

An Open-Loop Energy Neutral Power Manager for Solar Harvesting WSN

Andrea Castagnetti¹, Alain Pegatoquet¹, Cécile Belleudy¹, Michel Auguin¹

¹*University Nice Sophia Antipolis, LEAT-CNRS, 250 rue Albert Einstein, 06560 Valbonne, France*
{castagnetti, pegatoquet, belleudy, auguin}@unice.fr

Keywords: Wireless Sensor Networks, Solar Energy Harvesting, Power Management.

Abstract: Energy harvesting Wireless Sensor Networks are receiving increasing interest due to their potential to extend system lifetime. Because environmental energy availability is highly variable, an efficient power management is required. An energy harvesting power manager must adapt to different situations that depends on the energy that can be harvested from the environment. In this paper we propose a generic model for solar energy harvesting wireless sensor node, that we validate on real hardware. A novel power management architecture is then proposed. Simulation results show that up to 30% performance improvement can be achieved compared to a state of the art power management algorithm.

1 INTRODUCTION

Wireless Sensor Networks (WSNs) consists of distributed autonomous sensors that can operate without a pre-established infrastructure. WSNs are used for monitoring environment (temperature, humidity, etc.), fire surveillance, high precision agriculture and many other applications. Due to their physically embedded nature, WSNs are well suited for energy-harvesting. This offers an alternative solution to the problem of the energy limitation that arise from the fact that sensor nodes are equipped with small batteries. More often, due to the fact that nodes are deployed in harsh environment, batteries cannot be replaced. Virtually, the lifetime of an energy-harvesting WSN node is only limited by the hardware reliability. However we argue that the potential of an energy-harvesting WSNs node cannot be exploited without an efficient power management system.

In battery-powered system, the goal of a power manager is to minimize the power consumption in order to maximize the lifetime of the battery. At the opposite energy-harvesting systems can experience peak of energy availability. During those periods the system can provide peak quality of service (QoS) only using the energy coming from the environment. Those periods are interleaved with intervals in which the environment provides only small amount of energy, or even nothing. When an energy-harvesting node operates on battery, performance must be adjusted to extend the lifetime of the node until suffi-

cient energy will again be provided by the environment.

From the situations sketched above it becomes clear that different power management techniques can be used depending on environmental conditions. As an example, consider the case in which a sufficient amount of energy can be harvested from the environment. The power manager can then decide to operate the system in energy neutrality, that is balancing the energy consumed (to perform the task) and the harvested energy.

In this paper we address the problem of power management for energy-harvesting WSNs starting from a modeling perspective. We propose a generic model that permits to describe a solar energy-harvesting node in a compact form. The derived model is validated on a recent solar energy harvesting sensor platform. Based on this framework, we propose a novel architecture, called Open-Loop energy neutral Power Manager (OL-PM). This power manager is validated in simulation. We have also adapted the power manager presented in (Kansal et al., 2007) to our model and we compared its performance against our power manager. Experimental results are very encouraging as up to 30% data throughput improvement can be achieved. Moreover we have observed that the battery is never completely discharged using our power manager, that leads to an improvement on system reliability. Finally the impact of the power management reactivity on the system's performance has been studied.

2 RELATED WORK

In recent years several energy harvesting technology have emerged. Besides solar panels (Nelson, 2003) and wind generators, vibration scavengers, thermal-to-electricity (Toriyama et al., 2001) and electromagnetic converters (Yi et al., 2007) have also been proposed. Many energy harvesting systems also need energy storage because they have to continue operation even when no energy is available from the environment (i.e. during night for a solar harvesting system). Nowadays rechargeable batteries are used for long-term energy storage. However new technologies are emerging like supercapacitors (Shin et al., 2011).

