Energy Neutral Wireless Sensing for Server Farms Monitoring

Maurizio Rossi, *Student Member, IEEE*, Luca Rizzon, Matteo Fait, Roberto Passerone, *Member, IEEE*, and Davide Brunelli, *Member, IEEE*

***Abstract*—Energy harvesting techniques are consolidating as effective solutions to power electronic devices with embedded wireless capability. We present an energy harvesting system capable to sustain sensing and wireless communication using thermoelectric generators as energy scavengers and server's CPU as heat source. We target data center safety monitoring, where human presence should be avoided, and the maintenance must be reduced the most. We selected ARM-based CPUs to tune and to demonstrate the proposed solution since market forecasts envision this architecture as the core of future data centers. Our main goal is to achieve a completely sustainable monitoring system powered with heat dissipation of microprocessor. To this end we present the performance characterization of different thermal-electric harvesters. We discuss the relationship between the temperature and the CPU load percentage and clock frequency. We introduce a model to simulate the power characteristic of the harvester and a prototype has been realized to demonstrate the feasibility of the proposed approach. The resulting system achieves a minimum 5 min sampling frequency of environmental parameters such as temperature, humidity, light, supply voltage, and carbon-monoxide/volatile organic compounds gases using a MOX sensor mounted on a commercial wireless node with a power budget in the microwatt range.**

***Index Terms*—Embedded systems, energy neutral systems, environmental monitoring, thermoelectric harvesting.**

I. Introduction and State of Art

IN THIS work we present an energy neutral monitoring device supplied by electrical power harvested from thermal sources. In particular, the proposed system harvests heat dissipated by high performance devices (such as server boards), and it is designed to monitor environmental parameters and valuable data to ensure safety inside server rooms. The system works in symbiosis with the server, since it relies on the server to guarantee its power supply, and acts as a data collector to monitor interesting features, including temperature and other quantities related to thermal and resources wastage.

Technology advances enable the development of innovative architectures for embedded systems towards orthogonal directions. From one side, systems characterized by high computing performance exhibit lower power consumption and thermal dissipation compared to previous generation of personal computers. In fact, ARM-based devices are considered a promising solution for building computing server for cloud service providers and Web 2.0 applications [1], [2]. On the other side, technology advances offer also low-power micro-controllers that are able to manage multiple sensors and to communicate wirelessly to others with few millijoules, and, thanks to their characteristics, it is possible to supply them with nonconventional power sources [3].

Generally, data centers facilities are always turned on, and are sources of high thermal waste and energy consumption. In fact, only 0.5% of total fossil fuel power is spent in useful computation, while more than 65% is used for cooling or lost as heat [4]. Server rooms have plenty of wasted heat, which can be recovered to generate electric energy, and mitigate wastage [5]. Given the numbers involved in the data centers scenario, even a small improvement that diminishes the power acquisition from the grid can introduce environmental and economic benefits.

In a generic computing device there are some hot points that represent a thermal source for harvesting, for example CPUs or GPUs. Power and thermal management constitutes an important trade-off between the performance and lifetime characteristic of a computing device. In fact, CPUs are designed to avoid high temperature by adopting dynamic voltage and frequency scaling, or other techniques to evade the necessity of using bulky cooling systems [6], [7]. Moreover, to guarantee durability, monitoring system temperature, and eventually the use of strategies for passive or active heat dissipation (heat sinks, fans, liquid cooling, Peltier cells, etc.) are necessary.

In addition, for safety and reliability reasons, there is the need to monitor server rooms, and wireless sensor networks are a promising candidate to undertake this task [8]. Through the use of specific WSNs, it is possible to efficiently monitor many interesting information, such as: local room temperature, humidity, air flow, air quality, moisture leak, and intrusion detection.

The objective of this work is to take advantage of the heat dissipated by the servers to power the monitoring sub-system, thus

Manuscript received March 01, 2014; revised May 26, 2014; accepted June 26, 2014. Date of publication July 25, 2014; date of current version September 09, 2014. This work was supported by project GreenDataNet(Grant 609000) funded by the EU 7th Framework Programme. This paper was recommended by Guest Editor G. Wang.

M. Rossi and L. Rizzon are with the Information and Communication Technology International Doctoral School, University of Trento, I-38123 Trento, Italy (e-mail: maurizio.rossi@unitn.it; luca.rizzon@unitn.it)

M. Fait is with the Department of Information Engineering and Computer Science, University of Trento, I-38123 Trento, Italy (e-mail: matteo.fait@unitn.it).

R. Passerone is with the Department of Information Engineering and Computer Science, University of Trento, I-38123 Trento, Italy. (e-mail: roberto.passerone@unitn.it)

D. Brunelli is with the Department of Industrial Engineering, University of Trento, I-38123 Trento, Italy (e-mail: davide.brunelli@unitn.it).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JETCAS.2014.2337171

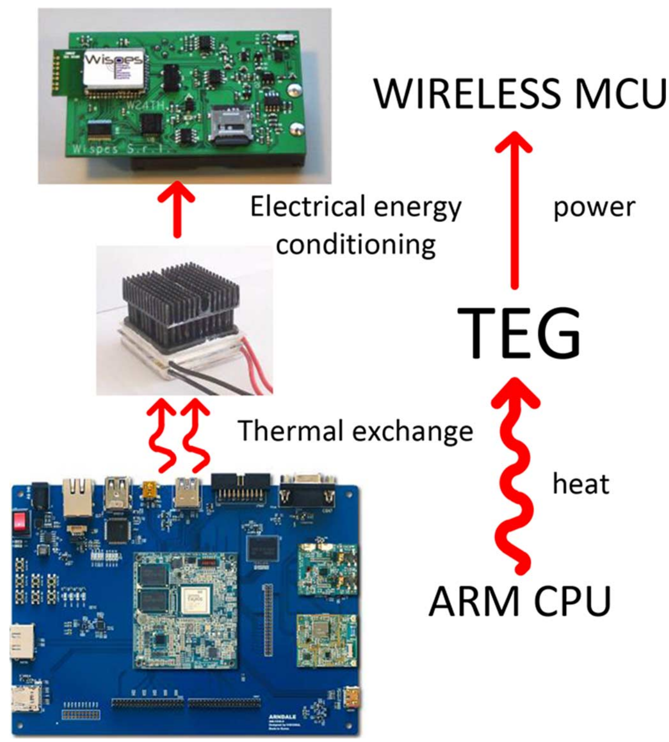

Fig. 1. Schematic view of the target application scenario. Energy wasted as heat can be collected and transformed into electrical energy to supply wireless embedded systems.

coupling technologies that are conceptually separated. Fig. 1 illustrates the basic components and the energy chain in our vision. The subsystem described in this paper provides monitoring service for the server and thanks to its low-power characteristics it is able to do this task for free. Contrary to battery powered systems, it does not require the intervention of an operator for battery replacement. Moreover, being independent from the standard power grid, it does not contribute to the energy bill, and can be mounted or detached easily.

The critical challenge in the design of such a system is to demonstrate that the presence of extra hardware does not affect the performance of the server. The proposed system composed of a thermoelectric harvesting module, an energy buffer and a wireless microcontroller (MCU) has been realized and demonstrated to be effective from both the monitoring and the heat dissipation points of view.

Our work focuses on a single ARM-based board—referred elsewhere as ARM CPU—that represents a single CPU inside a server farm. On top of it we place the thermo-harvesting device—TEG—to scavenge energy and supply a stand alone embedded wireless node—WSN NODE—through a dedicated conditioning circuit. The results provided obviously scale up with the dimension of the server farm. The final goal is to pack the whole node into a *System on Chip* (SoC depicted in Fig. 2) with the size and shape of a commercial heat sink, where harvesting, sensing and the radio are embedded in the base of the sink itself. This will be feasible thanks to the advancement in manufacturing processes, μ-scale TEGs [9], and microwatt radio transceivers [10].

