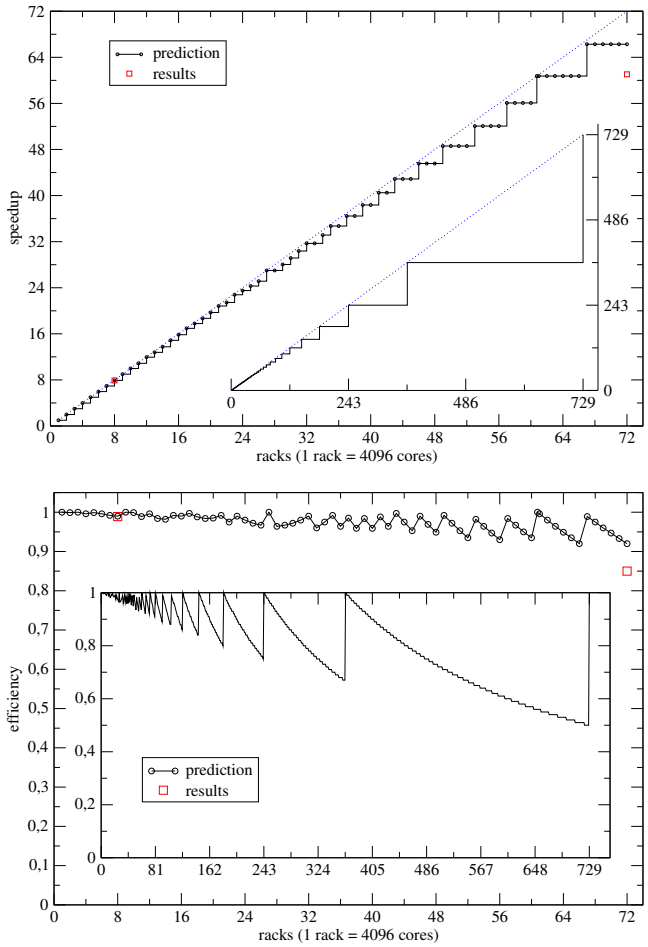

Fig. 4. The counterpart of the scaling plots drawn for application 2. The 8-rack run result depicted by red symbol fits the predictions plotted by the blue lines but the full-machine run data point indicates an underperformed behavior of the code. The insets depict the expected values calculated up to 729 racks (2 985 984 cores) – the highest fully efficient number of cores for this problem

IV. DISCUSSION AND CONCLUSIONS

The QTM technique is a powerful tool not only for simulations of ring-shape molecular nanomagnets but also for low-dimensional spin chains, both macroscopic and partitioned into finite-size segments [26, 27]. Its advantages comprise the computational aspects (such as significant reduction of the temporal complexity of the implementation) and a wide domain of applications. Within QTM it is possible to simulate the low symmetry models with anisotropic and non-uniform interactions which are hardly accessible for the ED approaches.

We have proved that our QTM codes can be scaled with the number of cores or racks over the whole JUGENE supercomputer resources. We have reached a nearly perfect scalability for application 1. The theoretical predictions matched the actual running time for optimally tuned cases. In the case of application 2, unexpected reduction of efficiency emerged for the full 72-rack run which was attributed to some intrinsic MPI performance rather than to a failure of the parallelized code.

For the purpose of testing the scalability, we have considered a simplified model of the physical problem in order to shorten the calculations of a single matrix element. To this end the number m of replicas was kept very small. To obtain accurate real physical characteristics, the constraint on the m value should be relaxed and the applications should be invoked for a number of varying parameters so that the total running time would increase significantly.

As proper implementation of the QTM method is scalable over to the biggest machines currently available, this feature predisposes our tool for simulation of relatively large molecular magnetic complexes which are the subject of current experimental research.

Acknowledgments

We thank Dr. P. Kozłowski for providing us with the code of his application (referred to as 1) which was subject to porting to the JUGENE environment. A partial support from the NCN grant Nr 579138 is also acknowledged. The computations were performed on JUGENE at the Juelich Supercomputing Center during the BlueGene/P Extreme Scaling Workshop in February 2011.

References

[1] D. Gatteschi, R. Sessoli, and J. Villain, *Molecular nanomagnets*, Oxford University Press, Oxford, 2006.

