



HAL
open science

An Analysis of the Feasibility of Energy Harvesting with Thermoelectric Generators on Petascale and Exascale Systems

Issam Raïs, Laurent Lefèvre, Anne-Cécile Orgerie, Anne Benoit

► To cite this version:

Issam Raïs, Laurent Lefèvre, Anne-Cécile Orgerie, Anne Benoit. An Analysis of the Feasibility of Energy Harvesting with Thermoelectric Generators on Petascale and Exascale Systems. Workshop Optimization of Energy Efficient HPC & Distributed Systems (OPTIM) in conjunction with HPCS, Jul 2016, Innsbruck, Austria. hal-01348554

HAL Id: hal-01348554

<https://inria.hal.science/hal-01348554>

Submitted on 9 Sep 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

An analysis of the feasibility of energy harvesting with thermoelectric generators on petascale and exascale systems

Issam Rais

Univ Lyon - Inria - LIP Laboratory - CNRS -
ENS de Lyon - Inria - UCB, France
{issam.rais}@inria.fr

Anne-Cecile Orgerie

CNRS, IRISA, Rennes, France
{anne-cecile.orkerie}@irisa.fr

Laurent Lefevre

Univ Lyon - Inria - LIP Laboratory - CNRS -
ENS de Lyon - Inria - UCB, France
{laurent.lefevre}@inria.fr

Anne Benoit

Univ Lyon - Inria - LIP Laboratory - CNRS -
ENS de Lyon - Inria - UCB, France
{anne.benoit}@inria.fr

Abstract—The heat induced by computing resources is generally a waste of energy in supercomputers. This is especially true in very large scale supercomputers, where the produced heat has to be compensated with expensive and energy consuming cooling systems. Energy is a critical point for future supercomputing trends that currently try to achieve exascale, without having its energy consumption reaching an important fraction of a nuclear power plant. Thus, new ways of generating or recovering energy have to be explored. Energy harvesting consists in recovering wasted energy. ThermoElectric Generators (TEGs) aim to recover energy by converting wasted dissipated energy into usable electricity. By combining computing units (CU) and TEGs at very large scale, we spotted a potential way to recover energy from wasted heat generated by computations on supercomputers.

In this paper, we study the potential gains in combining TEGs with computational units at petascale and exascale. We present the technology behind TEGs, the study of a typical supercomputer environment, and finally our results concerning binding TEGs and computational units in a petascale and exascale system. With the available technology, we demonstrate that the use of TEGs in a supercomputer environment could be realistic and quickly profitable, and hence have a positive environmental impact.

I. INTRODUCTION

Energy efficiency is nowadays a critical issue in high-performance computing [4]. There are two ways to obtain energy: production or harvesting.

Energy could be produced from transformation of raw materials (oil, gas, coal, petrol, etc). Every produced kiloWatt is accompanied with a generation of CO_2 [16]. Because energy generators are not 100% efficient, a significant part of energy is lost during the production and transport [6]. Only a slight percentage of the conversion of raw materials is transformed into electricity. For example, the ratio between useful energy output and energy input (as known as energy conversion efficiency) of a coal plant is 33-35% [5].

New ways to generate or at least recycle energy have to be explored. Every recycled Watt is a non produced Watt (thus

avoiding CO_2 generation). Energy harvesting is the process of extracting energy from surrounding environment [4]. Even if this recovered energy concerns few joules, it is recycled from a previous consumption, and it is thus increasing the energy efficiency.

Several energy harvesting solutions exist. Among them, photovoltaics is the one in vogue: it aims at transforming solar beams into electricity [13]. Piezoelectric generators are built to recycle pressure expressed on a component and generate an amount of energy from the distortion of the component [12]. Thermoelectricity aims at recycling wasted energy under heat form. Under a delta temperature, a thermoelectric generator converts lost energy from heat to electricity form [21]. Pyroelectricity follows the same principle but at higher delta temperature [15].