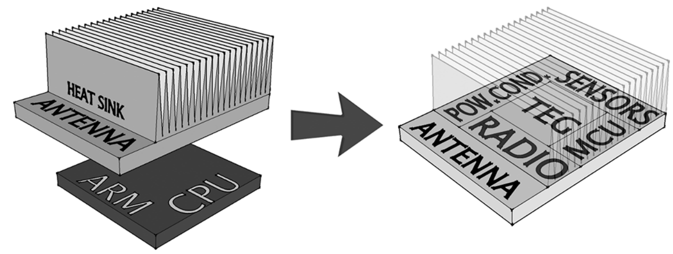

Fig. 2. Exploded view of the envisioned SoC. Harvesting, conditioning, and monitoring systems are embedded in the heat sink.

The key element for harvesting thermal energy is represented by thermoelectric generators (TEG) that exploit a phenomenon called the Seebeck effect to convert heat (or temperature gradient) into electrical energy [9]. They have been widely used for a variety of applications, including energy harvesting from heat dissipated in combustion engine vehicles. TEGs generally exhibit low efficiency, even in the high temperature context of automotive applications [11]. Some works studied the possibility of applying thermoelectric generators into computing systems [12]–[14]. In the literature it has been demonstrated that it is possible to partially harvest thermal energy from electronic computing systems and convert it into electric energy [15], [16]. However, the recovered energy is too small to be used to extend the battery lifetime of the main system [17], but it can be used to supply context aware computers (subsystem sustain, multi-modal application, and distributed sensors), or for cooling.

In this paper, we present multiple scientific contributions. First, we provide further details on the characterization of different models of microwatt TEGs placed on CPUs. The paper also shows the relationship between the temperature that originates on the CPU package with various computing loads and at various clock frequencies. As third contribution, we introduce a model for the estimation of the harnessed electrical power and we developed a simulator that helps us estimate the outcomes, and support the design of the proposed system. Finally, the fourth contribution of this work is the design and assembling of the prototype of an energy neutral wireless computing system for environmental monitoring.

This paper is organized as follows. In Section II, we describe our experimental setup and the devices used in our work, along with comments on the characterization of TEGs. The number of experiments and characterization we present is necessary to understand the performance and the behavior of the microwatt harvester. Section III describes the wireless sensor node and its power consumption characteristics. In Section IV, we present a numerical simulation that allows us to evaluate the feasibility of realization of a self sustain monitoring node knowing the amount of energy harnessed from high performance embedded boards. In Section V, we show in details the real application scenario and draw the conclusion of this work.

II. Thermoelectric Generation

Thermoelectric generators (TEG) are devices that convert heat into electrical energy, through a phenomenon called the Seebeck effect. Such a component is similar to a Peltier cooler

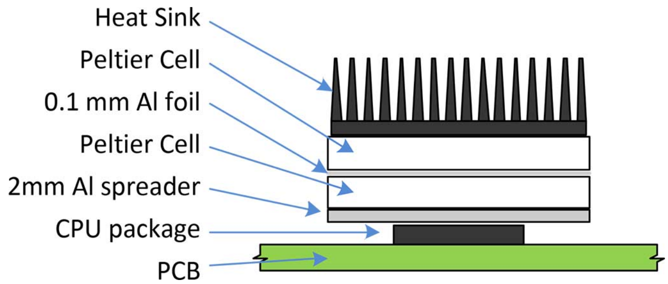

Fig. 3. Thermoharvester schematic made of two Peltier cells, spreaders, and sink in stacked fashion. We put thermal grease in between each layer of different materials to guarantee maximum thermal contact.

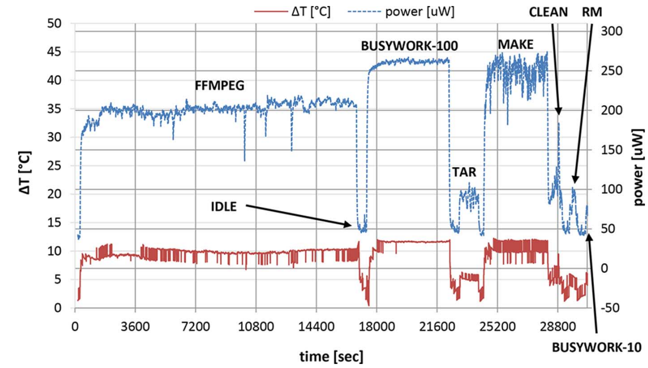

Fig. 4. Output power of the Nextreme TEG placed on top of the Pandaboard while running the benchmark application. Continuous line represents the thermal gradient as difference between the hot side and the heat-sink.

operated in the opposite way. When operated as a generator, one side of the device is heated to a temperature greater than the other side, and as a result a difference in voltage will be observed at the output. However, a well-designed Peltier cooler will be a mediocre thermoelectric generator, and vice versa due to different design and packaging requirements.

To estimate the order of magnitude of the power generated by thermoelectric generators in the context of computing systems, we have compared three different devices provided by different manufacturers. These devices differ in terms of aspect ratio, number of cells and their arrangement.

Two of these are specifically designed as thermoelectric generators, in fact they are made by microfabricated thermocouples optimized for the Seeback effect. The third one is a generic Peltier cell—generally used for cooling applications—that we used in two different configuration: a single cell, and two cells connected in series. The latter configuration is shown in Fig. 3, in the other case the Peltier/thin spreader/Peltier is substituted by a single element. The devices we chose as TEGs are as follows.

- *Nextreme* eTEG HV56 Thermoelectric Power Generator with 32 × 32 mm squared footprint, and a 31 mm × 33 mm power generator without its output power regulator [18];
- *Micropelt* TE-CORE7 TGP-751 ThermoHarvesting Power Module with a 43 × 33 mm heat sink, and a circular footprint with 5 mm radius [19].
- *Peltier-Cell* PE1-12706AC 40 × 40 mm squared cells.

We built the stacked structure presented here to optimize the thermal exchange between ARM CPUs and Peltier cells. The lower spreader (thickness 2 mm) allows the heat to distribute on the whole surface, while the one in between (thickness 0.1 mm) speeds up the exchange between cells, and spreads the heat once again. The spreaders are made of aluminum because of its thermal efficiency, low cost, and market availability. The last layer of the stack is a commercial heat spreader that optimizes the heat dissipation in the environment.

The characterization of harvesters performance are obtained using open circuit voltage and impedance matched power. The former one is used to evaluate the Seebeck coefficient of a given cell with N number of p-n junctions, using the equation $\alpha = V_{oc}/N\Delta T$. The latter states the maximum power that the cell can provide to a matched load.

Given the target of implementing an energy neutral monitoring system based on wireless embedded devices, we were mostly interested about performance in terms of maximum power and in different working conditions of the selected devices. To this end, in our experiment we use as "heater" two embedded system boards based on ARM processors.

- A Pandaboard with a 1 GHz ARM Cortex A9 CPU and 1 GB of low power DDR2 RAM [20].
- A Samsung Arndale equipped with an Exynos 1.7 GHz dual-core ARM Cortex A15 processor with 2 GB DDR3 RAM 800 MHz [21].

Both systems run Linux kernel version 3.10.

We conducted the characterization of the TEGs listed above by executing different tasks in sequence in a benchmark fashion. In particular the tasks were selected to achieve a range of time length and CPU loads (with fixed clock frequency set as the maximum of each ARM CPU board).

- **Video encoding** using `ffmpeg`[1] with four threads, which converts a 2-h-long movie.
- **Multithread application** that performs millions of algebraic and trigonometric computations of floating point numbers using a user defined number of threads (called `busywork` in pictures).
- **Kernel operations** in this task a Linux kernel is uncompressed, compiled and then cleaned, then the folder is removed.

Output voltage and current have been measured over a matched load with a 1 s period.

