

Inria Project Lab full proposal: ZEP: Zero Power computing systems

Kevin Marquet, Socrate Team, version of November 30, 2016

		1
1	Context and related works	2
1.1	Overview of the ZEP project	3
1.2	Related works	4
1.3	Partners	5
2	Scientific organization	5
2.1	Task 1: NVRAM-based Architecture	6
2.2	Task 2: Compiler Optimizations and Worst-Case Energy Consumption	8
2.3	Task 3: Runtime memory management	9
2.4	Task 4: Experimental platform development and building	9
2.5	Use cases	10
3	Administrative organisation	10
3.1	Organisation	10
3.2	Ongoing work on the themes of the IPL	12
3.3	Indicative requested resources for a 4-year project	12
4	Outcomes, Impact and risk	13
4.1	Outcomes and Impact of the project	13
4.2	Risk Analysis	14
5	Conclusion	14
	Bibliography	14

Abstract This proposal addresses the issue of designing tiny computing objects with no battery at all by combining non-volatile memory (NVRAM), energy harvesting, micro-architecture innovations, compiler optimizations, and static analysis. The main application target is Internet of Things (IoT) where small communicating objects will be composed of this computing part associated to a low-power wake-up radio system. The ZEP project will gather four Inria teams that have a scientific background in architecture, compilation, operating system and low power together with the CEA Lialp laboratory of CEA Leti. The major outcomes of the project will be a prototype harvesting board including NVRAM with demonstration running on it and the design of a new microprocessor associated with its optimizing compiler and operating system. Another important goal of the project is to structure the research and innovation that should occur within Inria to prepare the important technological shift brought by NVRAM technologies.

Teams involved

- INRIA Rhone Alpes: SOCRATE
- INRIA Rennes Bretagne Atlantique: CAIRN, PACAP
- CEA Leti
- CEA Lialp

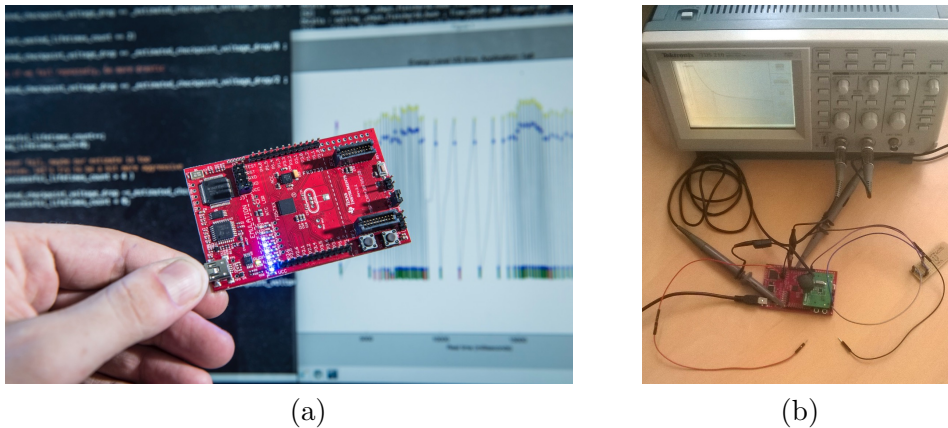


Figure 1: (a) The MSP-EXP430FR5739 board from Texas Instrument including 15kB of FRAM and (b) An experiment run by Socrate on intermittent power for this CPU (Photo ©Inria)

Keywords NVRAM, internet of things, energy harvesting, sensors, ultra-low power, micro-architecture, compilers, instant on-off and transient computing systems

1 Context and related works

The context for this proposal is the vision of ubiquitous computing, known as *Internet of Things (IoT)*: a scenario where we are surrounded by many intelligent objects continuously sensing their environment, talking to each other over the air, taking actions to react to the situation or providing information to end users. Application domains include environmental monitoring, smart clothes, building and home automation, automatic meter reading and control, industrial automation, health monitoring, smart on-street parking assistance, and many others¹.

Maintenance cost is a major bottleneck for ubiquitous computing. For example, if we have thousands of wireless sensors embedded in every building, the only task of changing batteries is already overwhelming. Moreover, many devices will be placed in hard-to-reach locations, inside walls for instance, where changing battery may be impossible. In addition, in many new application areas, sensors should be very small: with a size of the order of 1mm^3 . These are the main arguments for the current research activities on the use of **energy-harvesting** techniques to remove the need for battery.

Energy harvesting and ultra-low power make it possible to envision almost infinite lifetime. However, the energy collected is very low hence we need to save as much energy as possible by switching off the computing as often as possible. Even with sophisticated energy management techniques, the node will typically run out of power quite often, and the whole system has to expect frequent reboots. Existing compilers and operating systems are not designed for such a usage. Hardware manufacturers acknowledge this trend and they are beginning to produce new efficient **non-volatile memory** (NVRAM) chips, i.e. memory that does not lose its content when power is switched off. This technology is a major enabling factor for future low-power IoT devices. The ZEP project addresses the challenge of designing the **next-generation sensor platform computing parts** (hardware and software) organized around energy harvesting and non-volatile memory. It gathers researchers from various scientific fields: computer architecture, energy harvesting, operating systems and compilation. In addition, more and more actors recognize that NVRAMs will lead to huge changes in various fields of computer science. This project is a mean for INRIA teams to investigate NVRAM-based systems.