Besides the problem of designing an energy harvesting system that effectively extract and store energy the issue of an efficient power management policy must be faced. Research on energy harvesting power management has gained lot of attention in recent years mainly in the field of wireless sensor networks. One of the earliest work on energy harvesting power management has been presented in (Kansal et al., 2004). Duty-cycling is used to adjust the power consumption of a WSN node. During an initialization phase the characteristics of the energy source are learned. Then a fixed duty-cycle is applied. The effectiveness of this solution depends on the variability of the energy source. A more principled algorithm for dynamically adapting the the duty-cycle of a sensor node is discussed in (Kansal et al., 2007). The authors assume that the energy source is periodic. A period is then divided into intervals of equal duration. An estimate of the energy input, that is assumed to be constant over the course of a single block, is computed from historical data. Duty-cycle is then set for each block based on an initial estimate. Online changes are applied if there is a mismatch between the actual energy received and the expected energy computed by the model. The reactivity of this power management solutions, that is the ability to adapt to environmental variations, mainly depends on the initial choice of intervals duration. In the work of (Kansal et al., 2007) this duration was fixed to 30 minutes. The authors also present the notion of energy neutral operation, that is the ability to operate such that the energy used is always less than the energy harvested. They use this concept to develop the energy neutral power manager discussed above.

The work of (Kansal et al., 2007) is extended in (Vigorito et al., 2007). Here the authors use techniques from adaptive control theory to increase the adaptivity of an energy harvesting power manager. This power manager only uses the current battery level of the node to make duty-cycling decisions. This

approach is then model-free with respect to the energy source. The power manager is validated in simulation. Experimental results show that an average 16% performance improvement can be achieved compared to the power manager proposed by (Kansal et al., 2007).

3 SYSTEM MODEL

In this Section we present the system model that is composed of three main components: the platform load, the energy harvesting and the battery.

3.1 A Task Level Platform Load Model

The Figure 1 represents a formalization of the platform load. As shown, the platform load is first char-

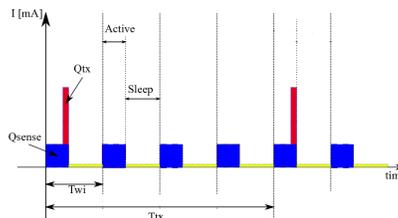


Figure 1: Formalization of the platform load

acterized by a wake-up period T_{wi} . During this period, the platform is first active, then inactive. An active period may be composed by several activities that correspond to different current consumption. A typical active period is composed of Sensor, CPU and RF activities. However, in order to get a platform model as much general as possible, different activities having their own period that can be different from the wake-up period can be defined. For example, in Fig. 1 two different kinds of activities have been defined: sensing and transmission. The sensing activity is executed every wake-up period (T_{wi}), while the transmission occurs with a period T_{tx} . During the inactive period the system enters sleep state for power saving. Q_{sense} is the charge delivered by the battery when the sensing task is executed, which accounts for the current consumed by the microcontroller and the sensor. When the transmission task is also executed, there is an increase in the current consumption and the battery delivers supplementary Q_{TX} Ampere-hour of charge compared to previous case. The rate at which the battery is discharged depends on the average discharge current. This parameter can be computed using our model and it is represented by the α factor that is defined as follows:

$$\alpha = \frac{Q}{T_{wi}} \quad (1)$$

α denotes the charge consumed by the platform over a wake-up period. In other words α is the average current delivered by the battery during a wake-up period. This parameter can be computed independently for each activity (e.g. sensing, RF transmission) enabling a modular and accurate characterization of the application.

3.2 A Battery and Energy Harvesting Integrated Model

We model the solar panel and the recharge circuit using a parameter, called β , which is a function of the light intensity. The β parameter models both the efficiency of the solar panel and the efficiency of the voltage regulator and the charge circuit that are used to recharge the battery. This parameter is expressed in Ampere and indicates the rate at which the solar panel can recharge the battery under a fixed light intensity. Figure 2 shows the β parameter as a function of the

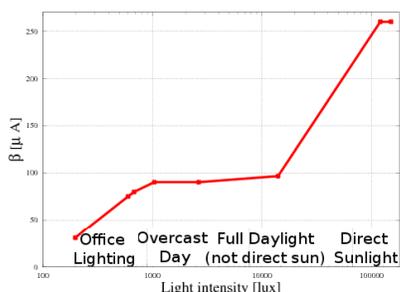


Figure 2: β for different light conditions

light intensity expressed in Lux using a logarithmic scale. The curve was extracted through experimentations for a 2.25in x 2.25in solar panel and a 100 μ Ah Li-Ion battery. More parameters can be taken into account in order to refine the model, e.g. the type of light (fluorescent, incandescent, sun light) and its spectral characteristics, the quantum efficiency of the solar cell and more. Although simple, our model provides a good approximation of the energy that can be harvested by the system and help defining the application load requirements and the power management policies.

