Possibilities for software development for energy-limited constrained devices

Mario Hoss and Jens-Peter Akelbein

Hochschule Darmstadt, University of Applied Sciences
{mario.hoss, jens-peter.akelbein} @h-da.de

Abstract. With the spread and rising complexity of IoT scenarios there are also new challenges emerging in planing and predicting the lifetime of application specific energy-limited resource constrained devices. To allow for a state of the art software development process, energy consumption as well available energy input will need to be considered during early stages of development. Depending on what information are available or predictable during which stage of the product lifecycle, adaptive behavior could also be used to supplemented or compensated for predictions. This paper presents an early stage work into a new research topic. It gives a review of the possibilities and challenges of predicting and determining energy consumption and input information throughout development and deployment of constrained devices. Contrasting existing approaches from related fields, the paper concludes in the outlook on the upcoming research questions.

Keywords: Energy Models · Adaptive Systems · IoT

1 Introduction

According to forecasts of exponential growth in the Internet of Things market [1], the demand for application-specific connected devices is expected to reach billions. This growth is driven by the emergence of a variety of new application scenarios, like Industrial IoT (IIOT) and Smart Cities. Many of these scenarios rely on masses of energy-limited constrained devices [2] that wirelessly transmit collected information for up to a decade or more without being connected to a permanent power supply. This poses new challenges in the design and development phases as it is critical for both battery operated and energy harvesting powered devices to optimize the power consumption of hard- and software to ensure that they can reach their targeted lifespan.

With more efficient power-saving modes for a System on a Chip (SoC), the potential complexity of the application has increased drastically. Until recently, commonly only class 0 and class 1 devices [2] have been used in energy-limited scenarios with such long lifespans. Both classes are not intended to use standard internet protocols. By now, class 2 devices are also a realistic option. Since common internet protocols are often used, state-of-the-art security is also required. Providing the necessary functionalities further increases software complexity, as well as resource- and power-consumption. This increased complexity is also a major driver in the increased use of operating systems for the IoT, which should accelerate further.

This shift will likely be accompanied by a change in requirements for the developer. While writing bare-metal applications requires a more in-depth understanding of the underlying hardware, using operating systems providing abstraction layers and reusable libraries is suitable for more general software developers.

In light of these developments, there is a need to provide software developers with tools and methods to accurately predict the life expectancy of a constrained device in the development phase. Optimally, developers should be able to select combinations of hardware, software and energy-source and immediately receive a life expectancy prognosis. Such a forecast supports early development phases and allows for an "energy by design" approach in product development. Developers are able to test and customize their software for different energy-sources, for example different illuminance levels for solar powered devices. Using such forecasts in test and verification allows automated detection of energy bugs so software defects that cause abnormal power consumption.

Whilst mobile device applications face similar problems, solutions in this area are generally not applicable to the field of energy-limited constrained devices due to vastly different hardware as well as energy consumption pattern. Current lifetime prediction research for constrained devices originates in the Wireless Sensor Network (WSN) field with assumptions for less complex software, stronger hardware bindings, and more complex network communication. Like in the are of mobile devices, research for constrained devices also focuses on adaptive in-situ behavior rather than using a predictive approach during development. Existing approaches also utilize simplified battery models or additional hardware for measurements in the field.

Verifying the functionality and life expectancy by measuring power consumption during HIL testing is also not a realistic option. In addition to the to be discussed practical problems, achieving a sufficient test coverage is problematic. With a life expectancy of a decade, there is just not enough time during a normal development phase to sufficiently test the devices behavior over this time span.

This paper explores the possibilities for energy predictions of energy-limited application specific constrained devices in product development and its synergies with adaptive in-situ functionalities. As such it will be a first look into early research and a first step towards sufficient tooling for developers of energy-limited constrained devices.