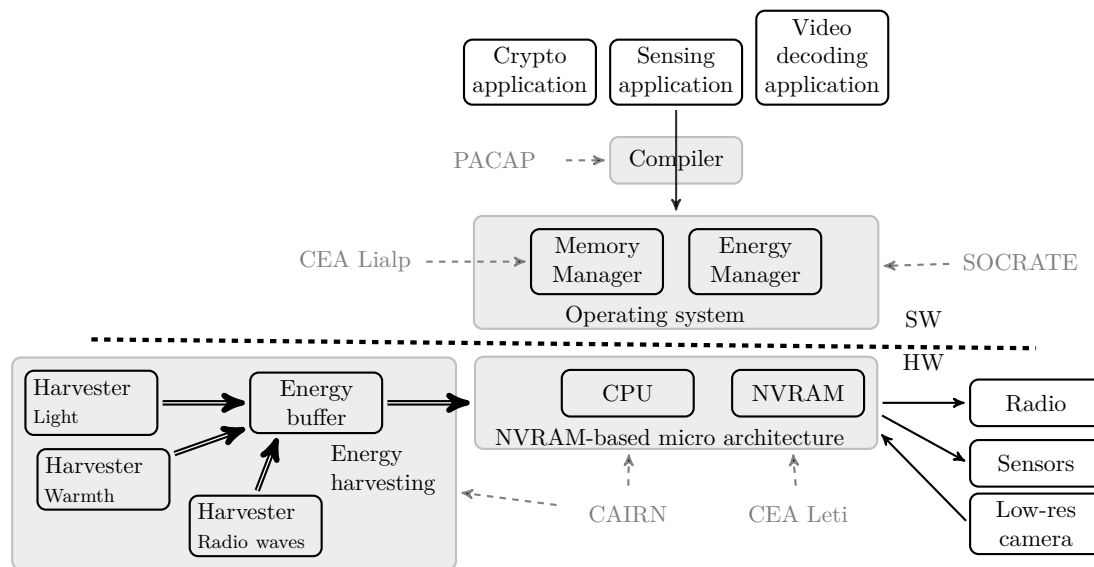


Figure 2: Overview of proposal

1.1 Overview of the ZEP project

The objective of the ZEP project is to provide communicating, sensing, computing systems able to work without battery and consuming orders of magnitude less energy than existing platform. In this goal, we will set up the new architecture sketched in Figure 2 where the key components targeted by ZEP, shaded on Figure 2, are the *energy harvesting system*, a *NVRAM-based sensing platform and micro-architecture* and the associated *operating system and compiler*. This figure exposes an *architectural view* of the project; the *scientific fields* will be detailed later in Section 2.

Energy harvesting system To power the whole platform, we will use a combination of different harvesting sources (light, heat, and radio waves), as well as a measurement infrastructure. This harvesting system will be based on the system proposed in the PowWow platform [5].

NVRAM-based sensing platform and micro-architecture The sensor platform will be studied both with a real prototype and with simulation:

- **A real hardware prototype** is necessary to convince industrial and scientists about the potential impact of the research. It is possible today: one can use for instance the last **FRAM MSP430 from Texas Instrument** (MSP430FR5739) which includes *Ferroelectric RAM* (FRAM, see our first experiment on Fig. 1). This platform embeds various sensors, and will be connected for instance to a low-definition camera and a radio-frontend. It will be integrated in a new version of the PowWow sensor in which an Actel Igloo FPGA, available for computing acceleration, also includes NVRAM. This new platform will constitute the **ZEP prototype**. A more advanced target will be the L-IOT prototype chip built by CEA Leti² and the ZEP project will work on NVRAM integration in L-IOT.
- **A simulation platform** will also be necessary as it is a very powerful tool to explore the design space composed with the many architectural choices and NVRAM technologies. We will use TI MSP430 simulation tool to build a complete platform simulator

¹L. Atzori, A. Iera, and G. Morabito. *The Internet of Things: A survey*. Computer Networks, 54, 2787-2805. 2010

²See <http://www-leti.cea.fr/fr/Prochain-evenement/2015/Leti-Days-Grenoble-2015>

for precise performance metrics measurement (power needed, energy consumption, computing power etc.). We will also have access to CEA-Leti NVRAM simulator models and L-IOT simulator to validate performance results.

NVRAM-based innovating micro-architecture We will explore new processor architectures to store the intermediate computation states before power failures occur, mainly through micro-architecture improvements and specific design techniques while taking into account nonvolatile memory devices already available at CEA (resistive RAMs). The objective is to propose a non-volatile architecture based on the RISC-V instruction-set architecture³ which can maintain its intermediate computation states during power failures and retrieve them when the power is recovered, thus achieving more forward progress.

Compiler In our approach, the compiler available for the chosen microprocessor (mspgcc for MSP430 on the ZEP platform or LLVM on ARM for L-IOT) will be completed with pre-deployment analyses that will provide the runtime system with information allowing it to manage resources more efficiently (in particular, worst case execution consumption WCEC). Moreover, we will develop specific optimizations and a new compiler linked to the NVRAM-based processor micro-architecture.

Operating system Although the term “Operating System” is a little too strong for this part of the work (“firmware” might be closer to what it will be), we still need to introduce low-level run-time mechanisms in order to handle many aspects related to NVRAM energy management, as well as making various applications to share resources.

1.2 Related works

There has been extensive research about energy-aware embedded systems but the scientific literature mostly focuses on two broad classes of devices. On one hand are mobile computing devices like smartphones and tablets. On the other hand are tightly resource-constrained platforms, such as wireless sensor networks or RFID tags. While both classes entail energy management considerations, the techniques proposed for each one of them are typically not relevant for the other.

A smartphone has plenty of resources: a large battery, a fast CPU, and large storage space. This makes it possible to propose heavyweight energy management techniques [23]. For example the LEAP architecture [20] advertised as “low-power” includes a micro-controller to measure power consumption of the main CPU. These approaches do not suit platforms powered by energy harvesting techniques which provide irregular, unreliable, low level of energy.

Harvesting techniques is in itself a huge research domain, usually handled by micro-electronics labs (such as CEA Leti for instance). The power-generating element must be selected after considering the type of energy to be collected from the surrounding environment, whether vibration, light, heat or radio waves. The most common elements used are solar, piezoelectric, or thermoelectric. It is also important that the power IC for use with the power-generating element efficiently collects the power from that element without loss, and that it supplies the stabilized power to a latter stage IC (see Table 1).

NVRAMs have been the focus of numerous recent works. A huge number of them aim at optimizing the mass storage system, typically by designing dedicated file systems [27][33]. Other work aim at providing orthogonal persistence, by using a persistent heap [26] [7, 10], checkpointing [25], but the vast majority of these works target big systems. Some work targeting embedded systems aim at optimizing dedicated applications [17], the few work that aim at allowing an energy harvesting system to survive power shortages use Flash memory [24].

³[The RISC-V Instruction Set Manual](#)

Energy source	Harvester	General Power	Output Voltage	Environnement of power generation
Light/Sun	Solar cell	$10mw/cm^2$	$> DC0.7CV(1cell)$	Outdoor (daytime)
		$10mw/cm^2$		Indoor (Nighttime)
Vibration	Piezoelectric	$500\mu W/cm^2$	$> AC80V$	Mechanical Vibration
		$10\mu W/cm^2$		Human action (press button)
Temperature	Thermoelectric	$0.5 W/cm^2$	several Volts	Thermal differentiation ($\Delta 100degC$)

Table 1: Generated power, output voltage, and the environmental factors that create energy differ for the common types of energy-harvesting technology⁴.