The growing demand for processing data and computing implies improvements of electronic components. At the level of computing capabilities or energy consumption, computational components of current machines are becoming more and more efficient. Despite technological and software breakthrough to increase computational power while reducing energy consumption, it became a real limiting factor for computers. This is especially true for supercomputers reaching petascale performance and aiming for exascale. Nowadays, sustained one megaWatt per year costs around one million dollars per year for a datacenter usage [8].

A great part of consumed energy is lost in an exothermic way. For the safety of data centers, this lost energy must be carried away. To do so, expensive air or water cooling systems are built upon data centers or supercomputers. None of this lost energy is reversed in a direct usable electricity for a direct use. Thermoelectricity seems to be the proper energy harvester for this problem.

Binding CPUs and TEGs has been mentioned in [24]. It deals with common computers built with only one computing

unit (CU), and it experiments only around the idle part of a CPU. Not stressing the CU implies lowering the delta temperature. Thus, the CPU is not under deep stress like it would be in an intensive supercomputer scenario.

This paper discusses the feasibility of using thermoelectric generators for petascale and future exascale machines. How much harvested energy can we hope for? Is recovered energy sufficient enough to be directly injected in CPUs or to be stored? How much time will it take to compensate the cost of combining thermoelectric generators and computational units?

This paper is structured as follows. Section II briefly presents the TEG technology. The supercomputer environment is introduced in Section III. Section IV presents a hypothetical application of TEG on hypothetical exascale machines. A hypothetical application of TEG on an existing machine, Tianhe-2 is presented in Section V. Finally, Section VI concludes this work.

II. THERMOELECTRIC GENERATORS (TEGs)

Thermo-electricity was discovered during the 19th century by Thomas Johann Seebeck. It is a physical phenomenon where a thermal delta seen through a component results in an electrical current [23]. A TEG can be used as a generator (when it is exposed to a delta temperature) or as a cooler (when a current passes through it). In this paper, we focus on the generator feature.

The actual increase concerns about climate and energy efficiency force researchers to consider all possible solutions to consume better. One of the main concerns is energy lost in an exothermic way. Indeed, this energy enforces expensive costs in cooling systems. By transforming a part of this dissipated energy, TEGs may reduce costs.

A thermo electrical material transforms a thermal delta in electricity. TEGs are composed by positively (p-type) and negatively (n-type) doped connected semiconductor couples (see Fig. 1). N-P couples are the charge carriers that can freely move through the metal. These carriers start to move under a delta temperature, according to the N-P couple properties. The delta temperature creates an excitation of the doped charge




Fig. 1: ThermoElectric Generator (TEG)

carrier, thus it creates a movement of the charge carrier, creating an electric current.

Ideally, a larger delta temperature creates a larger electrical current, but this statement is highly coupled with the fact that a semiconductor is effective only on a range of temperature, making the TEGs operational only on a limited delta temperature [21].

For years now, TEGs are used on spacial scientific programs like Voyager. TEGs are used to provide energy in space where solar panels are too far from the sun to generate energy. TEGs are not auto sufficient. Heat generated by plutonium isotope and cold space temperature (2 or 3 degrees above absolute zero) create a suitable environment for TEG power generation [1].

Other fields can have potential use of such a technique. A lot of heat is wasted from automotive, specially from exhaust pipes. A lot of exhaust pipes accompanied with TEG prototypes have emerged. Only 25 percent of the applied fuel energy made from combustion is used for mobility and accessories. About 40 percent of this energy is lost in exhaust pipe. Even if conversion percentage is low, it could lead to a nice ratio of CO_2 emission reduction [11].

Ultra-low power wireless sensors used for various measurements can also be combined with energy harvesting techniques. Thanks to the energy harvesting part of the TEGs and the ultra-low needs of power, it is nowadays possible to produce wireless sensors that are almost autonomous systems [18].

A TEG can be composed from different alloys. Every alloy has a temperature where it could be applied. The possible range of action on a TEG depends on a minimal and a maximal temperature. In fact, under this minimum temperature, no power can be generated and over this maximum temperature, the chosen TEG has an indeterminate behavior [21].

III. SUPERCOMPUTER ENVIRONMENT

Because of the intensity of usage, in the case of a supercomputer or a data-center, computational units (CUs) could be seen as generators of heat. We first detail the supercomputing thermal environment in Section III-A, and then we discuss TEG efficiency in Section III-B.