[2] M. Mannini, F. Pineider, P. Sainctavit, C. Danieli, E. Otero, C. Sciancalepore, A.M. Talarico, M-A. Arrio, A. Cornia, D. Gatteschi, and R. Sessoli, *Magnetic memory of a single-molecule quantum magnet wired to a gold surface*, Nature Mater. **8**, 194-197 (2009).

[3] A. Ardavan, O. Rival, J.J.L. Morton, S.J. Blundell, A.M. Tyryshkin, G.A. Timco, and R.E.P. Winpenny, *Will spin-relaxation times in molecular magnets permit quantum information processing?*, Phys. Rev. Lett. **98**, 057201 (2007).

[4] J. Lehmann, A. Gaita-Ariño, E. Coronado, and D. Loss, *Spin qubits with electrically gated polyoxometalate molecules*, Nature Nano. **2**, 312-317 (2007).

[5] G.A. Timco, S. Carretta, F. Troiani, F. Tuna, R.J. Pritchard, Ch. A. Muryn, E.J.L. McInnes, A. Ghirri, A. Candini, P. Santini, G. Amoretti, M. Affronte, and R.E.P. Winpenny, *Engineering the coupling between molecular spin qubits by coordination chemistry*, Nature Nano. **4**, 173-178 (2009).

[6] B. Georgeot and F. Mila, *Chirality of triangular antiferromagnetic clusters as qubit*, Phys. Rev. Lett. **104**, 200502 (2010).

[7] D. Gatteschi, R. Sessoli, and A. Cornia, *Single-molecule magnets based on iron (III) oxo clusters*, Chem. Commun., pages 725-732 (2000).

[8] S.G. Louie, *Nanoparticles behaving oddly*, Nature **384**(6610), 612-613 (1996), 10.1038/384612a0.

[9] O. Cador, D. Gatteschi, R. Sessoli, A-L. Barra, G.A. Timco, and R.E.P. Winpenny, *Spin frustration effects in an odd-membered antiferromagnetic ring and the magnetic Möbius strip*, J. Magn. Magn. Mater. **290-291**, 55-60 (2005).

[10] Y. Furukawa, K. Kiuchi, K. Kumagai, Y. Ajiro, Y. Narumi, M. Iwaki, K. Kindo, A. Bianchi, S. Carretta, P. Santini, F. Borsa, G.A. Timco, and R.E.P. Winpenny, *Evidence of spin singlet ground state in the frustrated antiferromagnetic ring Cr_8Ni*, Phys. Rev. B **79**, 134416 (2009).

[11] K. Bärwinkel, H.J. Schmidt, and J. Schnack, *Ground-state properties of antiferromagnetic Heisenberg spin rings*, J. Magn. Magn. Mater. **220**, 227-234 (2000).

[12] J. Schnack, *Properties of the first excited state of nonbipartite Heisenberg spin rings*, Phys. Rev. B **62**, 14855-14859 (2000).

[13] P. Kozłowski, M. Antkowiak, and G. Kamieniarz, *Frustration signatures in the anisotropic model of a nine-spin $s = 3/2$ ring with bond defect*, J. Nanopart. Res. (2011).

[14] M.L. Baker, G.A. Timco, S. Piligkos, J.S. Mathieson, H. Mutka, F. Tuna, P. Kozłowski, M. Antkowiak, T. Guidi, T Gupta, H. Rath, R.J. Woolfson, G. Kamieniarz, R.G. Pritchard, H. Weihe, L. Cronin, G. Rajaraman, D. Collison, E.J.L. McInnes, and R.E.P. Winpenny, *Spin frustration in molecular magnets - a classification: physical studies of large odd-numbered-metal, odd-electron rings*, Proc. Natl. Acad. Sci. USA**109**(47), 19113-19118 (2012).

[15] M. Antkowiak, P. Kozłowski, G. Kamieniarz, G.A. Timco, F. Tuna, and R.E.P. Winpenny, *Detection of ground states in frustrated molecular rings by the in-field local magnetization profiles*, Phys. Rev. B **87**, 184430 (2013).

[16] A. Caramico D'Auria, U. Esposito, F. Esposito, G. Kamieniarz, and R. Matysiak, *Exact simulations of quantum rings and characterization of hexanuclear manganese and dodecanuclear cyclic complexes*, J. Phys.: Condens. Matter **13**, 2017 (2001).

[17] G. Kamieniarz, R. Matysiak, A. Caramico D'Auria, F. Esposito, and C. Benelli, *Finite-temperature characterization and simulations of the molecular assemblies Mn_6 and Ni_{12}*, Eur. Phys. J. B **23**, 183 (2001).