1.3 Partners

Interactions between hardware and software people are mandatory. Problems related to hardware architecture design allow software programmers to know whether a solution is realistic or not. Also, software people can evaluate hardware architecture *with respect to the use* of it. This is particularly true with the use of NVRAM [4]. The teams involved in this proposal (Socrate, Cairn, Alf, and CEA-Leti) have complementary skills that allow this kind of collaboration. In order to keep the project with a reasonable number of partner, It was decided not to include the Infine Inria team (which work on RIOT which targets mW micro-controller while we would target μW devices), the ZEP project will however keep regular contact with this team.

INRIA Socrate People of the team involved in the project are Kevin Marquet (MdC), Guillaume Salagnac (MdC) and Tanguy Risset (PR) will bring expertise on embedded systems design [19], communicating systems [8], operating systems and memory management [18], non-volatile memory [1] and wireless sensor networks [11, 12].

INRIA Cairn The Cairn team works on energy-efficient system-on-chip and reconfigurable architectures, signal processing for mobile communications, cryptography and wireless sensor networks. People of the team involved in the project include Olivier Sentieys (DR Inria/Irisa), Steven Derrien (PR U. Rennes), and Arnaud Carer (IR U. Rennes) who brings expertise on low-power architecture design [31, 22, 21] and low-power sensor networks [28] protocols [3, 2], energy harvesting [14, 15, 13, 16], compilers and design tools [30].

INRIA Pacap Pacap will be mainly in charge of static analysis to evaluate the energy consumption of basic code blocks and compiler optimization. Erven Rohou and Isabelle Puaut will provide expertise on compilation-time energy evaluation [29] and WCET computation [29].

CEA Lialp CEA Lialp will mainly work on the link between compilation and runtime system. Henri-Pierre Charles is a specialist of just-int-time compilation.

CEA Leti The Lialp laboratory of CEA Leti, headed by Jerome Martin, has been working for many year of ultra-low power system (especially with FDSOI technology) and has recently proposed the flexible L-IOT platform to prototype ultra-low power IOT applications (footnote 2, page 3). People involved in ZEP will be Edith Beigne and Ivan Miro-Panades.

2 Scientific organization

We now present scientific program of ZEP. It is organized along the four scientific fields in which we will make contributions: a NVRAM-based architecture, a energy harvesting system, a runtime system dedicated to efficient low-power energy and non volatile memory management, and a compiler combined with static analysis bringing information concerning energy consumption

to this runtime system and optimizing software code accordingly. Then we briefly present the use cases chosen to validate our results. Main PhD and Engineer advisers are indicated but of course many discussions will occur between all the teams of the ZEP project. In addition to this mainstream PhD task, Master students will be hired implying advisers, as much as possible, from two teams of the ZEP project.

2.1 Task 1: NVRAM-based Architecture

The rising demand for low-power sensors and autonomous IoT systems is a promising era of new embedded processors powered by ambient energy sources and considered as fully autonomous. The intrinsically unstable, low-power, and intermittent nature of ambient energy sources prevents the direct deployment of conventional processors designed under stable power assumptions for use in these emerging applications without the risk of frequent loss of progress caused by power failures. This challenge has driven recent efforts to explore nonvolatile processor designs that are not subject to loss of progress. These processors are not only integrating non-volatile memories but also contains non-volatile registers and gates.

NVRAM technology study: In order to reduce leakage power and make power failures transparent for the user, we need a platform with non-volatile memory. To that end, we will study various emerging non-volatile memory technologies (e.g. MRAM, FRAM, and RRAM in particular but maybe others too) from the point of view of embedded systems design. As illustrated by Table 2, numerous NVRAMs technologies are emerging, each one having its specifics that make it usable for some usages but not others. Existing literature on NVRAM is largely focused on high-performance systems (e.g. data centers, or high-end systems-on-chip) with very different constraints and objectives. In a resource-constrained context, energy consumption is more important than access latency or performance limitations. Moreover, considering device constraints will be mandatory in our research work. In its technology divisions, CEA researchers have already proposed different non-volatile devices like Ox-RAM, CBRAM or MRAM. Models are already available and some basic cells like flip-flops and bitcells have been developed. This work will be also driven by this expertise to always consider that the proposed non-volatile processor is feasible in terms of technology. It will also be a good opportunity for CEA to compare and evaluate which technology will be best suited for this kind of approach considering energy efficiency, performance and area.

NVRAM-based sensor node platform: A first and quickly available prototype development is easily accessible thanks to CAIRN previous work on the PowWow platform [5]. Replacing the PowWow micro-controller (currently an MSP430) by an FRAM-based MSP430FR5739 will directly provide an experimental platform on which the partner of ZEP will be able to experiment their development.

NVRAM-based processor architecture: But the real breakthrough will occur by studying the upgrade of current state-of-the-art low power micro-controller with NVRAM technology. In this project, we propose to collaborate between CAIRN and CEA to explore new techniques at architectural and design level to store the intermediate computation states before power failures occur, through architecture improvements, specific design techniques while taking into account nonvolatile memory devices already available at CEA (resistive RAMs). We will consider micro-architecture issues (e.g. non-volatile flip-flops), NVRAM-based caches, and study modifications of existing architecture according to various metrics. The final objective is to propose a non-volatile architecture which can maintain its intermediate computation states during power failures and retrieve them when the power is recovered, thus achieving more forward progress.

The core of the micro-architecture will be based on a RISC-V processor⁵. At CEA, a platform based on a RISC-V processor – the L-IOT architecture from the CEA Lialp (see Fig. 3) – is

⁵<https://riscv.org/>

	Emerging Memory			Established Memory		
	STTMRAM	PCM	RRAM	SRAM	DRAM	Flash NAND
Non-Volatile	YES	YES	YES	NO	NO	YES
Endurance (Nb cycles)	High (10^{12})	Medium (10^8)	Low (10^6)	High (10^{15})	High (10^{15})	Low (10^5)
2014 latest technological node produced (nm)	90 nm	45 nm	130 nm	10 nm	30 nm	15 nm
Cell size (cell size in F ²)	Medium (6-12)	Medium (6-12)	Medium (6-12)	Very large (150)	Small (6-10)	Very small (4)
Write speed (ns)	High (10 ns)	Medium (75 ns)	High (20 ns)	High (5-10 ns)	High (10 ns)	Low (10,000 ns)
Power consumption	Medium/low	Medium	Low	Very low	Low	Very high
2014 price (\$/Gb)	High (\$100-\$50/Gb)	Medium (few \$/Gb)	Very High (\$5,000/Gb)	Low (\$1/Gb)	Low (\$1/Gb)	Very low (\$0.05/Gb)
Suppliers	Everspin	Micron, Samsung	Adesto	Qualcomm, Intel	Samsung, Micron, SK Hynix	Samsung, Micron, Toshiba, SK Hynix

Table 2: Various emerging non-volatile RAMs and their characteristics. Drawbacks are in red, advantages in green. (Source: Emerging NVM Market, InMRAM workshop, Yole)

currently developed and will be updated according to this project’s results. RISC-V processor family is freely available and appears to be a very good candidate for IoT computing applications. Moreover, it is possible to connect many different accelerators which will probably be part of our proposed non-volatile architecture and design proposal. The RISC-V architecture moreover disposes of mature tools, allowing easy emulation of the platform, which is important for the software part. One important part of the work will be to study architectural tradeoffs regarding to the type of NVRAM considered. Indeed, the *universal* memory does not exist, all technologies considered today suffer from some drawbacks (*e.g.* latency, energy consumption, price...). These tradeoffs will be studied thanks to the extension of existing emulation tools that will take into account the specifics of various NVRAM technologies.