A. Supercomputing thermal environment

For the safety of the machine composing the supercomputer, the heat generated by CUs must be carried away from CUs. This results in a huge amount of wasted energy, when taking into account the dissipated energy from computation and from the expensive cooling system.

Generated energy from a TEG depends on the temperature on the cold and hot side. Thus, it depends on the delta temperature between the two parallel sides. Table I presents the extreme thermal context that could be witnessed on a supercomputer environment.

We can assume that, because of the intensity of the usage, every CU is close to the maximum temperature.

TABLE I: Supercomputer extreme delta temperature in Celsius degree

Characteristic	Temperature
hotspot	95 - 110
cooling system	15 - 25
delta	95

The hotspot informations correspond to the maximum reachable temperature by a Xeon CPU and a Nvidia GTX Titan Z GPU. The cooling system temperature represents the recommended range of temperature expected in a room containing datacenter components [10]. The delta is the difference between these two extreme temperatures. We make the hypothesis that we put a TEG on every CU and that the temperature on top of this TEG is equal to the cooling system temperature. In other words, we consider that the cooling system is on top of every CU (like with water cooling system), thus creating the needed delta temperature for a TEG to generate electricity.

B. TEG maximum efficiency

The maximum efficiency, noted η_{max} , represents the maximum percentage that a TEG can convert from a thermal Watt to an electric Watt. The formula of η_{max} is described as follows:

$$\eta_{max} = \frac{\Delta T}{T_h} \frac{\sqrt{1 + ZT} - 1}{\sqrt{1 + ZT} + \frac{T_c}{T_h}}, \quad (1)$$

where

- T_c is the temperature on the cold side;
- T_h is the temperature on the hot side;
- ΔT is the difference between the hot and cold side, i.e., $\Delta T = T_h - T_c$;
- ZT is the figure of merit of the TEG.

The maximum efficiency of TEG is therefore governed by the figure of merit ZT , that is a parameter depending on the materials used to build the TEG, and by temperature of the hot and cold side of the TEG. It also depends on the Carnot efficiency (the left factor on Equation (1), i.e., $\Delta T/T_h$). This metric is needed to estimate the maximum usable power generated from a TEG on a specific environment.

Applying this metric to the previously described environment (Table I) allows us to obtain the maximum efficiency reference, which depends on the ZT of the chosen TEG.

Fig. 2 presents the evolution of the maximum efficiency, η_{max} , as a function of the ZT of TEGs with $T_c = 15$ Celsius degree, for three delta temperatures ($\Delta T = 50, 75, 100$). We consider a ZT between 0 and 5, given that the state of the art TEG figure of merit is nowadays around 2.2 [3]. In a very optimistic setting with $ZT = 5$ and $\Delta T = 100$, the maximum efficiency is around 12%.

The most widespread alloy for thermal exchanges between computing units and TEGs is the Bismuth Telluride (Bi_2Te_3) [9]. From [21], we can understand that this is the

only alloy that can be applied in an environment as described in Table I. In fact, in a range between 15 and 100 Celsius degree, only one N-P couple can be built to be used in a TEG (respectively Bi_2Te_3 and Sb_2Te_3). A TEG with these alloys has a device figure of merit ZT between 0.7 to 1.0 [17].

Considering Fig. 2 and the known ZT of a compatible TEG, we can now say that the maximum efficiency in our environment will be between 3 and 5%.

IV. TEGS ON EXASCALE INFRASTRUCTURES

In order to figure out if the usage of TEGs is pertinent in a petascale or exascale environment, we consider several scenarios.

A. Theoretical setup

For this analysis, we consider the state of art and potential CU candidates to an exascale machine. In this paper, a CU represents a computing node that dissipates a lot of heat while computing (CPU, GPU, MIC, etc).

Table II presents the studied computing units. We consider the in vogue KNL XeonPhi [19], the very promising RexNeo¹, and the mainstream Xeon CPU [14], and we determine how many megaWatts a machine reaching one exaFLOPS composed only with one of these computing units will consume (see Table III).

