R. Reussner, A. Koziolek, R. Heinrich (Hrsg.): INFORMATIK 2020,
Lecture Notes in Informatics (LNI), Gesellschaft für Informatik, Bonn 2021 923

A Decade of Energy Awareness Technology Evolution for Sensor Nodes

Marcus Thoss[1]

Abstract: Energy awareness is an important aspect of the design of sensor node hard- and software, particularly for battery-powered or energy-harvesting node architectures. Architectural design choices of such systems must regard a multitude of aspects, including size, weight, memory and processing power. Therefore, energy-related design aspects have in recent years become a feature that is being honoured throughout sensor node design. As a result, various technological solutions and strategies have evolved to facilitate energy awareness and energy management aspects.

This paper looks back at the evolution of sensor node technology during the 2010s. In the course of research activities, the state of the art of research and industrial solutions aiming at improving the support of energy awareness was being monitored. Advances were observed for various aspects and levels of relevant technological facets, including electronic measurement and control circuitry, harvesting facilities, power-saving mechanisms at both hard- and software level, energy management strategies and algorithms, networking aspects, and advances and extensions related to operating systems for sensor nodes. A conclusion of these observations, given in this paper, identifies technological increments, leaps and sidesteps that have occurred along the way.

For the overall time span of the decade observed, a short qualitative and quantitative analysis of the technological advances achieved is presented, including typical examples of actual sensor node designs. The paper concludes with an outlook on further evolution of advances in energy awareness technology for sensor nodes to be expected in the near future and to be desired in the long run.

Keywords: Energy Awareness; Sensor Node; Technology Review

1 Motivation

The Internet of Things (IoT) is being rolled out with increasing speed and in a multitude of application areas with competing and complementing communication technologies like 5G networks, massive satellite fleets, LoRA, WPANs and, ultimately, the classical Internet connections. Yet, current reports like Microsoft's "IoT Signals" [Mi19] clearly show that there is not only a strong willingness to adopt IoT technologies and many platforms to choose from, but also a huge percentage of projects (30% according to the report) failing to reach deployment.

It can be assumed that IoT scenarios relying on widely distributed small-scale self-powered nodes, with or without energy harvesting, are rather common. Engineering of this type of

[1] RheinMain University of Applied Sciences, Department of Design – Computer Science – Media, Unter den Eichen 5, 65195 Wiesbaden, Germany, marcus.thoss@hs-rm.de

 doi:10.18420/inf2020_84

device requires a focusing on the aspects of size and weight, node lifetime, connectivity and maintenance. All of those benefit from a better handling of the energy resources available, sometimes in an indirect manner, as when connectivity is being improved by designing a system around the goal to offer the power to transmit just when an application requires it.

When [Th10] was published 10 years ago, the IoT was still in its infancy. Looking back at this point in time, the decade that has passed since then is referred to in this paper as “the decade”. A precursor of the IoT had been the previous 10-year long increase in RFID technology usage during the 2000s. Subsequently, especially in Europe, numerous research programmes for then-topical Ambient Assisted Living were funded, multi-national research cooperations like the European Research Cluster on the Internet of Things[2] began to assemble and the first platform solutions like UniversAAL [Fe15] emerged.

From then on, interest in sensor network technology, both among industry and researchers, boomed, and technological advances, some of which will be described here, helped to increase reliability, applicability to more domains, and efficiency of the solutions. Besides connectivity improvements, both at sensor network and internet-based cloud levels, an increase of the regard for energy awareness issues was a major factor of the overall development observed.

2 Technological Facets of Energy Awareness

Design measures and system components contributing to (or hindering) energy awareness can be found throughout any systems architecture. Therefore, challenges and the potential of different technological facets should first be looked at separately to be considered for an overall, system-wide design strategy for energy awareness.