2 State of the art

2.1 Energy consumption

In order to predict energy consumption, it is generally necessary to determine both the CPU power consumption of individual instructions as well as hardware components. Because of the differences in usage scenarios and underlying hardware, energy consumption at an instruction level is much more relevant on mobile devices [3] then on constrained devices. Energy-limited constrained devices can also easily accumulate small inaccuracies, because unlike mobile devices, they are not regularly recharged. However, they do spend most of their lifetime in deep sleep modes, interrupted by repeating usage patterns, which reduces the overall complexity of the required models. Furthermore, due to significant difference in available resources, mobile devices can track the various hardware states in the field and are also able to measure the remaining battery charge [4], which due to per device cost is rarely an option for constrained devices. For instruction level models, determining data dependent dynamic power consumption of the worst case energy consumption was recently shown to result in a NP-hard problem, where an approximation cannot be made to a usable degree [5]. The examined class 2 constrained devices showed a variation in the data dependent power consumption of nearly half of a cores power dissipation. Another study also documented SoCs with a similar ratio [6]. More relevant for constrained devices are finite state machines (FSM) where the target hardware is modeled by different power states and their connecting transitions, determined through a measurement cycle. The parameters that influence the power consumption of a state and their behavior can then be identified by using regression analysis to create approximation functions for the parameter-dependent energy consumption of peripheral devices [7]. Transition triggers are identified by power bursts, since power states often do not line up with the utilization in software [8]. Automatically creating and refining such state machines is still an open research question [9]. A more detailed overview for both types of models can be found in earlier work [10].

There is also the established practice of using power consumption measurements during Hardware in the Loop (HIL) testing [11]. Woehrle et al. [12] validated WSN nodes utilizing HIL tests and the testbed Flocklab [13] also supports automated power tests for WSN nodes. This however puts practical limits on the test coverage, since the constrained devices are designed to run for years, making the coverage minimal in a normal development time frame. "Wearables" face this problem to a lesser degree. "Rocketlogger" is designed for in-situ measurements of the energy consumption through a normal wearable lifecycle of a few days [14].

One aspect that complicated past work on energy consumption models is the consideration of production irregularities and a variance in energy consumption in different power modes [15, 16]. These variances are usually considered in energy consumption models by factoring in an error margin during the creation. The variance has to be considered for any measurements on the actual hardware, be it HIL tests or in the creation of energy consumption models.

Energy-aware software engineering is used for mobile devices. This includes the identification of energy bugs and hotspots [17] and static analyses [18], which can also be used for general software development [19]. Energy consumption models and measurements are also utilized for compiler optimization [20].

2.2 Energy input

There is ample work on models for the remaining battery capacity [21], but the ambient temperature needs to be considered for more accurate results [22]. Battery capacity approximation also exists for energy harvesting powered constrained devices [23]. To predict the energy input of solar cells, both static and dynamic factors have to be considered. The static input consists of the efficiency of the cell for the different light sources and their respective wavelength. Information about the cells efficiency is readily available from the manufacturer and can be verified by confirmation measurements published in the biannual solar cell efficiency tables [24]. The dynamic input is dependent on environmental conditions during runtime, for example on the available illuminance level. A list of expected influencing factors indoors is listed in figure 1. A list for factors influencing the efficiency of photovoltaic systems in general can be found in [25]. Weather reports have also been used for short-term [26] and long-term [27] predictions of the energy input of solar powered constrained devices operating outdoors.

Positioning (direct)	Positioning (indirect)	Hardware properties	Hardware condition
Light composition	Ambient temperature	Efficiency of the solar cell at different wavelength	Cell age
Light level	Partially shaded cells	Efficiency of energy-harvesting IC	Dust and dirt
Cell orientation and angle	Weather and user interaction (e.g. shades)	Availability of maximum power point tracking	Hardware failure

Fig. 1. Factors influencing the energy input of solar powered constrained devices

Sufficient accurate data-sets for the prediction of indoor light levels in living- and functional buildings are currently lacking. Existing architectural information can not be used, as the field focuses on the indoor light level relative to the outdoor light level [28]. Consumer protection studies [29] can be useful but generally focus on living spaces. This information gap will be filled over the course of the LOEWE3 project "LONG MOVE" with continuous measurements over 18 month in a wide variety of functional buildings. Important aspects are consideration of both incoming sunlight over different seasons, as well as artificial light sources and information about the placement of the devices.