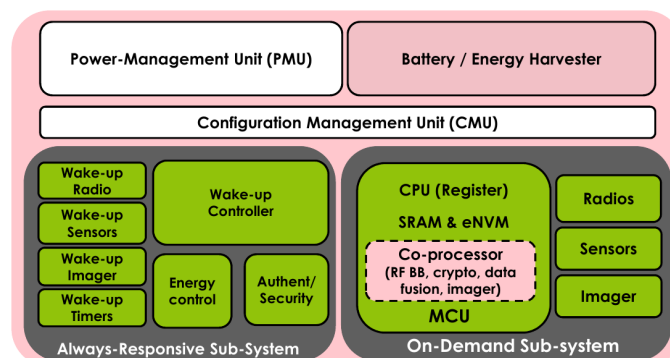


Figure 3: The L-IOT platform (extracted from <http://www-leti.cea.fr/en/content/view/full/2426>). The Always-Responsive Sub-System is used to wake up the on-Demand subsystem, NVRAM can be useful in both subsystems.

People involved: A PhD student will be supervised by Olivier Sentieys (CAIRN) and Ivan Miro Panadanés (CEA Lialp). This student will spend around 30% of its time in Grenoble at

the CEA, Grenoble. The organization envisioned today is that the first 6 months will be spent in Rennes, then 1 year in Grenoble, and the remaining time in Rennes, with video conferences every two weeks. Kevin Marquet (SOCRATE) will closely follow the work in order to give its *operating system* point of view and see the impact of the architecture on the software part. This PhD thesis will also be strongly linked to the work on the compiler described in Task 2. An Engineer for 1.5 year will build the ZEP prototype board, evolution of the PoWwoW board, and develop some of the use cases.

2.2 Task 2: Compiler Optimizations and Worst-Case Energy Consumption

The compiler can help the operating system to take decisions, by providing it with useful information such as the energy necessary to execute some pieces of code. Indeed, the amount of energy that the harvesting system will bring is very low. Some of the tasks that the platform will execute consume a lot of energy compared to the size of the energy buffer, for instance, sending a packet over the radio channel. If the energy level is too low, then the platform should not start such a computation, because it will not have the time to finish before power outage. If the operating system (OS) has indications on whether or not a given computation can be launched, it will optimize power failure.

Adapted from classical WCET (worst-case execution time) analysis, a compilation pass will be implemented, that computes a worst-case *energy consumption* (WCEC). This analysis will contain variables that will be defined by the runtime system because environmental conditions can change a lot the consumption estimation.

Energy Model: We anticipate that the energy model of the processor will be simple. Probably, the energy consumed will consist of two components: one proportional to the number of executed instructions, and another one proportional to the number of memory accesses. In the presence of NVRAM, memory accesses may behave as asymmetric: writes are more expensive than reads. This opens interesting opportunities to the code generator. Outside the processor itself, each device will be described by its own characteristics, such as energy needed to send a radio packet, etc.

Code Generation: We propose a WCEC-driven code generation. The compiler will decompose the code into fragments, separated by “check-points” such that a fragment is guaranteed to complete when it starts executing. For that purpose, each fragment will be annotated by the amount of energy necessary to reach the next check-point. This value may also be communicated to the OS. In case case, the OS will avoid scheduling the task until enough energy becomes available. By definition, WCEC estimates are pessimistic. We will propose a mechanism to query the battery level at runtime to assess the remaining slack, and let a task proceed if it can guarantee to reach another check-point.

Checkpoints: We anticipate that some registers may be implemented in non-volatile memory. The code generator could decide to save some values to non-volatile registers. When the code reaches a checkpoint, the runtime can decide, based on the remaining energy level, either to proceed, or to save only what is necessary. We plan to investigate trade-offs between saving data sporadically during execution, and saving larger chunks of data at checkpoints.

Quality of Service and Approximate Computing: We will also explore ways to reduce energy consumptions when the battery runs low:

- the code generator will provides several instances of some critical functions, using different levels of precision. As an example we might consider the same computation using types `double`, `float`, or some reduced-precision floating-point or fixed-point representation.
- guided by the programmer, the code generator may also rely of reduced quality, for example by processing fewer frames per second, reducing the resolution, etc. The most appropriate version will be selected at run-time based on dynamic information, such as battery level, but also domain-specific constraints.

People involved: 1 PhD student will work in Rennes, supervised by Erven Rohou from PACAP and Steven Derrien from CAIRN, on the static analysis. CAIRN will bring its expertise in energy consumption evaluation, while PACAP will work on the static analysis. The task will also involve Isabelle Puaut from PACAP and Olivier Sentieys from CAIRN.

2.3 Task 3: Runtime memory management

We will design a software stack able to manage efficiently the specific hardware architecture. We will build memory management techniques dedicated to persistent memory. First, check-pointing (a way to retain data at a given time) will allow us to answer the following questions. When is it necessary to save data? Which data is it necessary to save? How to ensure data structure consistency while reboots are expected in an unpredictable manner ?

Several challenges will be addressed by the OS :

Provide support to the NVRAM architecture The memory hierarchy will include RAM + NVRAM, as well as some mechanisms dedicated to this organisation. For instance, a DMA dedicated to the transferts between DRAM and NVRAM, or a protection unit dedicated to NVRAM, or programmable cache ? These mechanisms must be supported by the OS. Also, multi-applicative support will be necessary in constrained embedded systems, because they may be embedded in an inaccessible place. Therefore, the runtime system will have to share the non-volatile memory across the various applications, according to their needs and priorities.

Choose how energy is used The OS is responsible for choosing how energy is used. It cannot be the responsibility of the user because 1) energy is not exposed in programming languages 2) the amount of energy available is not known when programming. But if there are several applications, given an energy budget, an energy policy must define the way the system should prefer one or another.