3.3 A Battery State of Charge Model for Periodic Workload

Using the α and β parameters introduced in Sections 3.1 and 3.2, the state of charge (SoC) of a battery for the next n wake-up periods (T_{wi}) can be computed

using equation 2:

$$SoC(t + nT_{wi}) = SoC(t) + [\beta - (\alpha_s + \alpha_{Tx} \frac{T_{wi}}{T_x})] nT_{wi} - K_{leak} nT_{wi} \quad (2)$$

Where α_s and α_{Tx} are the current consumption for sensing and transmission respectively. The Eq. 2 is true if the conditions 3, 4 and 5 are respected.

$$SoC(t = 0) = SoC_{max} \quad (3)$$

$$SoC_{min} \leq SoC(t + nT_{wi}) \leq SoC_{max} \quad (4)$$

$$\beta, \alpha_s \text{ and } \alpha_{Tx} \text{ are constants on } [t, t + nT_{wi}] \quad (5)$$

SoC_{max} represents the capacity of the battery. We consider that the battery is fully charged at $t = 0$ (3). As the battery can neither store more charge than its capacity nor be discharged too deeply, its SoC is bounded between SoC_{max} and SoC_{min} (4). Finally and for sake of simplicity, Eq. 2 is true if the luminosity (β) and the average discharge current over a wake-up period (α) are constants (5). If these parameters are changing during this period, Eq. 2 can be easily divided into smaller period intervals. The leakage current is modeled with the K_{leak} parameter. This parameter models two effects: the battery self discharge and the current drawn by the platform during inactive periods (e.g. the low-power mode current consumption).

3.4 Model Validation Using the EZ430 Platform

We conducted a set of tests to validate our SoC estimation approach. The Texas Instruments EZ430 solar energy harvesting platform (Instrument, 2011) was used as experimental platform. This platform is equipped with a 2.25in x 2.25in solar panel optimized for operating indoor under low-intensity fluorescent light. The energy harvested by the solar panel is stored in a 100 μ Ah Li-Ion rechargeable battery. The application board is equipped with an MSP430 microcontroller and a CC2500 RF transceiver. The board is programmed with an End Device (ED) application that periodically sends a packet over the air to a base station. During the idle period both the microcontroller and the radio chip are put into a sleep state.

We used Equation 2 to predict the SoC, then the lifetime (LT) of the system under different light conditions. The estimated and measured life-times (LT) are compared for different wake-up periods. The results are shown in Tab. 1. As it can be observed, the model is accurate as the error always remains below

Table 1: Validation of the SoC estimation technique (LT in minutes)

T_{wi} [sec]	α [μ A]	0 lux		≈ 200 lux	
		LT	LT	LT	LT
		(exp)	(model)	(exp)	(model)
1	128.52	28	29	50	38
2	64.26	54	59	140	112
3	42.84	83	88	312	332
4	32.13	107	117	/	$+\infty$
6	21.42	162	176	/	$+\infty$
10	12.85	290	294	/	$+\infty$
20	6.43	585	587	/	$+\infty$

10% for 0 lux. It is interesting to note that with a T_{wi} of 20 seconds the tests lasts for 10 hours while the lifetime prediction error is only 2 minutes. The prediction error slightly increases for 200 lux, but the model still gives accurate results. Actually, during experimentations the luminosity fluctuates around 200 lux. This is why we observe a higher deviation of the model than without energy harvesting ($\beta = 0$). For a T_{wi} higher or equal to 4 seconds the energy harvested (β) is higher than the energy consumed (α), so the system is working in energy neutral conditions. Therefore the LT is theoretically infinite (∞). Despite our model has been validated on the TI EZ430 platform, it is applicable to a wide range of platforms.