2.3 Energy management and adaptive functionality

Energy management functionalities are typically used to track and manage hard to predict power consumption in the field of mobile devices. For mobile devices, this is largely caused by the impact user interaction and usage pattern have on the device power consumption. Such an approach was adopted by Tamkittikhun et al. [30] for solar powered embedded devices. Different to mobile devices, direct user interaction is rarely possible for these devices. However, since user intervention is not an option to achieve lifetime goals through regular charging, adaptive behavior to reduce functionality or increase sleep durations is often necessary. They are also faced with hard to predict power consumption in the form of processing incoming transmissions. For this Tamkittikhun et al. combined power measurements of individual functionalities before deployment with a functionality counter during runtime to estimate energy consumption on the device. This allows for lifetime predictions, as well as adaptation of the device functionality to meet lifetime goals. Since energy-limited constrained devices achieve their lifetime goal by remaining in sleep states, the transceiver is also powered off. As such energy consumption caused by processing incoming transmissions rarely has to be considered.

For energy-limited constrained devices there are also research activities into adapting the device behavior depending on the available energy [16,31,32]. Here the prediction problems often stem from hardware variances as well as the dependence of the energy-input on the environment.

Lachenmann et al. [31] proposed a programming abstraction for constrained devices, associating different sets of functionalities with different energy levels. For that the energy consumption of different functionalities is measured during development and associated with these levels. Depending on the remaining charge, the devices then adapt their behavior in the field to a functionality level with a lower energy consumption. For this a battery monitor hardware and a battery model mapping the voltage to the remaining battery capacity was used. Sieber [16] followed a similar approach but focused on more dynamic performance adjustments in the field based on an energy bucket concept in addition to the energy consumption model. For that application and system functionalities are defined with an energy priority and individual energy consumption limits. Instead of the use of measurement hardware on the device, the remaining battery charge is estimated based on a model of linearized remaining charge values based on measurements taken of the battery before deployment. This approach is motivated by profiting off and coping with production variances. Geissdoerfer et al. [32] utilized local energy input predictions on energy harvesting powered devices during run time, simultaneously taking into account the batteries current state of charge. The underlying model gets adjusted through state of charge feedback to the predictions.

In general, it is possible to identify several different approaches to adaptive energy management. There are those solutions that collect additional information about their environment, be it light intensity, temperature or state of charge, and act on that information. This often requires additional hardware, resulting in increased cost and software overhead for each device. Other approaches [16,30] do not collect environmental information at runtime, and are instead shipped with, and act on, energy-consumption- and -input-models created before deployment. These models can come in a wide variety of forms, from battery-models to simply measuring the power consumption of specific functions or features. While these approaches do not require additional hardware, they still produce software overhead on each device. However, if this approach is applied to constrained devices rather than to more complex embedded devices, then it should be possible to move large parts of the computation into the development phase, reducing the required overhead or even making the solution completely application-neutral. Lastly, there are managed devices that can communicate with a management server. In addition to the possibility to report environmental conditions and offload costly model based computations, it is also possible to provide a device with additional information, like weather forecasts to predict the energy input of solar powered devices [26]. This is however achieved by additional overhead for the management of the device as well as running costs for the management infrastructure. If however predictions about the necessity of "over the air" firmware updates come true, the additionally required overhead on the device would be minimal. Current management options for constrained devices, like "Lightweight Machine to Machine" (LWM2M) [33], have not found wide acceptance, likely due to the introduced overhead [34]. However, the current IETF draft for "Software Updates for the Internet of Things" (SUIT) [35] could provide such management options in the long term future, though likely first for devices with a permanent power supply. Another advantage of this approach is, that with the energy consumption model centralized in one location, the model could be adapted and refined over the lifetime. The categories listed are not mutually exclusive. Hybrid approaches like Lachenmann et al. [31] rely on both energy consumption models and additional measurement hardware on the device.

3 Outlook

This paper summarizes the literature review giving a comparison of current approaches. Based on this review, three research questions are derived to be answered in the next steps of research: *Q1:* How can energy consumption models be utilized for simulating power consumption of energy-limited constrained device? *Q2:* How can energy-input predictions be utilized for simulating the energy-input variability of energy-limited constrained devices? *Q3:* How can energy-input- and -consumption-information be utilized for lifetime prediction during the software development process for application specific constrained devices?