Ensure memory consistency The fact that part of the RAM is not non-volatile is far from solving the problem of ensuring memory consistency. Three problems are to be solved: 1) processor state is volatile 2) devices' states are volatile 3) when several applications execute concurrently, a priority between saving data of one or another must be established;

Refine worst-case execution The worst-case execution time depends on several dynamic parameters: variables of the program, temperature, primitives of the OS... It is necessary that the OS reconsider the WCEC with respect to these parameters that only itself knows.

A memory manager and an energy manager will be implemented in the OS and will allow to apply memory and energy policies in order to address these problems. An important part of the project will be to evaluate tradeoffs, in terms of performance and energy consumption, between (i) Storing data in NVRAM a given amount of time, (ii) Send it by radio with various duty cycles and (iii) process data and only send results rather than raw data.

People involved: 1 PhD student will work in Lyon, supervised by K. Marquet, T. Risset (Socrate) and H-P Charles from CEA Lialp. The design of the architecture by the CEA will be followed by K. Marquet in order to ensure the adequation between the architecture and the OS.

2.4 Task 4: Experimental platform development and building

There are two sub-tasks here: we need to associate this board with harvester that is adapted to each targeted application and we have to provide demonstrators for the chosen use cases (see Section 2.5). We will evaluate the power level provided by different harvesting technologies in

order to know if it is sufficient to power, for instance, a biomedical sensor. For example, it might be necessary to add energy harvesting from radio waves in order to communicate with the sensor and obtain its data, in the same manner as RFIDs. We will provide power consumption models and behavior for each harvesting source, as well as a study of the energy sources necessary to power to power our platform, *with respect to the application*.

People involved: The engineer in Rennes mentioned in Task 1 will participate to this task and one engineer will work with Socrate during one year for integrating the demonstrators on the platform chosen.

2.5 Use cases

Three use cases will be taken to validate the work on the platform:

Sensing Smart sensor are foreseen to appear for very many applications in metering: gaz metering, pollution, light, parking lots, etc. Following Socrate experience on smart metering [32, 9]. A generic sensing application will be provided parameterized by the characteristics of the chosen application: duty cycle, message sent, radio power, etc.

Cryptography no wireless communicating system should be designed without a solid cryptographic subsystem. But encrypting is time- and computation-consuming, and therefore necessitates energy. We will evaluate how different encryption/decryption system can effectively execute on the platform, based on previous works of the Socrate team [6].

Video decoding The potential of video monitoring is huge but also very energy-consuming. Regarding to the energy harvesting mean (solar, heat, etc.), the available power provides some strong and interesting challenges. Recent study on low resolution (video monitoring) shows that it might be possible to switch off the sensor between two recorded images and optimize other algorithmic parameters.