2.1 Energy Sources

Energy sources can be rated by many factors, the most relevant here being voltage range and maximum power output. They can be further classified into storage (primary and secondary cells, capacitors and solid state storage) and on-line sources (wired or harvesting sources).

The necessity to rate energy consumed by sensor nodes based on a cost factor, which decreases from primary cells to harvesting solutions, had been observed at least by the beginning of the decade [Th10]. With regard to cost as a general concept, rechargeable storage components can be neglected as they do not change the overall energy balance, energy harvesting supplies can be tagged with a no-cost factor, wired power sources add a moderate, constant cost, and primary cells would be considered the most expensive form of energy source both in terms of supply logistics and environmental impact.

[2] `http://www.internet-of-things-research.eu`

Consequently, the overall development over the decade moved from primary cells to swappable secondary cells to energy harvesting approaches. Sadly, many industry solutions are still deliberately designed to use primary cells because their lifetime often exceeds that of the application, and cost per cell is, at least monetarily, cheap.

As to the technology of harvesting sources, the most interesting development happened and is to be further expected among solar cell technology. With some cell types now nearing an efficiency of 30%, while 25% had been a good rating by the beginning of the decade (according to [Gr10], the report has been updated regularly since), size of the cells for a given node uptime could be decreased and harvesting nodes can now be utilised in a much broader range of environments.

2.2 Energy Storage

Storage-type energy sources introduce additional parameters like capacity, possibly recharge characteristics, wear-out, size, weight and operating temperature range. Regarding purely economic motivations, secondary cells of the Li-Ion and related types have become favoured over the decade, as small form factors like coin cells were available and the technology had advanced to allow energy harvesting to provide enough recurring energy input for many applications. The other trend – that did not gain as much momentum – is the use of supercapacitors as rechargeable storage element. From an environmental perspective, the use of carbon-electrode supercapacitors would be highly preferable compared to the more problematic secondary cell technologies often employing scarce or toxic substances. The problems of still minor capacities and more challenging output voltage curves, although partially outweighed by a high number of possible recharge cycles, leaves this technology so far a candidate for future advances.

2.3 Power Controllers

The evolution of microcontroller technology has led to single-supply voltage solutions, as modern controllers are capable of internally generating themselves diverse voltages required. Although this had been the case by the beginning of the decade already, input voltage ranges now become broader and single-supply architectures are even more common.

On the supply side, power management becomes more complex when storage and harvesting strategies are included. Here, technological evolution has produced an increase in efficiency of conversion and the availability of highly integrated power controllers. These can manage multiple harvesting and wired sources, intermediate energy storage, output voltage and power control, and require few external components. Thus, the complexity of node hardware designs could be reduced dramatically, while efficiency was increased.

2.4 Microcontrollers

In the case of a sensor node, the processor is almost certain to be a microcontroller. A consolidation of instruction set architectures could be observed among semiconductor manufacturers from a multitude of proprietary designs towards classical Intel 8051 Cores for 8-bit processors and ARM architectures covering 32-bit. Notable exceptions are Microchip (including former Atmel), remaining successful with their 8-bit PIC and AVR architectures, and Texas Instruments with their 16-bit MSP430 processor family.

Microcontrollers aid energy aware designs by the fact that power domains of processor core and peripherals are closely coupled and compatible, reducing transfer and conversion losses. A notable improvement has been introduced by controllable power gating of embedded peripherals and an increase of detail in sleep state modelling, both regarding the number of levels and the granularity of control that can be exerted.

2.5 Memory

The choice of RAM type for low-energy designs is of course static RAM offering negligible idle-state power losses. Thus, driven by the need for more complex sensor node firmware, RAM size offers have been increased from a few KB to some 100 KB without greater penalty, as relevant power drains occur only during RAM access, which is a matter of limiting the amount of code execution.