To find answers to these questions the design science research process is used by creating and evaluating proof of concept prototypes iteratively. These allow to determine the usability and precision of the individual models as well as forecasts based on a combination of both models. As a result, it will be possible to determine the prediction accuracy of each element within the forecast depending on the phase of the development process. Identifying the limitations of the approach allows defining where adaptive in-situ behavior can address deviations found over long time periods. This leads to determining a minimum requirement profile.

Based on these answers, it should become possible to further evaluate how the presented approach of forecasts can be integrated in todays product development processes.

This work was supported by the Hessenagentur within the Project "LONG MOVE - Leistungsadaption und ortsbezogene Verhaltensregeln fr eine nachhaltige IoT-Sensorik in der Gebudeausstattung zur modularen Vernetzung von Einheiten" (HA-Projekt-Nr. 802/19-122).

References

1. Middleton, P.: Forecast analysis: Internet of things endpoints, worldwide, 2017 update. Tech. rep., Gartner Research (2017), https://www.gartner.com/en/documents/3841268/forecast-analysis-internet-of-things-endpoints-worldwide
2. Bormann, C., Ersue, M., Keranen, A.: Terminology for constrained-node networks. RFC 7228, RFC Editor (May 2014), http://www.rfc-editor.org/rfc/rfc7228.txt
3. Shukla, N.K., Pila, R., Rawat, S.: Utilization-based power consumption profiling in smartphones. In: 2016 2nd International Conference on Contemporary Computing and Informatics (IC3I). pp. 881–886 (Dec 2016). https://doi.org/10.1109/IC3I.2016.7919046
4. Dong, M., Zhong, L.: Self-constructive high-rate system energy modeling for battery-powered mobile systems. In: Proceedings of the 9th International Conference on Mobile Systems, Applications, and Services. pp. 335–348. MobiSys '11, ACM, New York, NY, USA (2011). https://doi.org/10.1145/1999995.2000027
5. Morse, J., Kerrison, S., Eder, K.: On the limitations of analyzing worst-case dynamic energy of processing. ACM Trans. Embed. Comput. Syst. **17**(3), 59:1–59:22 (Feb 2018). https://doi.org/10.1145/3173042
6. Ascia, G., Catania, V., Palesi, M., Sarta, D.: An instruction-level power analysis model with data dependency. VLSI Design **12**(2), 245–273 (2001). https://doi.org/10.1155/2001/82129
7. Friesel, D., Buschhoff, M., Spinczyk, O.: Parameter-aware energy models for embedded-system peripherals. In: 2018 IEEE 13th International Symposium on Industrial Embedded Systems (SIES). pp. 1–4 (June 2018). https://doi.org/10.1109/SIES.2018.8442096
8. Pathak, A., Hu, Y.C., Zhang, M., Bahl, P., Wang, Y.M.: Fine-grained power modeling for smartphones using system call tracing. In: Proceedings of the Sixth Conference on Computer Systems. pp. 153–168. EuroSys '11, ACM, New York, NY, USA (2011). https://doi.org/10.1145/1966445.1966460
9. Buschhoff, M., Friesel, D., Spinczyk, O.: Energy models in the loop. Procedia Computer Science **130**, 1063 – 1068 (2018). https://doi.org/10.1016/j.procs.2018.04.154, the 9th International Conference on Ambient Systems, Networks and Technologies (ANT 2018) / The 8th International Conference on Sustainable Energy Information Technology (SEIT-2018) / Affiliated Workshops
10. Hoss, M., Westmeier, F., Akelbein, J.: Low cost high resolution ampere meter for automated power tests for constrained devices. In: Bleimann, U., Burkhardt, D., Humm, B., Löw, R., Regier, S., Stengel, I., Walsh, P. (eds.) Proceedings of the 5th Collaborative European Research Conference (CERC 2019), Darmstadt, Germany, March 29-30, 2019. CEUR Workshop Proceedings, vol. 2348, pp. 81–88. CEUR-WS.org (2019)