The benchmark applications allowed us to obtain an output power profile for several configurations (ARM CPUs plus TEGs). Detailed results of the experiments are reported in Figs. 4–11. These pictures present the output impedance matched power of each TEG under the ΔT produced by the heat dissipated by the CPUs, computed as the product of the measured voltage and current. ΔT in turn is strictly related with the CPU load. The air conditioning system of the lab kept the room temperature almost constant in the range 22 °C to 25 °C. Generally speaking, the longer the task the higher the CPU temperature as well as the number of memory accesses and switches (algebraic computations). Pandaboard's CPU builds up lower temperature with respect to the Arndale's (37 °C versus 80 °C), resulting in a 10x difference in the

[1]ffmpeg.org

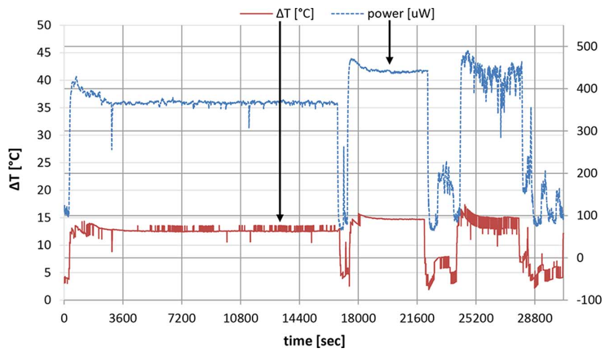

Fig. 5. Output power of the Micropelt on Pandaboard.

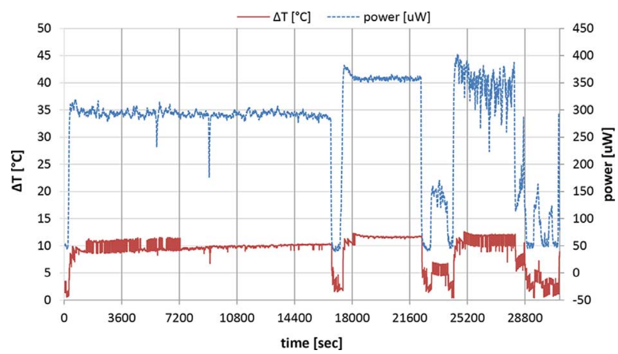

Fig. 6. Output power of a single Peltier cell on Pandaboard.

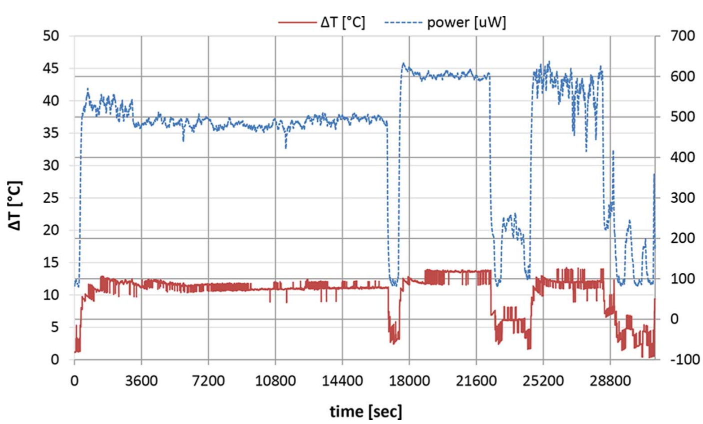

Fig. 7. Output power of a series of two peltier cells on Pandaboard.

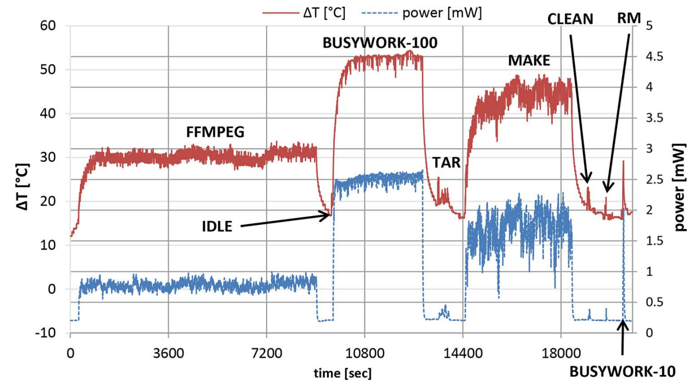

Fig. 8. Output power of the Nextreme on Arndale.

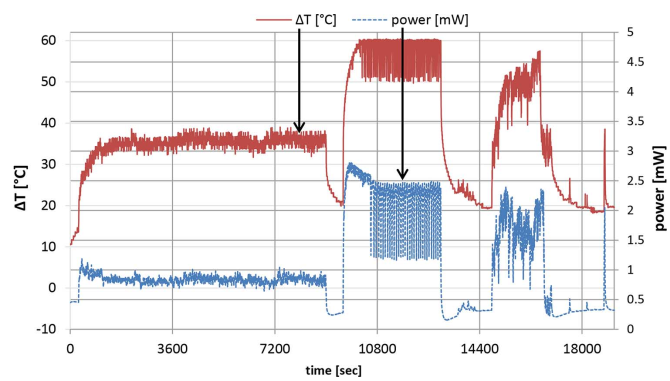

Fig. 9. Output power of the Micropelt on Arndale.

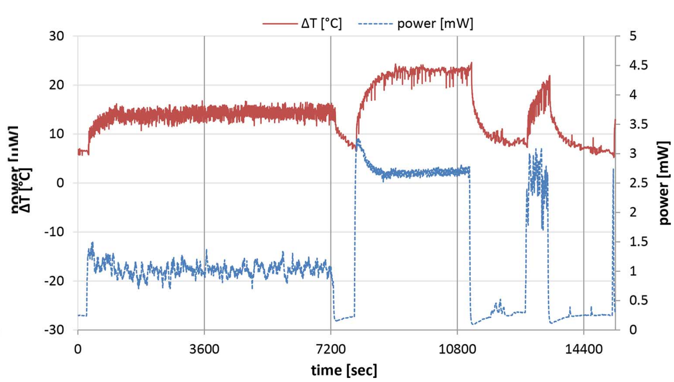

Fig. 10. Output power of a single Peltier cell on Arndale.

TEG measured output power in our experiments (microwatt range for Pandaboard while milliwatt for Arndale's). All but Nextreme exhibit a good reactivity to fast temperature spikes, this one produces smooth power profiles while others result in sharp peaks at the beginning of each task. Micropelt has the worst thermal efficiency with respect to the others: we can notice the effect of the voltage scaling in the middle of the benchmark—around 10800 s—clearly depicted in Fig. 9 and less evidently in Fig. 8, where the output power "flickering" reflects the temperature behavior, which is regulated automatically by the Linux kernel governor. The two Peltier cells almost double the output power of the single cell, the spikes move from 350 to 600 μW in Pandaboard's case and from 2.6 to 3.6 mW with Arndale.

The harvesting system composed by two Peltier cells in series, arranged in stacked fashion, exhibits higher performance with respect to other configurations. The comparative results, pictured in Fig. 12, demonstrate that the performance of this TEG setup are better for tasks that last more than 30 s: `idle`, `ffmpeg`, `busywork-100`, `tar`, `make`, and `clean`. Here, we can notice the maximum mean impedance matched power of almost 3.6 mW on Arndale and 600 μW in case of `busywork-100`. Only in case of very short tasks, `rm` and `busywork-10` last few seconds each, the Micropelt TEG performs better and only on the Arndale board.

The stacked TEG outperforms the other solutions also in terms of total energy. Considering the results of both Figs. 14 and 15 together, the conclusion for Arndale board is that, in lower time, the stacked architecture is able to harvest the

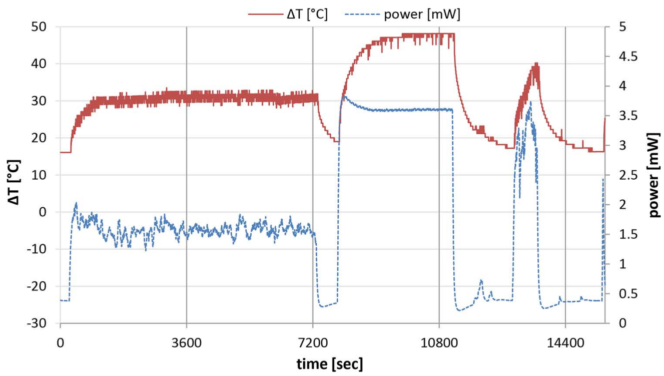

Fig. 11. Output power of a series of two Peltier cells on Arndale.