Not yet frequently available in common architectures are the follow-up non-volatile, near-RAM-access-speed technologies FRAM and MRAM, with FRAM being available in some MSP430 variants as the only mainstream platform. A major disadvantage of SRAM as main memory is the loss of information in power-off and deep sleep states. Since zero-power sleeping is a desirable state, though, energy aware designs must currently apply measures to persist or re-initialise RAM contents during these states. Should non-volatile memory technologies become the standard for main memory designs, powerless sleeping would become available with less overhead and complexity. In [Re19], energy consumption measurements for various sleep states and application aspects were analysed.

2.6 Energy Measurement

Energy measurements can basically be taken with an external meter or, the node can measure its own energy flows. In any case, voltage and current must be determined to calculate a momentary $P = U \cdot I$ and then integrate it over a time span Δt to get $E_{\Delta t} = \int_{t_0}^{t_0+\Delta t} P dt$. For external measurements, the problem of mapping samples taken to the network node's time domain poses a challenge. So, although external measurements can often afford more precise metering equipment, internal measurements are desirable, i.e. the note would measure itself.

Thus, initially, analog input pins of the microcontroller were, and still are, used to measure battery and operating voltage, and a shunt resistor is inserted in the supply path to determine the current by measuring the voltage drop created.

Over the decade, energy aware node designs emerged that were equipped with dedicated on-board power metering solutions, improving accuracy and sometimes allowing the microcontroller to sleep while measurements continue. Yet, such additional circuitry would consume power itself, degrading node uptime and falsifying measurement values.

There is a kind of “holy grail” of such measurement solutions beginning to be realised more frequently. It is the monitoring of the embedded switched-mode power supply (SMPS) circuits managing the supply voltages of the node, as noted in Section 2.3. An SMPS performs demand-based regulation by transferring chunks of energy from a supplier to the consumer. Thus, the amount of energy transferred is proportional to the number of chunks converted. The trick applied here is to have the microcontroller itself count the switching events by adapting the pulses through appropriate decoupling and probing circuitry.

Such solutions promise to minimise measurement overhead, both in terms of node circuit complexity and energy investment. As most modern microcontrollers are capable of counting events in deep sleep modes, even the requirement of continuous monitoring can be fulfilled.

2.7 Operating Systems

There has been a notable evolution of embedded operating systems (OS) meeting sensor network technology requirements. Foremost, strong connectivity support is obviously a necessity. Secondly, to be applicable to a broad range of projects, multi-platform support, also scaling from low-end architectures for small sensor nodes to those found in rather well-equipped router and gateway nodes, would contribute to the popularity of a sensor network OS.

Energy awareness did not seem to be of primary concern by the beginning of the decade, and has actually begun to be fully integrated into the OS architectures and APIs only recently. To some extent, the open source platforms Contiki OS [DGV04] and Tiny OS [Le05] pioneered the sensor node research community, and Contiki considered on-line, OS support for energy measurements early on, providing the important feature of attributing measurement data with the execution context of the measurement. During the decade, the open source embedded OS scene saw Mbed OS (2009), Apache MyNewt (2015), Linux Foundation’s Zephyr (2016), and only very recently, RIOT (2018) appear. Mbed is being driven by the ARM company, targeting their own architectures, while Texas Instruments has provided their own closed-source solution TI RTOS for their hardware platforms, notably the MSP430 family. They all show activities towards providing better support for energy awareness, with RIOT having a strong research background forstering its energy awareness aspects [Ha16].

2.8 Testbeds

By the middle of the decade, the survey published in [HH14] showed that a large number of approaches had accumulated, and researchers and industry had been busy investigating sensor networks through testbeds. The authors of the survey themselves had been part of the community from around 2010 on, and their SANDbed setup [HWM10] explicitly focused on energy monitoring. Testbeds continue to evolve regularly within the communities developing platform solutions, a recent, notable example being the RIOT testbed [GEN19].