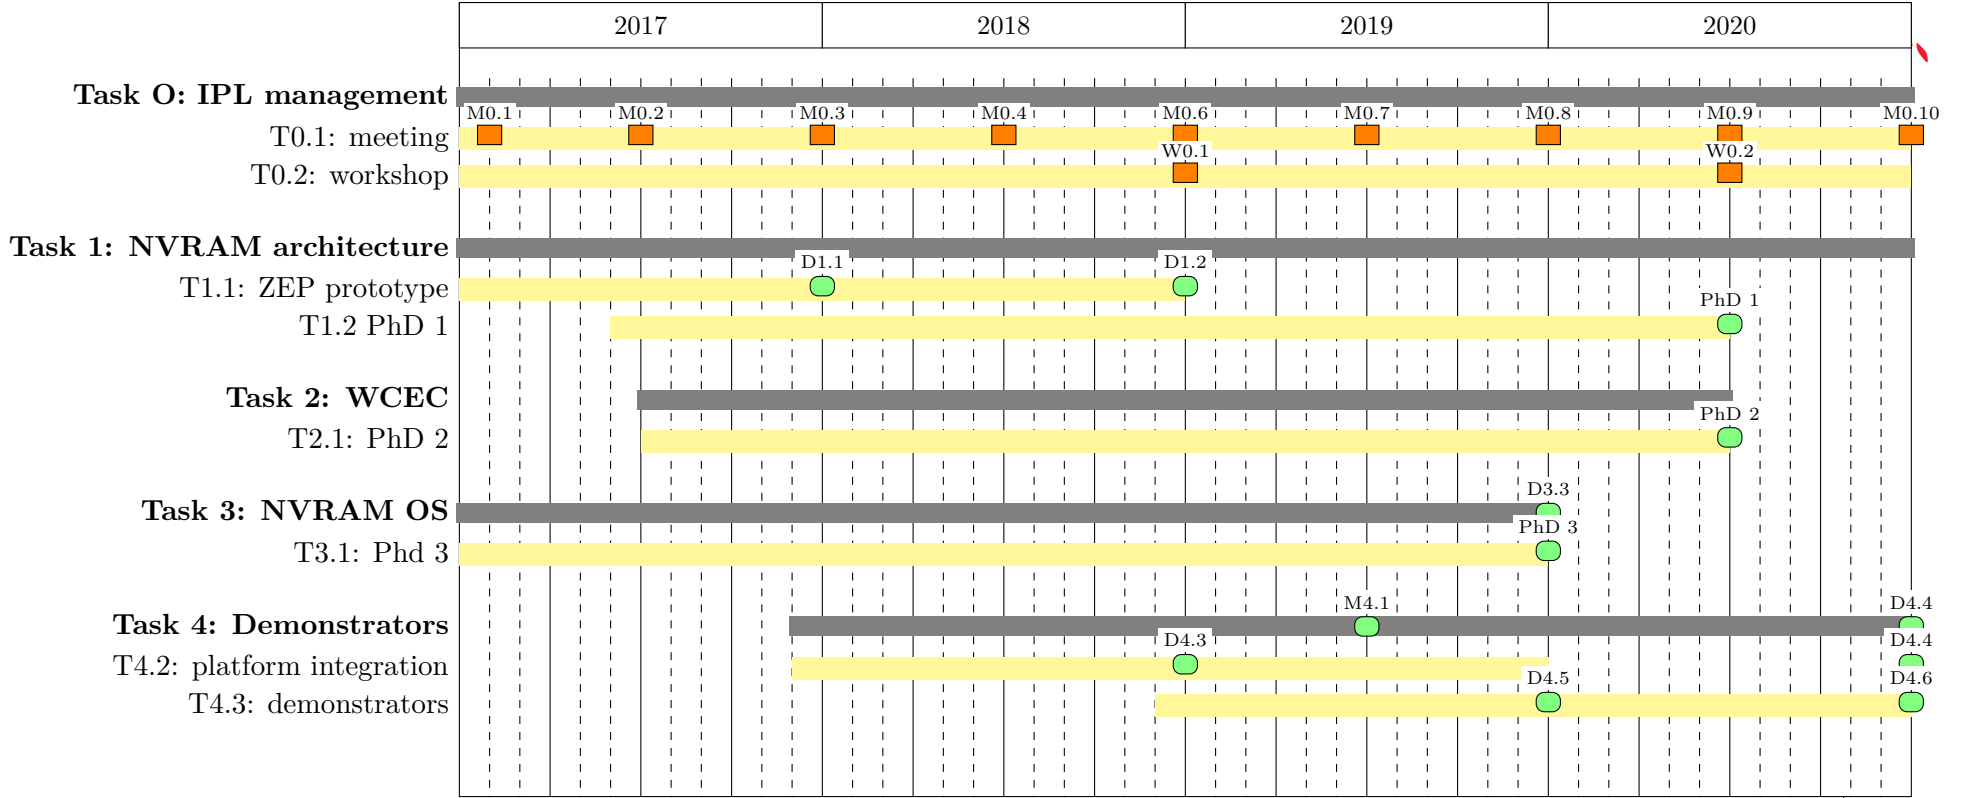
3 Administrative organisation

3.1 Organisation

The project is led by Kevin Marquet of Socrate, INRIA teams involved know each other very well, and Socrates and Cairn have long-standing collaboration with the laboratory of CEA Lialp. In addition to the tasks mentioned in Section 2, Socrate will be responsible for the overall management of the project with in particular:

1. A scientific meeting every six months where partners will present informally their work and joint decisions will be addressed.
2. Two official workshops will be organized during the project, it is planned to target “GDR day” and a workshop in an international conference.

The Gannt show on page 11 shows the temporal evolution of the tasks over the 4 years.



3.2 Ongoing work on the themes of the IPL

- In the SOCRATE team, a PhD student funded by the *Région Région Rhône-Alpes* (ARC6) will start designing an operating system for a NVRAM-based architecture. The hardware part will be defined by a spin-off of the CEA, and will be very different from the architecture of this IPL as it is likely to provide high-computational level (notably: ARM CPU and hundred of megabytes of NVRAM).
- In the SOCRATE team, the operating system currently developed thanks to the ADT Sytare will serve as a basis for the OS of this IPL. We will port it to the new platform, add the support for the WCET, the energy monitoring, etc.
- The cryptographic use cases will be taken from previous work of the team [6].
- In the CAIRN team, the PowWow platform, which includes energy harvesting features, will be extended to provide a prototyping platform to the whole project.
- The CEA Lialp has already designed the L-IOT platform, which will be used as a basis for the NVRAM-based architecture.
- Existing static analysis designed in the PACAP team for the evaluation of worst-case execution time will be modified to compute worst-case energy consumption.

3.3 Indicative requested resources for a 4-year project

Budget repartition

- 1 PhD student supervised by Pacap (and Cairn) on compiler and WCEC;
- 1 PhD student on primitive operating system, supervised by Socrate (and CEA Lialp);
- 1 PhD student on the NVRAM-based architecture, supervised by Cairn and CEA-Leti;
- 12 month engineer to integrate all demonstrators with Socrate.
- 18 month engineer to build the ZEP prototype with Cairn;
- Master students (2 per years)
- Equipment: 30.000 euros 5 personal computers, 10 TI MSP430FR5739 platforms, 10 radio-wave harvesters, PowWow V2 platform design and fabrication (20 nodes).
- Operating costs: 50.000 euros (conference fees, travel expenses for PhDs and permanents).

Other funding associate to ZEP topic Most of the teams involved in the project are already working on the NVRAM topic. Socrate in particular started in September a PhD funded by the *Région Région Rhône-Alpes* (ARC6) which intends to design an operating system for a NVRAM-based architecture with high power but low frequency activation. Socrate has also benefited from the ADT Sytare which started the development of an operating system that will serve as a basis for the OS of this IPL. Socrate and Cairn also participated to the (rejected) GRIOTE H2020 Ecsel proposal which hopefully will lead to other projects.

The CAIRN team, with the PowWow platform, has a good expertise in designing low-power board and already worked with energy harvesters. The Pacap team has an international appreciation in studying WCET (Worst Case Execution Time) theory and will apply it to provide WCEC (Worst Cases Execution Consumption). Socrate, CEA Lialp, Pacap and Cairn all have good background in compilation and LLVM compiler in particular.

4 Outcomes, Impact and risk

4.1 Outcomes and Impact of the project

We expect the project to validate the improvements provided by NVRAM technology and give guidelines for industry and research use it efficiently. As major actors of the digital world such as DELL, IBM, HP, Fujitsu, Oracle and others, are now recognizing⁶, it is clear that non-volatile memory will imply major changes in computer systems in the next few years. To the best of our knowledge, at time of writing, the prototype that we intend to build has no equivalent in the world but there are certainly many similar ideas in construction.

The definition of a completely new NVRAM-based processor architecture would be very long, and would need to design a new industry-ready compiler dedicated to this architecture, and would potentially introduce the need for completely new software mechanisms that are outside the scope of this IPL, but we hope to help CEA and ST-Microelectronics in designing ultra-low power sensor in a way to have efficient software stack as proposed by Inria.

After 4 years, the project should have achieved the following:

A platform prototype The ZEP platform will be built on the basis of the PowWow [5] platform designed and developed by the Cairn team, including a solution to harvest energy from light and heat. An RFID antenna, able to get energy from radio waves, will be added.

An innovating processor architecture Depending on the result of Task 1, a more advanced platform will possibly be delivered in collaboration with CEA.

A simulator for a generic NV-RAM platform to explore different solutions and the impact of each solution. This is mandatory because it is not possible to build a platform for each architecture. This simulator will be based on existing works by the Socrate team (WSIM) or existing available micro-controller simulator associated with CEA-Leti simulator for NV-Ram blocks.

A compiler leveraging WCEC analysis to provide the new NVRAM-based processor architecture with specific optimizations and a new compiler, a completed with pre-deployment analyses that will provide the runtime system with information allowing it to manage resources more efficiently, in particular, worst case execution consumption (WCEC).

A compilation pass providing the runtime system with information on the energy consumption of basic code blocks. The pass will be implemented inside the LLVM infrastructure taking advantage of Pacap, Socrate and CEA Lialp background on LLVM compiler.