11. Waeltermann, P.: Hardware-in-the-loop: The technology for testing electronic controls in vehicle engineering. Tech. rep., dSPACE (2016), https://www.dspace.com/files/pdf1/dspace-paper_hil_overview_waeltermann_e_160405.pdf
12. Woehrle, M.: Testing of Wireless Sensor Networks. Ph.D. thesis, TU Karlsruhe (2010). https://doi.org/10.3929/ethz-a-006206723
13. Lim, R., Ferrari, F., Zimmerling, M., Walser, C., Sommer, P., Beutel, J.: Flocklab: A testbed for distributed, synchronized tracing and profiling of wireless embedded systems. In: 2013 ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN). pp. 153–165 (April 2013). https://doi.org/10.1145/2461381.2461402
14. Sigrist, L., Gomez, A., Lim, R., Lippuner, S., Leubin, M., Thiele, L.: Measurement and validation of energy harvesting iot devices. In: Proceedings of the 2017 Design, Automation & Test in Europe Conference & Exhibition (DATE 2017). Lausanne, Switzerland (Mar 2017)
15. Hurni, P., Nyffenegger, B., Braun, T., Hergenroeder, A.: On the accuracy of software-based energy estimation techniques. In: Marrón, P.J., Whitehouse, K. (eds.) Wireless Sensor Networks. pp. 49–64. Springer Berlin Heidelberg (2011), https://link.springer.com/chapter/10.1007%2F978-3-642-19186-2_4
16. Sieber, A.: Energiemanagement fuer drahtlose tief eingebettete Systeme : Feingranulares, faires Management zum Erreichen von Lebenszeitzielen bei knappen Energieressourcen. doctoralthesis, BTU Cottbus - Senftenberg (2016), https://opus4.kobv.de/opus4-btu/frontdoor/index/index/docId/3928
17. Baek, D., Park, J.H., Lee, J.W.: An energy efficiency grading system for mobile applications based on usage patterns. The Journal of Supercomputing **74**(12), 6502–6515 (Dec 2018). https://doi.org/10.1007/s11227-018-2439-x
18. Banerjee, A., Chong, L.K., Chattopadhyay, S., Roychoudhury, A.: Detecting energy bugs and hotspots in mobile apps. In: Proceedings of the 22Nd ACM SIGSOFT International Symposium on Foundations of Software Engineering. pp. 588–598. FSE 2014, ACM, New York, NY, USA (2014). https://doi.org/10.1145/2635868.2635871
19. Gastel, P.v., Gastel, B.v., Eekelen, M.v.: Detecting energy bugs and hotspots in control software using model checking. In: Conference Companion of the 2Nd International Conference on Art, Science, and Engineering of Programming. pp. 93–98. Programming Companion, ACM, New York, NY, USA (2018). https://doi.org/10.1145/3191697.3213805
20. Roth, M., Luppold, A., Falk, H.: Measuring and modeling energy consumption of embedded systems for optimizing compilers. In: Proceedings of the 21st International Workshop on Software and Compilers for Embedded Systems. pp. 86–89. SCOPES '18, ACM, New York, NY, USA (2018). https://doi.org/10.1145/3207719.3207729
21. Rodrigues, L.M., Montez, C., Budke, G., Vasques, F., Portugal, P.: Estimating the lifetime of wireless sensor network nodes through the use of embedded analytical battery models. Journal of Sensor and Actuator Networks **6**(2) (2017). https://doi.org/10.3390/jsan6020008
22. Rodrigues, L.M., Montez, C., Moraes, R., Portugal, P., Vasques, F.: A temperature-dependent battery model for wireless sensor networks. Sensors **17**(2) (2017). https://doi.org/10.3390/s17020422
23. Buchli, B., Aschwanden, D., Beutel, J.: Battery state-of-charge approximation for energy harvesting embedded systems. In: Demeester, P., Moerman, I., Terzis, A. (eds.) Wireless Sensor Networks. pp. 179–196. Springer Berlin Heidelberg, Berlin, Heidelberg (2013)