Basically, there are three methods for assessing energy budgets, that is, model-based prediction (optionally based on previous off-line measurements), and internal and external on-line measurement. With an approach reasonably limiting the overhead of the measurement action, on-line measurements of the actual operation of the sensor network would be preferable. It requires a measurement infrastructure, though, that is able to retrieve and collect the measurement data and offers sufficiently comfortable testbed setup, experiment control and evaluation support. These requirements are likely to be the reason why many solutions still rely on off-line measurements and simulated or modelled energy assessments.

3 Example Designs

3.1 Harvesting Iris Mote

An early, rather simplistic energy aware harvesting sensor node solution was presented and analysed in 2011 in [RMT12], providing all basic ingredients based on the Iris mote platform from MEMSIC. The design and the analysis paper anticipate many problems and approaches that have been discussed in the community in the following years. The researchers added a solar cell, supercapacitor and power management to the base platform, considered energy measurement by counting switched mode power supply events and looked at the pros and cons of simulated off-line versus measurement-based on-line assessment strategies.

3.2 WieDAS

In this research project, the sensor node box and its solar panel measure 8 cm by 12 cm each, being voluminous by today's standards. The size was caused by the use of AAA Batteries and a single-layer carrier board for prototype production. Still these nodes were successfully operated in an AAL demonstrator scenario described in [Kr14]. Energy awareness was initially limited to battery voltage measurement using the main microcontroller and attaching a current probe in the supply line. A major step towards power awareness was the addition of a TI INA219 current sensor. This device is capable of determining power consumption measuring both voltage current and even averaging power values over programmable window sizes. Thus, with constant sampling frequency f and $N_{samples}$ per window, energy transfers can can be approximated by $\Delta E = P \cdot \Delta t$ with $\Delta t = N_{samples} \cdot \frac{1}{f}$.

Adding this external measurement component allowed leaving the microcontroller in deep sleep modes without stopping measurements, if the CPU was woken up in time to collect buffered measurement data. One drawback was the late addition of solar energy harvesting, which did not match the inflexible power domain scheme consisting of jumpers that had to be removed manually to disconnect single sensors of the board from the power supply.

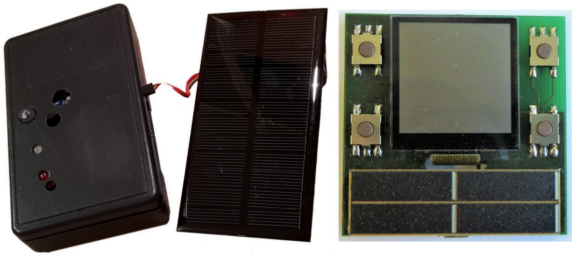

Fig. 1: WieDAS (left) and BASE MOVE (right) nodes compared

3.3 BASE MoVE

By the end of the decade, the BASE MoVE project [Ak20] had created custom sensor nodes, too, which measured less than half the size of a WieDAS node, but included comparatively tiny solar cells and a large display (shown in Figure 1, nodes are not to scale). The project created node hardware with programmable power gating of sensor peripherals, energy harvesting realisations for solar and thermal sources, and an automated testbed integrated with build automation for regression tests. Automated distributed energy measurements and energy management, that had been planned, finally exceeded the scope of the project.

4 Conclusion and Outlook

Within the given context, the project LONG MOVE[3], a successor to BASE MoVE, focuses on node-local and distributed energy management topics for sensor networks. Its mission statement calls for “energy awareness by design”, taking this aspect into account from the beginning on. It thus reflects the insight that previous designs could have benefited from earlier, consistent involvement of energy-related questions. With OS and hardware platforms offering better support in this area, it can be expected that this will become an overall trend. In fact, energy awareness could be an integral part of designs when self-adaptation features embedded at platform level can obliterate to some degree the need to regard energy within the design of the application proper.