3.5 Range of Applicability of the Model

In this section we discuss how our model can adapt to different architectures. We first analyze the platform load model proposed in Section 3.1. The power consumption usually changes according to the voltage supply. Nowadays electronics can be powered with a wide voltage range. If a switching regulator (e.g. a buck regulator) is used to provide regulated power to the load then the current that it draws from the battery will change according to its input and output voltage. In this case the α parameter must be rescaled by a factor that takes into account the efficiency of the switching regulator. At the opposite, if a linear regulator is used, as in our case, it is reasonable to consider α to be independent from the battery voltage.

In Section 3.2 we proposed to model the recharge process with a β parameter, which is a function of the light intensity. The solar cell and recharge circuit are then modeled as a simple current generator. Some assumptions must be made in order for this model to be valid. First of all the operating point of the solar panel must be forced by a proper circuit (i.e. MPPT control circuit). Moreover the efficiency of the recharge circuit will vary a lot depending on the output voltage of

the solar cell. Finally, the capacity of the battery will have a huge impact on the recharge process. In our model all the variables which have an impact on the efficiency of the recharge process are implicitly included in the β parameter. Of course β will change for a different solar panel and so the model would have to be re-calibrated.

4 OPEN-LOOP ENERGY NEUTRAL POWER MANAGER

In this Section we consider the problem of adapting the performance and the power consumption of the system to the available energy. Performance scaling is achieved by varying the wake-up period (T_{wi}) of the sensor node. Duty-cycling between active and low power states is very effective in reducing the power consumption. Moreover, almost all the WSNs nodes features one or more low-power states. Figure 3 shows the architecture of an Open-loop power manager (OL-PM). We use the name open-loop to under-

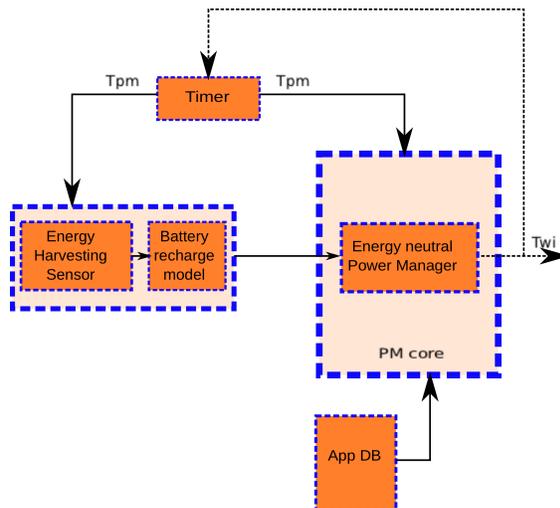


Figure 3: Open-Loop Power Manager Architecture

line the fact that the decision of the power manager is based only on the current value of β . The power manager does not use the SoC prediction to compute the current value of T_{wi} . The system is composed of four main blocks:

Power manager Timer (T_{pm}): The power manager is executed periodically, with a period that is a multiple of the wake-up period T_{wi} . Every time T_{wi} changes, this timer is set to n times T_{wi} . Limiting the execution of the power manager reduces the overhead but it also limits the reactivity. The trade-off between overhead and reactivity can be tuned by adjusting the n parameter.

Energy harvesting sensor and battery recharge model: An energy harvesting sensor is used to estimate the recharging rate of the battery. This value will then be used to compute the recharge rate using the battery recharge model, namely β . For a solar panel the harvesting sensor could be a current meter that measures the output current of the solar panel. If it is not possible to measure the output current of the solar cell, the energy harvesting sensor can be replaced by a light meter.

System parameters database (AppDB): The T_{wi} settings as well as the task level platform loads (α_i) are stored in a database.