TABLE II: Studied computing units - peak in TFLOPS and consumption in Watts

Computing unit	peak	consumption
XeonPhi (KNL)	3	200
RexNeo	0.256	4
Xeon	0.506	145

Following the Department of Energy (DOE) requirements, the expected exascale machine must achieve one exaFLOPS, with a consumption of 20 MW (or below) [20]. The machines

¹<http://www.rexcomputing.com/>




Fig. 2: Maximum TEG efficiency in the supercomputer scenario

TABLE III: Hypothetical wattage and number of computing units (in million units) for an exascale machine

Machine	MegaWatts	# CU
100% RexNeo	15	3.9
100% XeonPhi	66	0.3
100% Xeon	286	1.9

in Table III are the hypothetical machines made with only one of the computing units presented in Table II.

The first interesting observation is the fact that an exascale machine made only with RexNeo computing units can achieve one ExaFLOPS with 15 MW, which is under the recommended limit fixed by the DOE. The second observation is that a machine reaching one ExaFLOPS made only with KNL XeonPhi or Xeon is unconceivable when considering power consumption and the required limit imposed by the DOE. However, because these CUs are very common in the supercomputer field, we will keep them in the study.

B. Recovered energy and return on investment

The followed hypothesis is to put a TEG on every computational unit and to estimate how much energy would be recovered with the previous described environments (Table I) extended to these exascale machines (Table III).

Table IV shows the hypothetical gain with a TEG with an efficiency η_{max} between 3 and 5% on every CU on the hypothetical machines that we chose for our exascale machines. In addition to the realistic values, we study the best hypothetical efficiency (ZT= 5, η_{max} around 12 % on Figure 2) in the supercomputer case. Even if the percentage is low, the energy recovered is equivalent to a non negligible computing unit number for every exascale machine.

Such an installation is beneficial only if the price of the deployment of TEGs is quite accessible and if the return on investment is under the lifetime of the supercomputer on top of which it was built.

In Table V, we display the price of applying TEGs on every CU. Prices are from websites for common usage and TEG vendors (\$4² and \$10³), so we could imagine that the cost for such a machine, per TEG, will be way lower. As a lower bound, we consider an extreme case where a TEG costs \$1. Even with common usage prices, applying a TEG on every CU stays reasonable, considering the price of a supercomputer such as Tianhe-2 as a reference. Note however that this cost only integrates the raw cost of TEG, without considering manpower and other deployment costs.

The return on investment is an important parameter to consider. In fact, as motivated earlier, TEGs offer a way to recover wasted energy. Therefore, this recovered energy is an amount of energy that will not be asked by the machine, thus it does not have to be produced elsewhere and bought. To compute the return on investment, we consider a new metric

$yReturn$: it represents the number of years after which having installed TEGs on all CUs becomes profitable, and it is defined as:

$$yReturn = \frac{TEGsOverallCost}{SustainedWattCostOneYear \times recoveredEnergy},$$

hence by dividing the overall cost of TEGs by the cost saved within one year thanks to the generated energy. In fact, the TEG costs will be amortized on the recovered energy costs. This generated energy could be directly used in the supercomputer or sold back to the provider, and it corresponds to an amount of energy that would have been asked to the provider. Recovered energy becomes a direct benefit past $yReturn$ years.

Based on the hypothetical exascale machines described in Table III, the cost of having a TEG on every CUs of this machine (as described in Table V), and the cost of one sustained Watt per year (fixed at \$1 in our case), Table VI shows the return on investment of the different scenarios (i.e under various TEG cost and η_{max}).

Because of the low FLOPS of RexNeo, a very high number of CUs are needed to achieve one exaFLOPS. Thus, the return on investment for this machine is drastically high, overcoming the usual life time of a supercomputer. Except for this specific configuration, other hypothetical exascale machines benefit from TEGs between 8 months and 1 year and 4 months.

One can ask if the actual return on investment does not overcome the life time of a TEG. In fact, TEGs could last more than 15 years [2], so there should be no need to replace them during the lifetime of the supercomputer. In comparison, the return on investment of an exascale machine made only with XeonPhi or Xeon computing unit is very low.