[18] D. Gatteschi, A. Caneschi, L. Pardi, and R. Sessoli, *Large clusters of metal ions: The transition from molecular to bulk magnets*, Science **265**(5175), 1054 (1994).

[19] A. Caneschi, D. Gatteschi, C. Sangregorio, R. Sessoli, L. Sorace, A. Cornia, MA Novak, C. Paulsen, and W. Wernsdorfer, *The molecular approach to nanoscale magnetism*, Journal of magnetism and magnetic materials **200**(1-3), 182-201, (1999).

[20] G. Kamieniarz and R. Matysiak, *Simulations of the low-dimensional magnetic systems by the quantum transfer-matrix technique*, Computational Materials Science **28**(2), 353-365 (2003), Proceedings of the Symposium on Software Development for Process and Materials Design.

[21] G. Kamieniarz and R. Matysiak, *Transfer matrix simulation technique: effectiveness and applicability to the low-dimensional magnetic spin systems*, J. Comput. Appl. Math. **189**, 471-480 (May 2006).

[22] G. Kamieniarz, P. Kozłowski, M. Antkowiak, P. Sobczak, T. Ślusarski, D.M. Tomecka, A. Barasiński, B. Brzostowski, A. Drzewiński, A. Bieńko, and J. Mroziński, *Anisotropy, Geometric Structure and Frustration Effects in Molecule-Based Nanomagnets*, Acta Physica Polonica A**121**(5-6), 992-998 (2011).

[23] G. Kamieniarz, W. Florek, and M. Antkowiak, *Universal sequence of ground states validating the classification of frustration in antiferromagnetic rings with a single bond defect*, Phys. Rev. B **92**, 140411(R) (2015).

[24] G. Kamieniarz, P. Kozłowski, G. Musiał, W. Florek, M. Antkowiak, M. Haglauer, A.C. D'Auria, and F. Esposito, *Phenomenological modeling of molecular-based rings beyond the strong exchange limit: Bond alternation and single-ion anisotropy effects*, Inorg. Chim. Acta **361**, 3690-3696 (2008).

[25] P. Kozłowski, G. Kamieniarz, M. Antkowiak, F. Tuna, G.A. Timco, and R.E.P. Winpenny, *Phenomenological modeling of the anisotropic molecular-based ring Cr_7Cd*, Polyhedron **28**, 1852-1855 (2009).

[26] R. Matysiak, G. Kamieniarz, P. Gegenwart, and A. Ochiai, *Field-dependent specific heat of Yb_4As_3: Agreement between a spin-1/2 model and experiment*, Phys. Rev. B **79**, 224413 (2009).

[27] R. Matysiak, P. Gegenwart, A. Ochiai, M. Antkowiak, G. Kamieniarz, and F. Steglich, *Specific heat of segmented Heisenberg quantum spin chains in $(Yb_{1-x}Lu_x)_4As_3$*, Phys. Rev. B **88**, 224414 (2013).

Michał Antkowiak received his MSc degree in computer science at the Umeå University, Sweden, in 2006. In the same year he received MSc degree in applied computer science at the Adam Mickiewicz University in Poznań, Poland. He defended his PhD in physics at the same university in 2013. His research interests include modelling of molecular nanomagnets and parallel processing.

Łukasz Kucharski is a PhD student at the Faculty of Physics, Adam Mickiewicz University in Poznań. He had the opportunity to work on the biggest supercomputers in Europe and is an author of the software dedicated to the parallel quantum chemistry calculations using genetic algorithms.

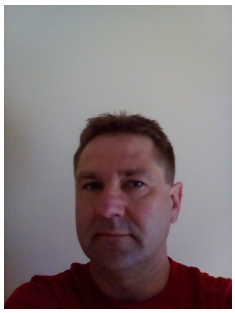

Ryszard Matysiak received the PhD degree in Physics Science at the Adam Mickiewicz University in Poznań, in 1999. His field of interest covers statistical mechanics and simulations of the low-dimensional magnetic systems.

Grzegorz Kamieniarz is a professor in theoretical and computational physics at the Institute of Physics of the Adam Mickiewicz University in Poznań. His field of interest covers statistical mechanics, phase transitions, phenomenological modeling and simulations of the low-dimensional magnetic systems including molecular nanomagnets.