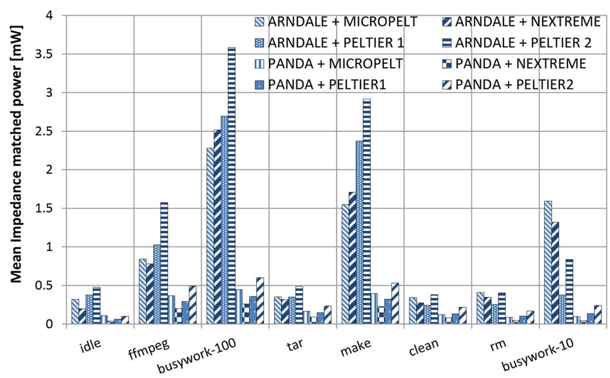

Fig. 12. Average output power per benchmark's stage.

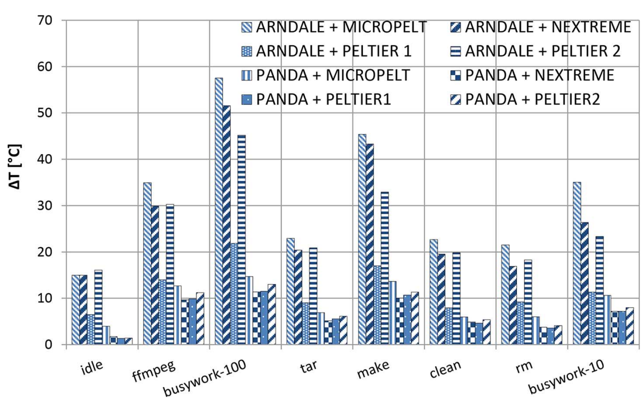

Fig. 13. Average ΔT per benchmark's stage.

highest amount of energy (26 J in 4^h 30^m with double Peltier cell versus 24 J in 5^h 45^m with Nextreme and 20 J in 5^h 20^m); while for Pandaboard the execution time is almost constant (26 J with double Peltier, 6 J with Nextreme, and 11 J with Micropelt in 8^h 20^m). The difference in the Arndale's results is due to the effect of the voltage scaling performed by the kernel governor, which reduces the execution speed to reduce the CPU temperature. As stated above, Nextreme and Micropelt TEGs are not able to dissipate the heat generated by the Arndale's CPU and also to efficiently exploit the high variation in the resulting ΔT.

Having a device on top of the CPU package limits the natural heat exchange of the microprocessor. If the system overheats, high temperature may cause structural damages to the ARM

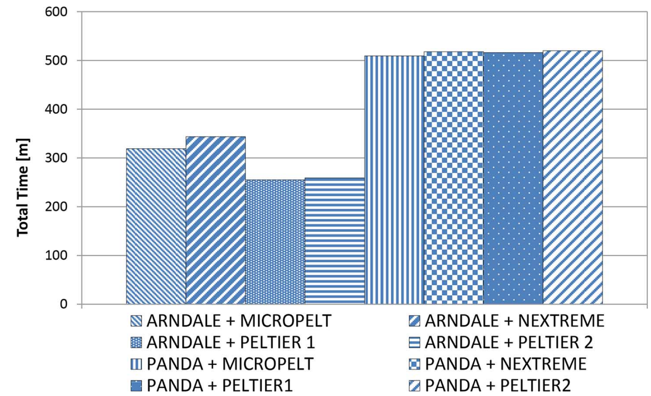

Fig. 14. Total benchmark execution time per configuration.

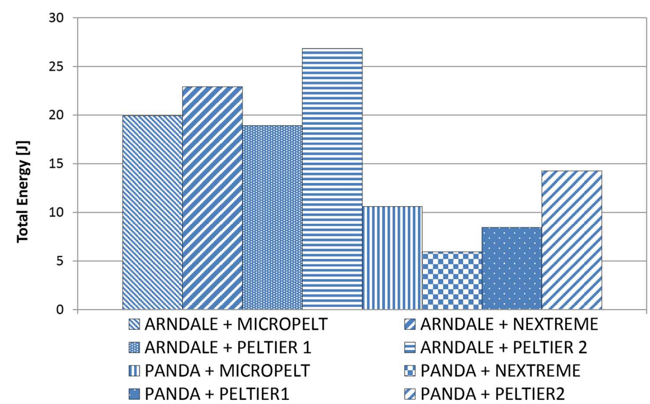

Fig. 15. Total impedance matched energy per configuration.

CPU board, and may reduce the device lifetime. Also the TEG may suffer structural damages if exposed to excessive thermal gradients [22]. We evaluated the impact of the harvesting system on the CPU temperature. The results are presented in Fig. 16 where the CPU temperature is depicted against time for three different scenario: 1) natural dissipation (CPU used as it is), 2) heat sink on top of the CPU package and 3) the proposed harvester. Without any thermal dissipation the CPU immediately reaches a dangerous temperature and the voltage scaling decreases its performance for safety (dash–dot line, the temperature reaches immediately the critical threshold of 82 °C and then lays bounded around 80 °C). In the other two cases the rate of increase is smoother and the presence of the harvester slows down the thermal exchange; at steady state both cases exhibit comparable performance (stability around 70 °C). This once again underlines the good performance of the double Peltier stacked structure specifically built for this work, which demonstrates the best thermal dissipation among all, as evident from the comparison in Fig. 13.

We can state that the stacked series of two Peltier is the most suited solution to build an energy neutral embedded system to monitor the state of the server and its environment. We refer to this configuration for the rest of this paper.

III. Microwatt Sensor Network

We realized a distributed system which is completely self sustained by microwatt harvested power. It is based on WSN nodes widely used to build distributed systems of portable electronics

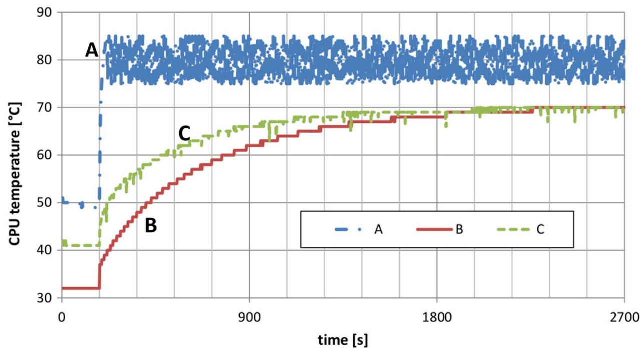

Fig. 16. Comparison of CPU temperature rise with 100% CPU load @1.7 GHz: (A) natural heat dissipation, (B) heat sink, (C) double Peltier stacked structure above CPU package.

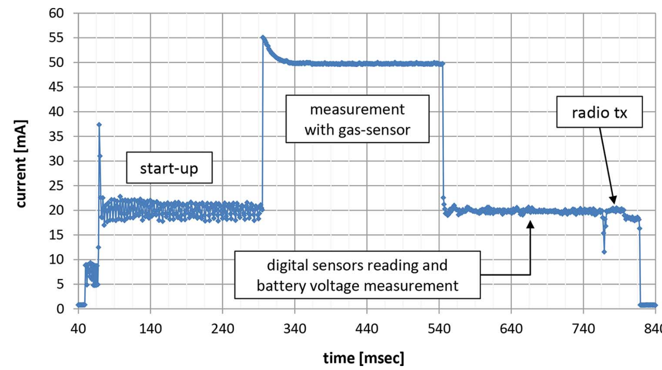

Fig. 17. Wireless sensor node, current consumption during monitoring activity.

devices targeted to the realization of pervasive electronic environments. Each node in the network embeds sensors to collect data that are in turn routed to a central gateway (aggregator or coordinator) responsible for data processing, storage, and for the actuation of predefined policies. Nodes can also be designed as active devices to implement the control instruction sent by the policy manager. Key feature and enabling characteristic of the nodes is the very low power required both in active (sensor sampling) and sleep modes of operation. The main limit is the power required by the communication task. Reducing the power needed by a transceiver to microwatt range definitely results in a lower range of communication. The introduction of energy harvesting techniques from TEG provides state-of-art powering solutions replacing disposable batteries with dedicated electronics; thus lowering the costs and the maintenance, and slightly increasing the design complexity. We used a commercial wireless node to implement a self-sustained environmental monitoring system and evaluate its feasibility.