Energy neutral power manager: This block implements the power management strategy. By balancing the α and β parameters, the system works with only the energy coming from the environment. In this case the SoC of the battery is constant over the time. Its value depends on the past history, i.e. how much time the system has worked in both negative and positive energy conditions. Using Equation 2 the condition for energy neutrality can be expressed as follows:

$$SoC(t) = SoC(t + nT_{wi}) \quad (6)$$

This condition leads to the following energy neutral wake-up period:

$$T_{wi} = \left\lceil \frac{Q}{\beta - K_{leak}} \right\rceil \quad (7)$$

As T_{wi} is an integer the result of Eq. 7 must be approximated (the decimal part must be removed). The ceil function is used to round toward the nearest energy neutral.

5 SIMULATION RESULTS

We compared the OL-PM with the power manager presented in (Kansal et al., 2007) that we have adapted to the model discussed in Section 3. We simulated the two power managers on a five days solar energy data set. This Section presents the simulation setup as well as the simulation results.

5.1 Simulation Setup and Energy Data Set Profile

The values for the α parameter were measured for the TI EZ430 platform (some are reported in Tab. 1). In order to get realistic values for the β parameter we measured the light intensity in an office. We connected a luxmeter to a PC and repeatedly measured the light intensity at an interval of five seconds during five days. The values for the β parameter were then

computed using the function of Fig. 2. The results are plotted in Fig. 4. The system parameters are shown in

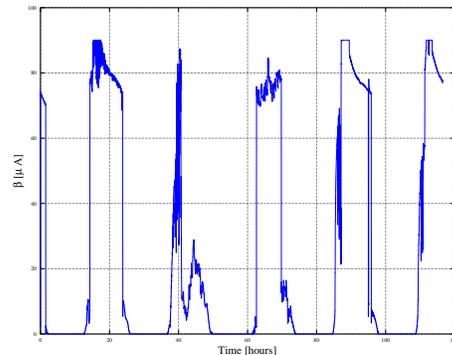


Figure 4: The values of β during five days in an office

Table 2. The overhead of a power management execution (Q_{pm}) has been considered equal to the charge consumed for a RF transmission. The battery param-

Table 2: System parameters

	T_{wi} [min, max]	Q_{tx}	Q_{pm}	n
	[sec]	[μ As]	[μ As]	
OL-PM	[1, 120]	128.52	128.52	10

eters are reported in Tab. 3.

Table 3: Battery parameters

SoC_{min}	SoC_{max}	C_d
[μ Ah]	[μ Ah]	[μ Ah]
37	100	63

5.2 Evaluation Metrics

In order to evaluate both power managers we have defined different metrics:

- Average data-Rate ($\langle Rd \rangle$): It is a metric of the performance of the power manager. It is defined as the ratio between the size of the packet that is sent (we use a payload of 33 bytes) and the wake-up period, T_{wi} . The average is computed over the five days, including the periods of time where the battery is fully discharged (in this case the data-rate is zero).
- Maximum and minimum data-rate (Rd_{max} , Rd_{min}): Those metrics give an indication of the peak and minimal performance achievable by a node using a given power manager.
- Average SoC ($\langle SoC \rangle$): To assess if the power management algorithm drift. Actually, the average

SoC can be used to determine how far the decisions of the power manager are from the energy neutral condition.

- Battery failures (B_f) : This metric gives a measure of the correctness and reliability of the power manager’s choices. A value of 0 means that the battery is never fully discharged and the node is always operational.

5.3 Experimental Results

Simulation results for the power manager are reported in Tab. 4.

Table 4: Performance comparison of Kansal algorithms (Kansal et al., 2007) and the OL-PM

	$\langle Rd \rangle$	Rd_{max}	Rd_{min}	$\langle SoC \rangle$	B_f
	[bits/s]	[bits/s]	[bits/s]	[μ Ah]	
Kansal	29.55	132	0	65.81	9
OL-PM	44.61	132	2.2	88.8	0

From Tab. 4 it can be observed that for the data set used the OL-PM improves the data-rate by more than 30% compared to the algorithm proposed by (Kansal et al., 2007). This can be mainly explained by the fact that the data set has a high energy variability as it can be observed from Fig. 4. The OL-PM is able to exploit this variability thanks to of its reactivity (that can be tuned through the n parameter). The activation period of this power manager is in the order of tens of seconds. The algorithm proposed by (Kansal et al., 2007) takes a different approach. It samples the energy source each 30 minutes and it also uses an EWMA filter to predict the energy availability. The combination of those two effects excessively smoothes the data from the energy harvester (β), and the power manager cannot exploit the peak of energy availability to increase the data-rate (Rd). The differences in reactivity also explains that the battery is never fully discharged for the OL-PM, while it happens nine times in five days for Kansal. As a consequence the node stops sending data, for a variable amount of time, on average two times a day.