A primitive operating system including *i*) NVRAM management techniques *ii*) power management policies. The software will be based on previous work of the Socrate [1] and Cairn[22] teams in this field. The integration in an embedded OS such as RIOT, OpenWSN or Contiki is out of the scope of ZEP but it is an important topic that will be investigated as soon as the basic components of the platform will be available.

Demonstrators It is important to have demonstrators on several use cases to convince industrial and investors that this NVRAM technology provides a real breakthrough.

In addition, other direct consequences of ZEP will include:

⁶<http://www.snia.org/about/news/newsroom/pr/snia-announces-non-volatile-memory-nvm-programming-technical-work-group>

Publications, Workshop and conference On each scientific topic detailed above, we will work on the emergence on dedicated workshop in important conferences (DATE for instance) on this thematic which is currently studied mostly from the electronic point of view.

Lobbying, communication, new research direction An important goal of the ZEP proposal is to have INRIA emerging as a major actor of the technological revolution of NVRAM. A constant lobbying in that direction will be operated by the members of the ZEP project (European project, National initiative: ANR, GDR, international publication, ST-microelectronics and SMEs etc.). We hope to open perspective for Inria to open research teams or start-ups that will define tomorrow's IOT (hardware and) software.

4.2 Risk Analysis

The main risk of the project is related to the design of the NVRAM-based platform because the design space is enormous. Our approach is to mitigate the risk of designing a useless or unrealistic platform by:

- starting from the existing, convincing L-IOT platform of the CEA;
- working on a specific but realistic scenario in the field of constrained embedded devices. We do not claim that we will build the universal NVRAM-based CPU. But we hope to make one step in the direction of this kind of NVRAM-based platform.

Also, research and prototypes related to operating systems can be tedious and lead to lots of time lost in pure engineering. Our approach concerning this field should mitigate this risk because we won't try to adapt an existing . Still, we hope that one interesting result of the project will be to gain experience on the changes in the software part required to manage such a NVRAM-based architecture. This should be the work of a dedicated master student at the end of the project : attempting to port the mechanism designed for our prototype OS in a real OS like RIOT or FreeRTOS.

Another risk is related to the Worst-Case Energy Consumption. While the static part does not worry us to much because it will be based on existing works, it is uncertain what we will . We hope the project will at least allow to evaluate to which extent this scientific track is promising.

5 Conclusion

Our project intends to prototype next-generation of sensors, based on energy harvesting and NVRAM management. These sensors will be able to live without batteries, which is a strong requirement for some applications of the IoT. In particular, non-volatile memories will increasingly change the computing device paradigm and this project investigates specific challenges related to that.

Bibliography

- [1] Fayçal Ait Aoudia, Kevin Marquet, and Guillaume Salagnac. Incremental checkpointing of program state to NVRAM for transiently-powered systems. In *ReCoSoC - 7th International Workshop on Reconfigurable Communication-centric Systems-on-Chip*, Montpellier, France, June 2014.
- [2] Mahtab Alam, Olivier Berder, Daniel Menard, and Olivier Sentieys. Tad-mac: traffic-aware dynamic mac protocol for wireless body area sensor networks. *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, 2(1):109–119, March 2012.
- [3] M.M. Alam, O. Berder, D. Menard, and O. Sentieys. Latency-energy optimized mac protocol for body sensor networks. In *Ninth International Conference on Wearable and Implantable Body Sensor Networks (BSN)*, pages 67–72, May 2012.
- [4] Katelin Bailey, Luis Ceze, Steven D. Gribble, and Henry M. Levy. Operating system implications of fast, cheap, non-volatile memory. In *13th Workshop on Hot Topics in Operating Systems, HotOS XIII, Napa, California, USA, May 9-11, 2011*. USENIX Association, 2011.