The WSN node is the W24TH from Wispes[2] (Wireless MCU in Fig. 1), and is specifically designed to reduce the most the power consumption and leakage. It is based on the powerful JN5148 microcontroller unit (MCU) from NXP-Jennic[3], which embeds a 32 bit MCU and a IEEE 802.15.4 compliant radio transceiver with its antenna. The board hosts a number of sensors that can be individually switched on and off, reducing the power wasted. The available digital sensors allow the node to measure temperature, relative humidity and light, while an internal ADC permits to gather information from an analog Gas sensor and of the voltage supply of the node. The ad-hoc monitoring application developed uses all the listed devices since all those parameters can be useful in the detection of damages in the server.

The faster the execution the lower the energy required, so we put most of the effort in the design of this application to reduce the execution time, while assuring reliability of the monitoring. We configured the device to use a very simple network protocol, built on top of the IEEE 802.15.4 MAC layer. Simply, the device sends data to the coordinator as soon as enough energy is available to complete the task. Sensors require some specific time to reach a steady state response. We turn them on at once and perform the measurement as soon as they are stable. This approach reduces the time the node is active, and hence reduce the power consumed by the node. Fig. 17 shows the current drained by the WSN node while executing the firmware described above.

The required energy has been measured using a 1 Ω shunt resistor using an oscilloscope. The whole execution requires almost 800 ms and the maximum required power is around 160 mW (@3 V), 70 mW on average. The target issue of the proposed work is to match an unstable and unpredictable microwatt range power source with a device that requires milliwatt, as shown above. The conditioning circuit, presented in the next section, fulfills this gap by adaptively choosing a suitable duty-cycle, based on the charge collected using TEGs.

Particular care must be given to the presence of a chemoresistive gas sensor (MOX) which is responsible for the current increment in the middle of the firmware execution (from 300 to 550 ms in Fig. 17). This increment is higher than the radio transmission that occurs at the end of the monitoring task (from 760 to 800 ms in Fig. 17) and requires almost 51 mW (17 mA @3 V). The profile depicted is due to the MICS-5121[4], a MOX gas sensor targeted to carbon-monoxide and volatile organic compounds measurement, which requires 76 mW of power to work. Despite their high consumption, MOX sensors are widely used in monitoring applications—for example for the realization of smoke detectors—since they offer long term stability, and do not require maintenance. Instead of keeping it always on, we used the gas monitoring strategy presented in [23]–[25]. This strategy allows the energy required by MOX sensors to be reduced by a factor of 20. Depending on the interval between MOX sampling phases it is possible to characterize its response reducing the sensor switch-on time from some seconds to 250 ms.

Lastly, we evaluated the impact of this kind of monitoring device in terms of energy taken from the grid. We compared two situations, (a) continuous, when energy is provided by means of batteries or dedicated wires, also during sleep; versus (b) discontinuous supply, the proposed and implemented choice where the node is switched on only when the reserve can sustain it. In case of 0.25% duty-cycle (one measure every 5 min) the W24TH requires 2.7 mW on average in case (a) against the 225 μW on average in case (b). When duty cycling is 0.02% (one measure per hour) we get 2.4 mW for (a) and 18.6 μW for (b). Translating these into energy per year (Wh) we can estimate a saving of 92% (1.4 KWh against 117 Wh) with 0.25% duty-cycle, and 99.25%

[2]Produced by Wispes s.r.l. http://www.wispes.com

[3]http://www.jennic.com/products/wireless_microcontrollers/jn5148

[4]Produced by E2V http://www.e2v.com.

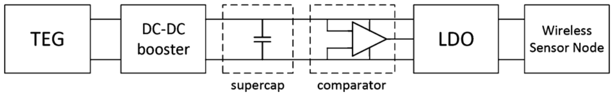

Fig. 18. Block diagram of the conditioning circuit.

with 0.02% duty-cycle (1.3 KWh against 10 Wh). These are tiny numbers compared to the energy required by server farms but they must be scaled with the numbers of nodes in the WSN. The resulting energy saving combined with the lack of need for maintenance makes the WSN monitoring system inexpensive.

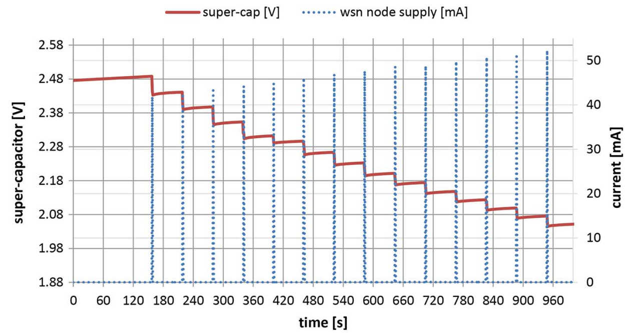

Fig. 19. WSN node lifetime when supplied with direct connection to the supercap charged at $V_{\max}$ (bypassing comparator and LDO).

IV. Simulation

The preliminary results and information presented in previous sections are used to evaluate the feasibility of an energy neutral monitoring device with wireless connectivity, and to design it. The system we designed and developed is made of three parts: 1) the TEG, 2) the conditioning circuit, and 3) the WSN node. We decided to keep only the Arndale Board as testing environment from now on because 1) it has higher computational performance, 2) it generates more heat, and 3) its CPU architecture is more recent with respect to the Pandaboard, hence more realistic to simulate an ARM-based data center. Finally, 4) to speed up the development due to the length in time of our tests.

Simulations were used to predict the amount of energy harvested from the heat dissipated through the CPU package of the device under different loads, and various clock frequencies. Due to the high variability of the output quantities, and the uncertainty of the cooling capabilities, it is likely that the simulation would overestimate the real output power profile (cf. [26]), thus we have taken it into account. We evaluated the characteristic of the system at a settled clock frequency and load percentage.

The output power provided by TEGs depends on the ΔT as previously shown, so we decided to operate in terms of server CPU load and clock frequencies ($f_{\rm CLK}$) that are directly related with the dissipated heat. The available Arndale Board's $f_{\rm CLK}$ ranges from 200 MHz up to 1.7 GHz, with 100 MHz steps. In our simulations we selected the range 1.2–1.7 GHz because not enough heat is dissipated below this interval. In addition, we chose a sequence of CPU loads of 0%, 30%, 50%, 70%, and 100%, for a total of 30 different working conditions.

The electric power generated by the TEG fluctuates, therefore it needs to be conditioned so that it can be stored, or used. We build a conditioning and storage circuit based on commercial components, arranged as shown in the block diagram of Fig. 18.

A supercapacitor (3 F, 2.7 V) was chosen as storage unit because of the energy required by the monitoring wireless node. Fig. 19 depicts the lifetime of the WSN node supplied by direct connection to this supercap. The reason for using a supercapacitor stems from its low internal resistance, which results in a high surge current, as well as its low self-discharge characteristic that make it desirable in energy harvesting application.

A comparator is used to enable the current flow only when the charge stored in the supercap is enough to power the node. The comparator has been realized with an integrated circuit and some components to tune the range. The high threshold $V_{\max}$ is set to 2.48 V, that is also the WSN node supply voltage. The lower threshold $V_{\min} \simeq 2.41$ V guarantees to completely disable output power, so that the node is isolated and unable to drain current. This choice does not affect the node network functioning when the node is turned back on. The above two thresholds were chosen for multiple reasons. First, with 2.48 V the super capacitor has enough energy to supply more than a single node activity (that has been demonstrated sufficient in Fig. 19), and it is charged up to a value lower than the nominal. Secondly, the tiny gap (≈70 mV) allows the comparator to switch off the LDO immediately after the node completes its duty.