Rd_{max} is the same for the two algorithms, as it depends on the peak energy harvested by the solar panel, which is the same in both cases. It is interesting to note that both power managers’ goal is to operate the system in energy neutrality. In perfect conditions the average SoC of the battery should be equal to the initial SoC (in our case SoC_{max}). Inaccuracies in β estimation and lack of reactivity cause a SoC drift. The drift is only 10μ Ah, that is 16% of the discharge capacity (C_d) for the OL-PM. It can accurately track the variation of β and it can operate the system near the

energy neutrality. At the opposite the Kansal algorithm produces a drift of 34μ Ah, that is 54% of the discharge capacity C_d . This is due to the assumption made by the authors that the energy harvested is constant over a period of 30 minutes. Actually this assumption is violated many times in our data set. We therefore conclude that considering the harvested energy constant over a long period (i.e. 30 minutes) could lead to performance degradation.

The rest of our experiments investigated the impact of varying the reactivity of the power manager, on the overall performance of the OL-PM. Figure 5 shows the average SoC, the average data rate and the SoC_{min} obtained for different values of n . As previ-

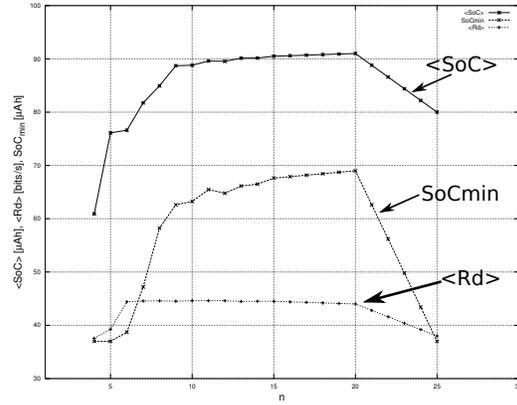


Figure 5: The effects of varying the power manager reactivity (n) on the $\langle SoC \rangle$, SoC_{min} and $\langle Rd \rangle$

ously discussed the more the SoC diverges from the SoC_{max} the more the energy neutrality condition is violated. Experimental results show that for a n smaller than 10 the system tends to perform poorly. Actually the gain induced by the increased reactivity is counterbalanced by the overhead. The impact of the power management overhead is higher when β is lower (i.e. during night). During those periods the system cannot actually work in energy neutrality even with the smallest value of T_{wi} . As it can be seen, the average SoC as well as the average data-rate tends to become stable for a n comprised between 10 and 20. For values greater than 20 the lack of reactivity of the power manager deteriorates the global performance.

Finally, the Fig. 6 depicts the instantaneous values of the SoC, the data-rate (R_d) and the energy harvested (β) during the second and third day for both approaches. As shown, the OL-PM reactivity allows the data-rate to accurately follow the energy harvested profile. In the meantime, for Kansal the data-rate is less correlated to the β and leads to a larger SoC drift than with our OL-PM.