V. APPLICATION ON AN EXISTING MACHINE

To see the potential impact of TEGs on a real architecture, we chose to lean on the number one machine on the top500, Tianhe-2 [22]. For a few years now, Tianhe-2 has been the number one machine in terms of peak performance.

TABLE IV: Energy recovered in MegaWatts and equivalent computing unit consumption compensation

Machine	3% (MW)	5% (MW)	12 % (MW)
100% RexNeo	0.47	0.78	1.87
100% XeonPhi	2	3.3	8.0
100% Xeon	8.5	14	34.38
Machine	3% (#CU)	5% (#CU)	12 % (#CU)
100% RexNeo	117 188	195 313	468 750
100% XeonPhi	10 000	16 667	40 000
100% Xeon	59 289	98 814	237 154

TABLE V: TEG on CU - overall cost in million dollars

Machine	TEG \$4	TEG \$10	TEG \$1
100% RexNeo	15.6	39	3.9
100% XeonPhi	1.3	3.3	0.3
100% Xeon	7.9	19.7	1.97

²\$4 TEG example

³\$10 TEG example

TABLE VI: Return on investment in years in function of TEG cost and η_{max}

Machine	\$4, 3%	\$10, 5%	1\$,12%
100% RexNeo	33.3	50	2
100% XeonPhi	0.6	1	0.04
100% Xeon	0.9	1.3	0.05

A. Tianhe2 - Architecture and energy consumption

As exposed in [7], every CU of Tianhe-2 is composed of two Xeon processors and three Xeon Phi accelerators. There are 16000 computational nodes accompanied with 4096 front end nodes. It costs around 2.4 billion yuan, i.e., around 390 million US\$ [22].

TABLE VII: Tianhe-2 configuration

Chipset	TDP	Number
Xeon E5-2690	115	32000
Intel Xeon Phi 31S1P	270	48000
Galaxy FT-1500	65	4096

Table VII exposes the TDP (Thermal Dissipation Power) and quantity of the CUs composing Tianhe-2.

As shown in Table VII, Tianhe-2 is mainly composed of Xeon E5 processors along with Xeon Phi accelerators. The Galaxy FT-1500 is the main front end nodes CPU component of Tianhe-2.

TABLE VIII: Tianhe-2 power consumption in MegaWatts

	Consumption
Total facility energy	24
IT equipment energy	17

Table VIII presents the total facility energy and IT equipment energy exposed consumption of Tianhe-2 in [7]. The Power usage effectiveness (PUE) of such an infrastructure is around 1.41.

B. Tianhe-2 - Potential gain of TEGs, cost and return on investment

As explained in Section III, the maximum efficiency in a supercomputer environment and with current TEGs technology is estimated to be between 3 and 5%, and the most optimistic efficiency is around 12%.

Even if η_{max} is equal to 3% of 5%, Table IX shows that the energy recovered is equivalent to a non negligible number of CUs. Such an installation could be beneficial if and only if the return on benefit is inferior to the supercomputer lifetime.

In Table X, we made the assumption that TEGs recover energy from every CU dissipating heat.

Table XI presents how much time, in years, it would take to recover the cost of one TEG on every CU for Tianhe-2. This table presents various TEG configuration results. The worst case concerns the front end node with a TEG that costs \$10. If

TABLE IX: Energy recovered in MegaWatts and equivalent computing unit consumption compensation in the Tianhe-2 case

Machine	3% (MW)	5% (MW)	12 % (MW)
XeonPhi	0.38	0.648	1.55
Xeon	0.11	0.184	0.44
Galaxy	0.0079	0.0133	0.032

Machine	3% (#CU)	5% (#CU)	12 % (# CU)
XeonPhi	960	1600	5760
Xeon	1440	2400	3840
Galaxy	122.88	204,8	491.2

TABLE X: TEG on all Tianhe-2 CUs - cost

	TEG \$4	TEG \$10	TEG \$1
XeonPhi	128000	320000	32000
Xeon	192000	480000	48000
Galaxy	16384	40960	4096
Total	336384	840960	84096

TABLE XI: TEG on all Tianhe-2 CUs - Profitability (years)