The influence of the CPU load in the generation of thermal gradient must be understood to predict the availability of electrical power to supply the small computing device. The profile, based on the knowledge of the behavior of the TEG exposed to different thermal gradients driven by the CPU activity, is determined by the measurements described in Section III. We processed the data of Figs. 7 and 11 with different regression scheme (provided by commercial software) to obtain the input-ΔT versus output power. Our approach differs from the ones proposed in the literature [26]–[29] since we consider both TEG and dc-dc boost circuit efficiency as a black-box. Fig. 20 depicts the empirical input-output relation of our harvesting device. The resulting model is a composite piece-wise continuous function of the temperature, and switches between an initial exponential segment between 5° and 20° and 20° ΔT, followed by a polynomial segment between 21° and 31°, and finally by a linear segment from gradients ranging from 32° up to 40°. This arrangement was required due to the highly nonlinear regulation between the input temperature and the efficiency of both the TEG and the booster. The estimated parameters of the model are $K = 0.08$ and $\alpha = 0.16$ for the exponential scheme in the 5 °C to 20 °C of ΔT range; $K_1 = -0.000005$, $K_2 = 0.00001$, $K_3 = 0.0104$, $K_4 = -0.7$ for the polynomial component in the 21 °C to 31 °C of ΔT range; and finally $\beta = 0.025$ and $\gamma = 2.8$ from 32 °C up to 40 °C of ΔT. We fed this model into the simulation software along with empirically estimated thermal gradients for a range of CPU loads and frequencies, obtained with further experiments, not reported due to space reasons.

The simulation is based on a simple energy budget evaluation for each fixed time-slot (1 s). The input power is provided to the harvesting circuit by the TEG, and the conditioning circuitry manages the wireless node. We fixed the efficiency of the complete harvesting system equal to 5%. This value takes into account the mismatch between the TEG and the dc-dc booster,

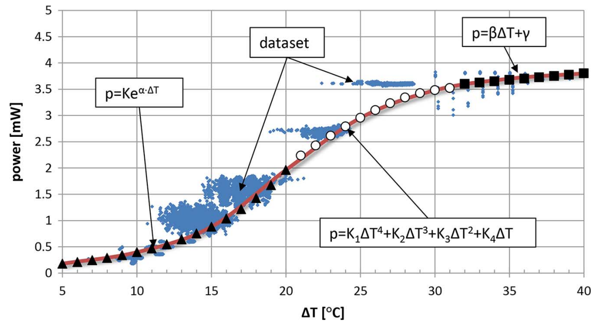

Fig. 20. Input output relation of the proposed harvester consisting of booster and supercap.

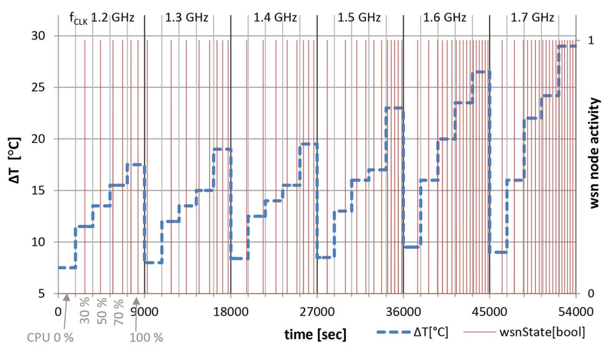

Fig. 21. Energy budget simulation, wsn node completed tasks as function of the TEG's ΔT, tuned plying with CPU loads and $f_{\rm CLK}$.

(based on an inductive transformer [30]), the efficiency of the super-capacitor during charge and the overall performances of the integrated circuits and components (Comparator, LDO, etc.).

The result of the simulation is reported in Fig. 21. In this case we used a fixed time intervals of 1800 s (30 min) for each CPU load, hence 2.5 h for each $f_{\rm CLK}$. The graph shows the occurrence of monitoring tasks that the wireless node performs as the time elapses, in binary fashion, on top of the ΔT sequence, function of the CPU activity. The simulation clearly shows the strong correlation between temperature and frequency of monitoring, that is summarized in Fig. 23. These results justify the effort of building a real energy neutral system since it would be possible to achieve 5 min monitoring frequency at most, when the CPU is under a heavy load, and the monitoring can be crucial to prevent failures.

V. Experimental Results

We realized the harvesting circuit described above to validate the theoretical results of the simulation. We used the same sequence of CPU loads and $f_{\rm CLK}$ presented in the simulation's description (Section IV). We achieved the target load activities by interleaving mathematical computations with sleep intervals to generate variable load profiles. The results of the 15 h running of the experiment are presented in Fig. 22. In this picture the average current drained by the node is depicted on top of the TEG's ΔT. We can immediately notice the reduced number of monitoring tasks performed by the node with respect to simulated estimate. However, during high demanding tasks, we observe a tight correspondence between simulated and real occurrences.

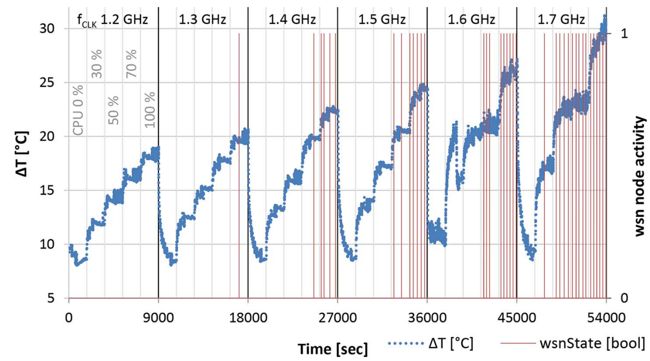

Fig. 22. Real test of energy neutral monitoring using WSN nodes. Same CPU load and $f_{\rm CLK}$ used in simulation.

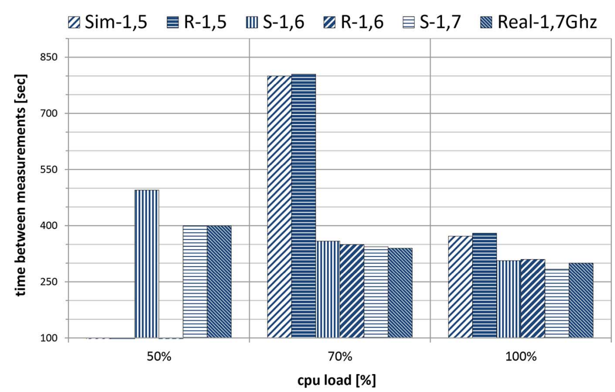

Fig. 23. WSN monitoring intervals, real (R) versus simulated (S) results, as function of load and $f_{\rm CLK}$.

Comparative results are summarized in Fig. 23. In this picture we compared the resulting monitoring frequencies, in terms of time between two sampling, for $f_{\rm CLK}$ in the 1.5–1.7 GHz range. The extracted TEG and circuit model does not fit well the real efficiency for ΔT lower than ≈ 20 °C, but it perfectly fits in the higher range, where the worst case mismatch is within 15 s. The difference is due to two reasons: 1) the fixed thermal gradients used in the model, which constitutes a very strong simplification of the CPU behavior; and 2) the fixed efficiency used to describe the control circuit (comparator and LDO), which is overestimated, since also here efficiency depends on the working point. CPU frequencies lower than 100%@1.5 GHz results in monitoring frequencies that are very difficult to evaluate due to extremely large time requirements. The efficiency of the control circuit is definitely lower than the expected with very low input power and result in very slow supercap recharging below $\approx$20°C. This arises from system nonidealities, expected as stated in Section IV; whose estimation is out of the scope. As a consequence, the efficiency of the real control circuit is lower than what is predicted by the model for temperatures below 20 °C.