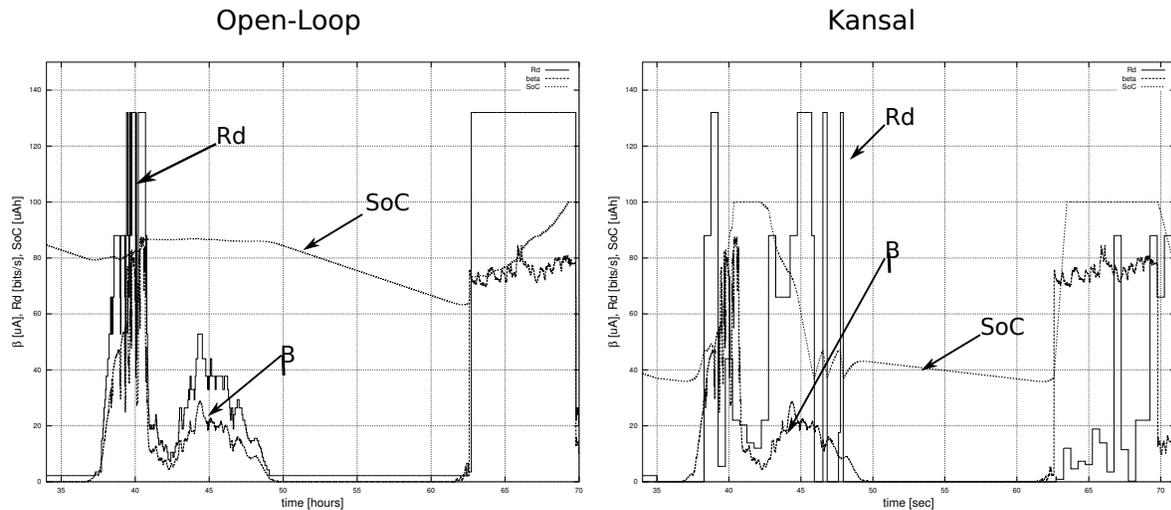


Figure 6: Execution traces of the OL-PM ($n = 10$) and Kansal during the second and third day

6 CONCLUSIONS

In this paper, we have proposed a modeling approach for an energy harvesting wireless sensor node. Based on the derived model we have developed a power manager that can dynamically adapt the duty-cycle of the node to operate in energy neutral condition. Experimental results have shown that up to 30% performance improvement can be achieved compared to a state of the art power management algorithm. We have also studied the impact of the power management reactivity on the global system performance. Our analysis has shown that the deviation from the energy neutrality condition is related to the reactivity of the power manager. Therefore, this outcome can be used to increase the power management efficiency by adding an adaptive control of the reactivity. Future works will also focus on using SoC estimation to adapt the power manager policies. Finally, we aim to extend our approach for different energy harvesting system (piezoelectric, thermoelectric) and batteries.

ACKNOWLEDGEMENTS

This work is supported by the French National Research Agency (ANR) project GRECO bearing reference ANR-2010-SEGI-004-04.

REFERENCES

- Instrument, T. (2011). ez430-rf2500 solar energy harvesting kit.
- Kansal, A., Hsu, J., Zahedi, S., and Srivastava, M. (2007). Power management in energy harvesting sensor networks. *ACM Transactions on Embedded Computing Systems (TECS)*, 6(4):32–es.
- Kansal, A., Potter, D., and Srivastava, M. B. (2004). Performance aware tasking for environmentally powered sensor networks. *SIGMETRICS Perform. Eval. Rev.*, 32:223–234.
- Nelson, J. (2003). *The physics of solar cells*. Imperial College Press London.
- Shin, D., Kim, Y., Seo, J., Chang, N., Wang, Y., and Pedram, M. (2011). Battery-supercapacitor hybrid system for high-rate pulsed load applications. In *DATE*, pages 875–878. IEEE.
- Toriyama, T., Yajima, M., and Sugiyama, S. (2001). Thermoelectric micro power generator utilizing self-standing polysilicon-metal thermopile. In *Micro Electro Mechanical Systems, 2001. MEMS 2001. The 14th IEEE International Conference on*, pages 562–565. IEEE.
- Vigorito, C., Ganesan, D., and Barto, A. (2007). Adaptive control of duty cycling in energy-harvesting wireless sensor networks. In *Sensor, Mesh and Ad Hoc Communications and Networks, 2007. SECON'07. 4th Annual IEEE Communications Society Conference on*, pages 21–30. IEEE.
- Yi, J., Ki, W., and Tsui, C. (2007). Analysis and design strategy of uhf micro-power cmos rectifiers for micro-sensor and rfid applications. *Circuits and Systems I: Regular Papers, IEEE Transactions on*, 54(1):153–166.