TEG \$10			
	3%	5%	12%
XeonPhi	2,89	1,73	0,72
Xeon	1,23	0,74	0,30
Galaxy	5,12	3,07	1,28
TEG \$4			
	3%	5%	12%
XeonPhi	1,15	0,69	0,28
Xeon	0,49	0,29	0,12
Galaxy	2,05	1,23	0,51
TEG \$1			
	3%	5%	12%
XeonPhi	0,28	0,17	0,07
Xeon	0,12	0,07	0,03
Galaxy	0,51	0,30	0,12

we consider only CUs, the worst case is still the 3% TEG that costs \$10, which will become profitable after approximately 3 years of usage.

Tianhe-1A is 6 years old, Titan is 4 years old, while Tianhe-2 is 3 years old. Even in the worst case, the time in which our hypothetical TEG deployment on Tianhe-2 CUs will become profitable is way under the usage of a supercomputer. In fact, as shown in the top500 list, a supercomputer usage could last a few decades.

It is interesting to note that with a reasonable and realistic TEG cost and a recycling efficiency, it becomes profitable, concerning CUs, only after 0.69 year, approximately 9 months (for example a \$4 TEG with an efficiency of 5%).

VI. CONCLUSION

Energy efficiency is nowadays a critical issue in high-performance computing. At the petascale and exascale level,

it will become an even more critical issue. Every potential percentage gained will have a huge impact on the overall consumption. Energy harvesting techniques are digging the idea of recovering and transforming wasted energy from different unusable to usable forms, hence making this technique a potential solution to recover wasted dissipated energy.

In this paper, we show that the usage of TEGs in a supercomputer environment could be realistic and quickly profitable, even if it is not the best environment to exploit the potential of TEGs. In fact, such an environment implies a low delta temperature between hot and cold side of the TEG, a low hot temperature on the computing unit, and thus a low maximum efficiency. However, theoretical applications show that even though only a small fraction of the heat could be transformed into energy, the return on investment could be short, as shown on the application on the number one supercomputer of the top500 list, Tianhe-2. Thus, at the petascale or exascale level, where every hypothetical energy gain is important, it could be a new way to recover energy. We present the feasibility of energy harvesting on a petascale and exascale environment using ThermoElectric Generators (TEGs).

VII. ACKNOWLEDGMENTS

This work is integrated and supported by the ELICI project, a French FSN ("Fond pour la Société Numérique ") project that associates academic and industrial partners to design and provide software environment for very high performance computing.