When demanding CPU tasks are running we can reach a very good 5-min interval ($\approx$300 s in case of 100%@1.6 and 1.7 GHz). In some conditions (for example 70%@1.6 and 1.7 GHz) we have better performance in real tests than in simulation. This phenomenon can be related to the fact that very fast thermal

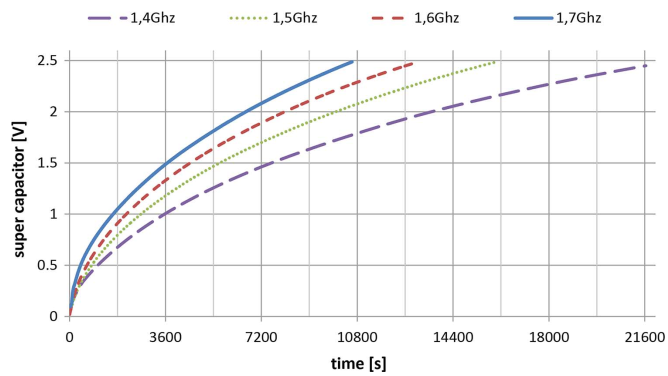

Fig. 24. 3 F supercapacitor charge time with the proposed architecture consisiting of TEG, booster and supercap.

variations result in higher output power with respect to a constant thermal profile, as can be noticed considering the ΔT behavior in Fig. 22. Lastly, we can state that the limit of the prototype implementation is 5 min.

The presence of the TEG in the proposed configuration does not affect the CPU performances at all. Considering the initial benchmarks, in Figs. 8 and 9 we can notice the very high ΔT and the effect of the frequency scaling on the power harvested in case of demanding tasks. We can notice a higher efficiency of the proposed thermo harvester for lower thermal gradients compared with the result shown in Fig. 11. In this last case there is no need for frequency regulation since the thermal dissipation provided by our solution overcomes also the standard package, as shown in Fig. 16 in Section II.

For the sake of completeness, we evaluated also the "*boundary conditions*" of the proposed system: 1) the cold-start time and 2) the autonomy after CPU switch-off.

With cold-start time we refer to the amount of activity time of the embedded board required to charge the supercap from 0 V up to $V_{\max}$ that corresponds to the first activity of the WSN node. Fig. 24 presents the cold start time for four different CPU conditions, ranging from 1.4 GHz up to the maximum clock frequency of 1.7 GHz, always with 100% of CPU load. These curves show how much time the ARM CPU board is required to run in order to have the monitoring subsystem working. We can notice that the time required at 1.4 GHz is roughly double the time required at a clock frequency of 1.7 GHz. To improve the cold-start time, different solutions can be adopted: on the storage side, the usage of bootstrap capacitors [31], [32], or reconfigurable capacitor bank block (CBB or flying capacitors) [33], [34], allow the rate of the transferred energy to be increased, using multiple capacitors charged in parallel, that are then switched into series configuration to achieve the required voltage on the load. On the boosting interface side, both the capacitance of the charge pump can be decreased to speed up the switching frequency of the oscillating circuit [35], and a maximum power point tracking stage can be introduced to maximize the power transferred [36].

In case of switch-off or, equivalently, very low CPU load, the supercap can sustain the monitoring node for only one extra measurement with the $V_{\min}$ threshold set as above. To be able to completely drain the energy from the supercap (as in the case depicted in Fig. 19 where it is demonstrated that it is possible to sustain the node for more than 10 min with monitoring interval of 1 min), extra hardware would be required to bypass the comparator that is actually off after the last node activity. Otherwise, rising the supercap to the maximum rated voltage, or replacing it with another rated to higher voltage can withstand this limitation.

The energy neutral system presented can also be used as a temperature sensor itself since the sampling frequency and data transmission rates depend on the amount of energy scavenged. This intrinsic feature paves the way to the development of a kind of passive temperature sensors that could be used where the thermal energy is very low, space limitation can not afford the presence of external sensors and temperature changes exhibit slow variation in time.

These considerations are left open for further improvements and to the development of a more complex monitoring system since it is expected that the technological improvements will reduce the power required by wireless transceivers to the microwatt scale.

VI. Conclusion

We presented the design and validation of an energy neutral embedded system in the microwatt class for environmental monitoring in server farms. Even in small amount, having computing devices detached from the grid contributes to green energy initiatives. Moreover, having no need for batteries decreases the use of toxic substances.

The proposed system consists of a thermoelectric harvester, a conditioning circuit and a wireless sensor network board. Disposable batteries were replaced by the harvester that extract electrical energy from the heat dissipated by the server's CPU to supply the microcontroller.

We characterized commercial TEGs as well as custom ones to select the most suited to scavenge electric energy from ARM CPUs, the current technology trend in server's industry. The selected TEG was then used to simulate and then realize the prototype energy neutral WSN application. Our system achieves a minimum 5 min interval between successive measurement and transmission of environmental parameters (temperature, humidity, light intensity, self supply voltage, and air safety through a MOX gas sensor).

We fit a microwatt power supply with a wireless enabled embedded platform with high energy demand. The ultimate goal is to scale down the whole platform into a SoC the size of a standard heat sink and provide real-time plug-&-play monitoring system for free.

References

[1] R. Merritt, "Dell, IBM give thumbs up to ARM servers," *EE Times*, 2010 [Online]. Available: http://www.eetimes.com/document.asp?doc_id=1256326
[2] Z. Yeraswork, "AMD launches first ARM-based server CPU," *EE Times*, 2014 [Online]. Available: http://www.eetimes.com/document.asp?doc_id=1320985
[3] D. Carli, D. Brunelli, L. Benini, and M. Ruggeri, "An effective multi-source energy harvester for low power applications," in *Design, Automat. Test Eur. Conf. Exhibit.*, Grenoble, France, Mar. 2011, pp. 1–6.

[4] D. C. S. Group, "Meeting the energy efficiency and financial challenges in it," in *BCS, Tech. Rep.*, 2007 [Online]. Available: http://www.oxon.bcs.org/downloads/BCS-DCSG 20070927.pdf

[5] A. Haywood, J. Sherbeck, P. Phelan, G. Varsamopoulos, and S. K. Gupta, "Thermodynamic feasibility of harvesting data center waste heat to drive an absorption chiller," *Energy Conversion Management*, vol. 58, no. 0, pp. 26–34, 2012.

[6] C. Bolchini, A. Miele, A. Das, A. Kumar, and B. Veeravalli, "Combined DVFS and mapping exploration for lifetime and soft-error susceptibility improvement in mpsocs," in *Design, Automat. Test Eur. Conf. Exhibit.*, Dresden, Germany, Mar. 2014, pp. 1–6.

[7] H. Yu, R. Syed, and Y. Ha, "Thermal-aware frequency scaling for adaptive workloads on heterogeneous mpsocs," in *Design, Automat. Test Eur. Conf. Exhibit.*, Dresden, Germany, Mar. 2014, pp. 1–.

[8] L. M. Oliveira and J. J. Rodrigues, "Wireless sensor networks: A survey on environmental monitoring," *J. Commun.*, vol. 6, no. 2, pp. 143–151, 2011.

[9] R. Vullers, R. van Schaijk, I. Doms, C. V. Hoof, and R. Mertens, "Micropower energy harvesting," *Solid-State Electron.*, vol. 53, no. 7, pp. 684–693, 2009.

[10] A. Moradi and M. Sawan, "A new FSK-based transmitter dedicated for low-power wireless medical transceivers," in *Proc. 24th Can. Conf. Electr. Comput. Eng.*, May 2011, pp. 1238–1241.

[11] K. Zorbas, E. Hatzikraniotis, and K. Paraskevopoulos, "Power and efficiency calculation in commercial TEG and application in wasted heat recovery in automobile," in *Proc. 5th Eur. Conf. Thermoelectr.*, 2007, p. T8.

[12] R. Denker, A. Muhtaroglu, and H. Kulah, "Empirically based methodology for thermoelectric generation in notebook systems," in *Int. Conf. Energy Aware Comput.*, Istanbul, Turkey, 2011, pp. 1–5.