24. Green, M.A., Dunlop, E.D., Levi, D.H., Hohl-Ebinger, J., Yoshita, M., Ho-Baillie, A.W.: Solar cell efficiency tables (version 54). Progress in Photovoltaics: Research and Applications **27**(7), 565–575 (2019). https://doi.org/10.1002/pip.3171
25. Rao, R.R., Mani, M., Ramamurthy, P.C.: An updated review on factors and their inter-linked influences on photovoltaic system performance. Heliyon **4**(9), e00815 (2018). https://doi.org/https://doi.org/10.1016/j.heliyon.2018.e00815
26. Kraemer, F.A., Ammar, D., Braten, A.E., Tamkittikhun, N., Palma, D.: Solar energy prediction for constrained iot nodes based on public weather forecasts. In: Proceedings of the Seventh International Conference on the Internet of Things. pp. 2:1–2:8. IoT '17, ACM, New York, NY, USA (2017). https://doi.org/10.1145/3131542.3131544
27. Buchli, B., Sutton, F., Beutel, J., Thiele, L.: Towards enabling uninterrupted long-term operation of solar energy harvesting embedded systems. In: Proceedings of the 11th European Conference on Wireless Sensor Networks - Volume 8354. pp. 66–83. EWSN 2014, Springer-Verlag New York, Inc., New York, NY, USA (2014). https://doi.org/10.1007/978-3-319-04651-8_5
28. Mardaljevic, J., Christoffersen, J.: climate connectivity in the daylight factor basis of building standards. Building and Environment **113**, 200 – 209 (2017). https://doi.org/https://doi.org/10.1016/j.buildenv.2016.08.009, advances in daylighting and visual comfort research
29. Wold, C.: Analysis of background illuminance levelsduring television viewing. Tech. rep., Collaborative Labeling and Appliance Standards Program (CLASP) (2011), https://www1.eere.energy.gov/buildings/appliance_standards/pdfs/analysis_of_background_illuminance_levels_during_television_viewing.pdf
30. Tamkittikhun, N., Hussain, A., Kraemer, F.A.: Energy consumption estimation for energy-aware, adaptive sensing applications. In: Bouzefrane, S., Banerjee, S., Sailhan, F., Boumerdassi, S., Renault, E. (eds.) Mobile, Secure, and Programmable Networking. pp. 222–235. Springer International Publishing, Cham (2017). https://doi.org/10.1007/978-3-319-67807-8_17
31. Lachenmann, A., Marrón, P.J., Minder, D., Rothermel, K.: Meeting lifetime goals with energy levels. In: Proceedings of the 5th International Conference on Embedded Networked Sensor Systems. pp. 131–144. SenSys '07, ACM, New York, NY, USA (2007). https://doi.org/10.1145/1322263.1322277
32. Geissdoerfer, K., Jurdak, R., Kusy, B., Zimmerling, M.: Getting more out of energy-harvesting systems: energy management under time-varying utility with preact. In: Proceedings of the 18th International Conference on Information Processing in Sensor Networks, IPSN 2019, Montreal, QC, Canada, April 16-18, 2019. pp. 109–120 (2019). https://doi.org/10.1145/3302506.3310393
33. Rao, S., Chendanda, D., Deshpande, C., Lakkundi, V.: Implementing lwm2m in constrained iot devices. In: 2015 IEEE Conference on Wireless Sensors (ICWiSe). pp. 52–57 (Aug 2015). https://doi.org/10.1109/ICWISE.2015.7380353
34. Schmelzer, P., Akelbein, J.: Evaluation of hardware requirements for device management of constrained nodes based on the LWM2M standard. In: Proceedings of the 5th Collaborative European Research Conference (CERC 2019), Darmstadt, Germany, March 29-30, 2019. pp. 103–110 (2019), http://ceur-ws.org/Vol-2348/short02.pdf
35. Moran, B., Meriac, M., Tschofenig, H., Brown, D.: A firmware update architecture for internet of things devices. Internet-Draft draft-ietf-suit-architecture-05, IETF Secretariat (April 2019), http://www.ietf.org/internet-drafts/draft-ietf-suit-architecture-05.txt