- [5] Olivier Berder and Olivier Sentieys. PowWow: Power Optimized Hardware/Software Framework for Wireless Motes. In *Workshop on Ultra-Low Power Sensor Networks (WUPS), co-located with Int. Conf. on Architecture of Computing Systems (ARCS 2010)*, pages 229–233, Hannover, Germany, February 2010.
- [6] Mickaël Cazorla, Kévin Marquet, and Marine Minier. Survey and Benchmark of Lightweight Block Ciphers for Wireless Sensor Networks. In *SECRYPT 2013 - Proceedings of the International Conference on Security and Cryptography*, pages 543–548, Reykjavik, Iceland, July 2013. SciTePress.
- [7] Joel Coburn, Adrian M. Caulfield, Ameen Akel, Laura M. Grupp, Rajesh K. Gupta, Ranjit Jhala, and Steven Swanson. Nv-heaps: Making persistent objects fast and safe with next-generation, non-volatile memories. In *Proceedings of the Sixteenth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS XVI*, pages 105–118, New York, NY, USA, 2011. ACM.
- [8] M. Dardaillon, K. Marquet, T. Risset, J. Martin, and H. P. Charles. A compilation flow for parametric dataflow: Programming model, scheduling, and application to heterogeneous mp soc. In *Compilers, Architecture and Synthesis for Embedded Systems (CASES), 2014 International Conference on*, pages 1–10, Oct 2014.
- [9] ETSI. *Electromagnetic compatibility and Radio spectrum Matters (ERM) ; System Reference document (SRdoc) : Spectrum Requirements for Short Range Device, Metropolitan Mesh Machine Networks (M3N) and Smart Metering (SM) applications*. Sept 2011, TR 103 055.
- [10] Taeho Hwang, Jaemin Jung, and Youjip Won. Heapo: Heap-based persistent object store. *Trans. Storage*, 11(1):3:1–3:21, December 2014.
- [11] Raja Jurdak, Peter Corke, Alban Cotillon, Dhinesh Dharman, Chris Crossman, and Guillaume Salagnac. Energy-efficient localization: Gps duty cycling with radio ranging. *ACM Trans. Sen. Netw.*, 9(2):23:1–23:33, April 2013.
- [12] Matthieu Lauzier, Antoine Fraboulet, Jean-Marie Gorce, and Tanguy Risset. Live group detection for mobile wireless sensor networks. In *Ninth International Conference on Body Area Networks*, London, Great Britain, September 2014.
- [13] T.N. Le, A. Pegatoquet, Olivier Berder, and Olivier Sentieys. A Power Manager with Balanced Quality of Service for Energy-Harvesting Wireless Sensor Nodes. In *International Workshop on Energy Neutral Sensing Systems (ENSsys) - ACM SenSys Conference*, pages 19–24, Memphis, United States, November 2014.
- [14] Trong-Nhan Le, Michele Magno, Alain Pegatoquet, Olivier Berder, Olivier Sentieys, and Emanuel Popovici. Ultra low power asynchronous mac protocol using wake-up radio for energy neutral wireless sensor networks. In *International Workshop on Energy-Neutral Sensing Systems (ENSsys) organized in conjunction with 11th ACM SenSys Conference*, page Paper 10, Rome, Italy, November 2013.
- [15] Trong-Nhan Le, Alain Pegatoquet, Olivier Berder, and Olivier Sentieys. Energy-Efficient Power Manager and MAC Protocol for Multi-Hop Wireless Sensor Networks Powered by Periodic Energy Harvesting Sources. *IEEE Sensors Journal*, 15(12):7208 – 7220, December 2015.
- [16] Trong Nhan Le, Alain Pegatoquet, Olivier Berder, Olivier Sentieys, and Arnaud Carer. Energy Neutral Design Framework for Supercapacitor-based Autonomous Wireless Sensor Networks. *ACM Journal on Emerging Technologies in Computing Systems*, 12(2):19:1–19:21, August 2015.
- [17] X. Luo, D. Liu, L. Liangy, Y. Li, K. Zhong, and L. Long. Mobilock: an energy-aware encryption mechanism for nvram-based mobile devices. In *Non-Volatile Memory System and Applications Symposium (NVMSA), 2015 IEEE*, pages 1–6, Aug 2015.
- [18] Kevin Marquet and Gilles Grimaud. A dsl approach for object memory management of small devices. In *Proceedings of the 5th International Symposium on Principles and Practice of Programming in Java, PPPJ '07*, pages 155–164, New York, NY, USA, 2007. ACM.
- [19] Kevin Marquet and Matthieu Moy. Pinavm: A systemc front-end based on an executable intermediate representation. In *Proceedings of the Tenth ACM International Conference on Embedded Software, EMSOFT '10*, pages 79–88, New York, NY, USA, 2010. ACM.
- [20] Dustin McIntire, Thanos Stathopoulos, Sasank Reddy, Thomas Schmidt, and William J. Kaiser. Energy-efficient sensing with the low power, energy aware processing (leap) architecture. *ACM Trans. Embed. Comput. Syst.*, 11(2):27:1–27:36, July 2012.
- [21] Adeel Pasha, Steven Derrien, and Olivier Sentieys. System level synthesis for wireless sensor node controllers: A complete design flow. *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, 17(1):2:1–2:24, January 2011.
- [22] M. A. Pasha, S. Derrien, and O. Sentieys. A complete design-flow for the generation of ultra low-power wsn node architectures based on micro-tasking. In *Design Automation Conference (DAC), 2010 47th ACM/IEEE*, pages 693–698, June 2010.
- [23] Abhinav Pathak, Y. Charlie Hu, and Ming Zhang. Where is the energy spent inside my app?: Fine grained energy accounting on smartphones with eprof. In *Proceedings of the 7th ACM European Conference on Computer Systems, EuroSys '12*, pages 29–42, New York, NY, USA, 2012. ACM.
- [24] Benjamin Ransford, Jacob Sorber, and Kevin Fu. Mementos: System support for long-running computation on rfid-scale devices. In *Proceedings of the Sixteenth International Conference on Architectural Support for Programming Languages and Operating Systems, ASPLOS XVI*, pages 159–170, New York, NY, USA, 2011. ACM.
- [25] Jinglei Ren, Jishen Zhao, Samira Khan, Jongmoo Choi, Yongwei Wu, and Onur Mutlu. Thynvm: Enabling software-transparent crash consistency in persistent memory systems. In *Proceedings of the 48th International Symposium on Microarchitecture, MICRO-48*, pages 672–685, New York, NY, USA, 2015. ACM.
- [26] David Schwalb, Tim Berning, Martin Faust, Markus Dreseler, and Hasso Plattner. nvm malloc: Memory allocation for nvram. In Rajesh Bordawekar, Tirthankar Lahiri, Bugra Gedik, and Christian A. Lang, editors, *ADMS@VLDB*, pages 61–72, 2015.

- [27] P. Sehgal, S. Basu, K. Srinivasan, and K. Voruganti. An empirical study of file systems on nvm. In *2015 31st Symposium on Mass Storage Systems and Technologies (MSST)*, pages 1–14, May 2015.
- [28] Olivier Sentieys and Olivier Berder. *Energy Autonomous Micro and Nano Systems*, chapter Optimizing Energy Efficiency of Sensor Networks, pages 325–360. Wiley-ISTE, 2012.
- [29] Cristina Silvano, Giovanni Agosta, Andrea Bartolini, Andrea R. Beccari, Luca Benini, João Bispo, Radim Cmar, João M. P. Cardoso, Carlo Cavazzoni, Jan Martinovič, Gianluca Palermo, Martin Palkovič, Pedro Pinto, Erven Rohou, Nico Sanna, and Kateřina Slaninová. AutoTuning and Adaptivity appRoach for Energy efficient eXascale HPC systems: the ANTAREX Approach. In *Design, Automation, and Test in Europe*, Design, Automation, and Test in Europe, Dresden, Germany, March 2016.
- [30] Timo Stripf, Oliver Oey, Thomas Bruckschloegla, Juergen Becker, Gerard Rauwerda, Kim Sunesen, George Goulas, Panayiotis Alefragis, Nikolaos S. Voros, Steven Derrien, Olivier Sentieys, Nikolaos Kavvadias, Grigoris Dimitroulakos, Kostas Masselos, Dimitrios Kritharidis, Nikolaos Mitas, and Thomas Perschke. Compiling Scilab to high performance embedded multicore systems. *Microprocessors and Microsystems: Embedded Hardware Design (MICPRO)*, 37(8):1033–1049, November 2013.
- [31] V. D. Tovinakere, O. Sentieys, S. Derrien, and C. Huriaux. Low power reconfigurable controllers for wireless sensor network nodes. In *IEEE 22nd Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*, pages 230–233, May 2014.
- [32] Mathieu Vallérian, Guillaume Villemaud, Benoit Miscopein, Tanguy Risset, and Florin Hutu. SDR for SRD: ADC specifications for reconfigurable gateways in urban sensor networks. In *Radio and Wireless Symposium (RWS), 2014 IEEE*, pages 178–180. IEEE, 2014.
- [33] Haris Volos, Sanketh Nalli, Sankarlingam Panneerselvam, Venkatanathan Varadarajan, Prashant Saxena, and Michael M. Swift. Aerie: Flexible file-system interfaces to storage-class memory. In *Proceedings of the Ninth European Conference on Computer Systems*, EuroSys '14, pages 14:1–14:14, New York, NY, USA, 2014. ACM.