[13] Y. Zhou, S. Paul, and S. Bhunia, "Harvesting wasted heat in a microprocessor using thermoelectric generators: Modeling, analysis and measurement," in *Design, Automat. Test Eur.*, Mar. 2009, pp. 98–103.

[14] R. Denker, A. Muhtaroglu, and H. Kulah, "Empirical proof of concept for TE generation in mobile computers," in *Int. Conf. Energy Aware Comput.*, Guzelyurt, Cyprus, 2012, pp. 1–6.

[15] T. T.-T. Lai, W.-J. Chen, Y.-H. Hsieh, K.-H. Li, Y.-Y. Su, P. Huang, and H.-H. Chu, "Why blow away heat? Harvest server's heat using thermoelectric generators," presented at the 17th Int. Conf. Archit. Support Program. Languages Operat. Syst., Poster Session, London, U.K., Mar. 2012.

[16] L. Rizzon, M. Rossi, R. Passerone, and D. Brunelli, "Wireless sensor networks for environmental monitoring powered by microprocessors heat dissipation," in *Proc. 1st Int. Workshop Energy Neutral Sensing Syst.*, Rome, Italy, 2013, pp. 1–6.

[17] R. Rocha, J. Carmo, L. Goncalves, and J. Correia, "An energy scavenging microsystem based on thermoelectricity for battery life extension in laptops," in *Industrial Electronics, 2009. IECON'09. 35th Annual Conference of IEEE*, 2009.

[18] Nextreme thermal solutions, Thermomobility wireless power generator (WPG-1).

[19] Micropelt, TE-CORE6—TE-CORE7 ThermoHarvesting Power Module, Oct. 2012.

[20] pandaboard.org OMAP4 Pandaboard systemreference manual, Sep. 2010.

[21] Arndale Board [Online]. Available: http://www.arndaleboard.org/wiki/index.php/WiKi, 2013

[22] Thermoelectric handbook Laird Technol. [Online]. Available: www.lairdtech.com/temhandbook/

[23] M. Rossi and D. Brunelli, "Analyzing the transient response of MOX gas sensors to improve the lifetime of distributed sensing systems," in *Proc. 5th IEEE Int. Workshop Adv. Sensors Interfaces*, Jun. 2013, pp. 211–216.

[24] A. Somov, A. Baranov, A. Savkin, D. Spirjakin, A. Spirjakin, and R. Passerone, "Development of wireless sensor network for combustible gas monitoring," *Sens. Actuators A: Phys.*, vol. 171, no. 2, pp. 398–405, Nov. 2011.

[25] M. Rossi and D. Brunelli, "Ultra low power wireless gas sensor network for environmental monitoring applications," in *Proc. IEEE Workshop Environ. Energy Structural Monitor. Syst.*, Sep. 2012, pp. 75–81.

[26] M. Chen, L. A. Rosendahl, T. J. Condra, and J. K. Pedersen, "Numerical modeling of thermoelectric generators with varing material properties in a circuit simulator," *IEEE Trans. Energy Convers.*, vol. 24, no. 1, pp. 112–124, 2009.

[27] X. Niu, J. Yu, and S. Wang, "Experimental study on low-temperature waste heat thermoelectric generator," *J. Power Sources*, vol. 188, no. 2, pp. 621–626, 2009.

[28] E. Ortiz-Rivera, A. Salazar-Llinas, and J. Gonzalez-Llorente, "A mathematical model for online electrical characterization of thermoelectric generators using the PI curves at different temperatures," in *Proc. 25th Annu. IEEE Appl. Power Electron. Conf. Exposit.*, 2010, pp. 2226–2230.

[29] M. L. Anand and C. R. Ramesh, "Assessment and impact study of surface soil on thermoelectric generator energy efficiency," *Int. J. Sci. Eng. Res.*, vol. 4, pp. 980–983, Aug. 2008.

[30] M. Pollak, L. Mateu, and P. Spies, "Step-up DC-DC converter with coupled inductors for low input voltages," *Fraunhofer IIS*, 2008.

[31] P. H. Chou and S. Kim, "Techniques for maximizing efficiency of solar energy harvesting systems," in *Proc. 5th Conf. Mobile Comput. Ubiquitous Netw.*, Seattle, WA, 2010, vol. 10, pp. 26–28.

[32] C.-Y. Chen and P. Chou, "DuraCap: A supercapacitor-based, power-bootstrapping, maximum power point tracking energy-harvesting system," in *Proc. ACM/IEEE Int. Symp. Low-Power Electron.Design*, Aug. 2010, pp. 313–318.

[33] J. Olivo, D. Brunelli, and L. Benini, "A kinetic energy harvester with fast start-up for wearable body-monitoring sensors," in *Proc. 4th Int. Conf. Pervasive Comput. Technol. Healthcare (PervasiveHealth)*, 2010, pp. 1–7.

[34] D. Shang, X. Zhang, F. Xia, and A. Yakovlev, "Asynchronous design for new on-chip wide dynamic range power electronics," in *Proc. Conf. Design, Automat. Test Eur.*, 2014, p. 138.

[35] L. Mateu, C. Codrea, N. Lucas, M. Pollak, and P. Spies, "Energy harvesting for wireless communication systems using thermogenerators," in *Proc. XXI Conf. Design Circuits Integrated Syst.*, Barcelona, Spain, 2006, pp. 1–6.

[36] L. Mateu, M. Pollak, and P. Spies, "Analog maximum power point circuit applied to thermogenerators," *PowerMEMS 2008*, pp. 9–12, 2008.

Maurizio Rossi (S'12) received the M.Sc. degree in telecommunication engineering, in 2012, from the University of Trento, Trento, Italy, where he has been a Ph.D. degree student at the Information and Communication Technology International Doctoral School, since 2012.

His research interests includes the modeling and design of energy neutral embedded systems and wireless nodes for environmental monitoring.

Luca Rizzon received the B.Sc. degree and the M.Sc. degree in telecommunication engineering, in 2009 and 2012, respectively, from the University of Trento, Trento, Italy, where he is currently a Ph.D. candidate at the Information and Communication Technology International Doctoral School.

His research interests include the design of cyber-physical systems and design methods.

Matteo Fait received the B.Sc. and M.Sc. degrees in telecommunication engineering from the University of Trento, Trento, Italy, in 2011 and 2014, respectively.

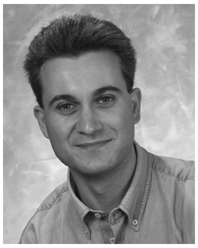

Roberto Passerone (M'97) received the M.S. and Ph.D. degrees in electrical engineering and computer sciences from the University of California, Berkeley, CA, USA, in 1997 and 2004, respectively.

He is Assistant Professor at the Department of Information Engineering and Computer Science, University of Trento, Trento, Italy. Before joining the University of Trento, he was Research Scientist at Cadence Design Systems. He has published numerous research papers on international conferences and journals in the area of design methods for systems and integrated circuits, formal models and design methodologies for embedded systems, with particular attention to image processing and wireless sensor networks.

Prof. Passerone was Track Chair for the Real-Time and Networked Embedded Systems at ETFA from 2008 to 2010, and General and Program Chair for SIES from 2010 to 2015. He was Guest Editor of the IEEE TRANSACTIONS ON INDUSTRIAL INFORMATICS for numerous special issues on embedded systems.

Davide Brunelli (M'10) received the M.S. (*cum laude*) and the Ph.D. degree in electrical engineering from the University of Bologna, Bologna, Italy, in 2002 and 2007, respectively.

He has been an Assistant Professor at the University of Trento, Trento, Italy, since 2010. From 2005 to 2007 he was a Visiting Researcher at ETHZ studying to develop methodologies for energy harvesting aware embedded design. He was scientific supervisor of several EU FP7 projects and he was leading industrial cooperation activities with Telecom Italia. His research interests concern smart grids and the development of new techniques of energy scavenging for wireless sensor networks (WSNs) and embedded systems, the optimization of low-power and low-cost WSNs, and the interaction and design issues in embedded personal and wearable devices.