REFERENCES

- [1] G. L. Bennett, J. J. Lombardo, R. Hemler, and J. Peterson. The General-Purpose Heat Source Radioisotope Thermoelectric Generator: Power for the Galileo and Ulysses Missions. In *Proceedings of the 21st intersociety energy conversion engineering conference*, 1986.
- [2] G. L. Bennett, J. J. Lombardo, R. J. Hemler, G. Silverman, C. Whitmore, W. R. Amos, E. Johnson, A. Schock, R. W. Zocher, T. K. Keenan, et al. Mission of daring: the general-purpose heat source radioisotope thermoelectric generator. In *4th International Energy Conversion Engineering Conference and Exhibit (IECEC)*, San Diego, California, AIAA, volume 4096, page 2006, 2006.
- [3] K. Biswas, J. He, I. D. Blum, C.-I. Wu, T. P. Hogan, D. N. Seidman, V. P. Dravid, and M. G. Kanatzidis. High-performance bulk thermoelectrics with all-scale hierarchical architectures. *Nature*, 489(7416):414–418, 2012.
- [4] S. Chalasani and J. M. Conrad. A survey of energy harvesting sources for embedded systems. In *Southeastcon, 2008. IEEE*, pages 442–447. IEEE, 2008.
- [5] B. Clark, W. Graves, J. Lopez-de Cardenas, M. Gurfinkel, and A. Peats. Working Document of the NPC Global Oil and Gas Study. *Topic Paper*, 22, 2007.
- [6] M. A. Delucchi and M. Z. Jacobson. Providing all global energy with wind, water, and solar power, Part II: Reliability, system and transmission costs, and policies. *Energy Policy*, 39(3):1170–1190, 2011.
- [7] J. Dongarra. Visit to the national university for defense technology changsha, china. *Oak Ridge National Laboratory, Tech. Rep.*, June, 2013.
- [8] J. Dongarra et al. The international exascale software project roadmap. *International Journal of High Performance Computing Applications*, page 1094342010391989, 2011.
- [9] H. J. Goldsmid. Bismuth telluride and its alloys as materials for thermoelectric generation. *Materials*, 7(4):2577–2592, 2014.
- [10] S. Greenberg, E. Mills, B. Tschudi, P. Rumsey, and B. Myatt. Best practices for data centers: Lessons learned from benchmarking 22 data centers. *Proceedings of the ACEEE Summer Study on Energy Efficiency in Buildings in Asilomar, CA. ACEEE, August*, 3:76–87, 2006.
- [11] G. Griepentrog, S. Hüttinger, and G. Vidiella. 21 EU HeatReCar Project: Waste Heat Recovery by Thermoelectric Power Conversion in Light-Duty Trucks. *Thermoelektrik*, 94:267, 2009.
- [12] E. Lefeuvre, A. Badel, C. Richard, L. Petit, and D. Guyomar. A comparison between several vibration-powered piezoelectric generators for stand-alone systems. *Sensors and Actuators A: Physical*, 126(2):405–416, 2006.
- [13] G. Li, V. Shrotriya, J. Huang, Y. Yao, T. Moriarty, K. Emery, and Y. Yang. High-efficiency solution processable polymer photovoltaic cells by self-organization of polymer blends. *Nature materials*, 4(11):864–868, 2005.
- [14] A. Nikolaev, I. Burylov, and S. Salahuddin. Intel® version of STAC-A2 benchmark: toward better performance with less effort. In *Proceedings of the 6th Workshop on High Performance Computational Finance*, page 7. ACM, 2013.
- [15] S. Ravindran, T. Huesgen, M. Kroener, and P. Woias. A self-sustaining micro thermomechanic-pyroelectric generator. *Applied Physics Letters*, 99(10):104102, 2011.
- [16] E. S. Rubin, A. B. Rao, and C. Chen. Comparative assessments of fossil fuel power plants with CO2 capture and storage. In *Proceedings of the 7th International Conference on Greenhouse Gas Control Technologies*, 2005.
- [17] M. Saleemi, M. S. Toprak, S. Li, M. Johnsson, and M. Muhammed. Synthesis, processing, and thermoelectric properties of bulk nanostructured bismuth telluride (Bi 2 Te 3). *Journal of Materials Chemistry*, 22(2):725–730, 2012.
- [18] D. Salerno. Ultralow voltage energy harvester uses thermoelectric generator for battery-free wireless sensors. *LT Journal*, 2010:1–11, 2010.
- [19] E. Saule, K. Kaya, and Ü. V. Çatalyürek. Performance evaluation of sparse matrix multiplication kernels on Intel Xeon Phi. In *Parallel Processing and Applied Mathematics*, pages 559–570. Springer, 2013.
- [20] J. Shalf, S. Dossanjh, and J. Morrison. Exascale computing technology challenges. In *High Performance Computing for Computational Science-VECPAR 2010*, pages 1–25. Springer, 2010.
- [21] G. J. Snyder and E. S. Toberer. Complex thermoelectric materials. *Nature materials*, 7(2):105–114, 2008.
- [22] B. TSAY. The Tianhe-2 Supercomputer: Less than Meets the Eye? *SITC Bulletin Analysis*, 2013.
- [23] K. Uchida, S. Takahashi, K. Harii, J. Ieda, W. Koshibae, K. Ando, S. Maekawa, and E. Saitoh. Observation of the spin Seebeck effect. *Nature*, 455(7214):778–781, 2008.
- [24] Y. Zhou, S. Paul, and S. Bhunia. Harvesting wasted heat in a microprocessor using thermoelectric generators: modeling, analysis and measurement. In *Proceedings of the conference on Design, automation and test in Europe*, pages 98–103. ACM, 2008.