[3] https://www.hs-rm.de/de/fachbereiche/design-informatik-medien/forschungsprofil/long-move

Bibliography

[Ak20] Akelbein, J.-P.; Beckmann, K.; Hoss, M.; Schneider, S.; Seyfarth, S.; Thoss, M.: BASE MoVE - A Basis for a Future-proof IoT Sensor. In: INFORMATIK2020. Gesellschaft für Informatik, Bonn, 2020. to be published.

[DGV04] Dunkels, A.; Gronvall, B.; Voigt, T.: Contiki - A Lightweight and Flexible Operating System for Tiny Networked Sensors. In: 29th Annual IEEE International Conference on Local Computer Networks. pp. 455–462, 2004.

[Fe15] Ferro, E.; Girolami, M.; Salvi, D.; Mayer, Ch.; Gorman, J.; Grguric, A.; Ram, R.; Sadat, R.; Giannoutakis, K.; Stocklöw, C.: The UniversAAL Platform for AAL (Ambient Assisted Living). Journal of Intelligent Systems, Jan 2015.

[GEN19] Günes, M.; Engelhardt, F.; Nothnagel, K.: Technical report - Designing a Testbed for Wireless Communication Research on Embedded Devices. In: 18. GI/ITG KuVS FachGespräch SensorNetze, FGSN 2019. pp. 41–44, 2019.

[Gr10] Green, M. A.; Emery, K.; Hishikawa, Y.; Warta, W.: Solar Cell Efficiency Tables (v. 35). Progress in Photovoltaics: Research and Applications, 18(2):144–150, 2010.

[Ha16] Hahm, O.: Enabling Energy Efficient Smart Object Networking at Internet-Scale : Experimental Tools, Software Platform, and Information-Centric Networking Protocols. PhD thesis, University of Paris-Saclay, Dec 2016.

[HH14] Horneber, J.; Hergenröder, A.: A Survey on Testbeds and Experimentation Environments for Wireless Sensor Networks. IEEE Communications Surveys & Tutorials, 16(4):1820–1838, April 2014.

[HWM10] Hergenröder, A.; Wilke, J.; Meier, D.: Distributed Energy Measurements in WSN Testbeds with a Sensor Node Management Device (SNMD). In: Workshop Proceedings of the 23th International Conference on Architecture of Computing Systems. VDE Verlag, pp. 341–438, February 2010.

[Kr14] Kröger, R.; Lux, W.; Schaarschmidt, U.; Schäfer, J.; Thoss, M.; von Fragstein, O.: The WieDAS AAL Platform: Architecture and Evaluation. In: 7. Deutscher AAL-Kongress mit Ausstellung, 21.-22. Jan. 2014, Berlin. VDE Verlag GmbH, Berlin/Offenbach, 2014.

[Le05] Levis, P.; Madden, S.; Polastre, J.; Szewczyk, R.; Whitehouse, K.; Woo, A.; Gay, D.; Hill, J.; Welsh, M.; Brewer, E.; Culler, D.: TinyOS: An Operating System for Sensor Networks. In: Ambient Intelligence. Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 115–148, 2005.

[Mi19] Microsoft Corporation: IoT Signals. Technical report, Microsoft Corporation, 2019.

[Re19] Reinhard, J.: Energiebetrachtungen für Sensorknoten mit FRAM (Energy analysis of FRAM-equipped sensor nodes). Lab report, Embedded Systems Course 2018/19, Rhein-Main Univerity of Applied Sciences, 2019.

[RMT12] Renner, C.; Meier, F.; Turau, V.: Holistic Online Eenergy Assessment: Feasibility and Practical Application. In: Ninth International Conference on Networked Sensing, INSS 2012, Antwerp, Belgium, June 11-14, 2012. IEEE, pp. 1–8, 2012.

[Th10] Thoss, M.: Supporting Energy Awareness in Distributed Embedded Systems. In (Dörner, Ralf; Krömker, Detlef, eds): Self Integrating Systems for Better Living Environments: First Workshop, Sensyble 2010. Shaker Aachen, pp. 117–124, Nov 2010.