

A Holistic Approach to Energy Harvesting for Indoor Robots: Theoretical Framework and Experimental Validations

THÈSE N° 6522 (2015)

PRÉSENTÉE LE 20 MARS 2015

À LA FACULTÉ DES SCIENCES ET TECHNIQUES DE L'INGÉNIEUR
LABORATOIRE DE SYSTÈMES ROBOTIQUES
PROGRAMME DOCTORAL EN SYSTÈMES DE PRODUCTION ET ROBOTIQUE

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Florian Christopher VAUSSARD

acceptée sur proposition du jury:

Dr A. Karimi, président du jury
Prof. F. Mondada, directeur de thèse
Dr M. Bonani, rapporteur
Prof. L. M. Gambardella, rapporteur
Prof. A. Ijspeert, rapporteur



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2015

古池や

蛙飛びこむ

水の音

芭蕉

To Juliette, who gave me her Love.

Acknowledgements

This work was made with the help of numerous people. First and foremost, I wish to thank my thesis committee members: Prof. **Auke Ijspeert**, Prof. **Luca Gambardella**, Dr. **Michael Bonani** and Dr. **Karimi Alireza**. They carefully read my manuscript and gave me several insightful remarks to enhance the quality. I also thank my thesis advisor Prof. **Francesco Mondada** for funding my work and offering me a pleasant place where I worked during the last four years. This work was funded by the **Swiss National Science Foundation (SNFS)** through the **National Center of Competence in Research Robotics (NCCR Robotics)**. I am also grateful to all the members of the NCCR Robotics.

David Hamel and **Alain Vuille** greatly contributed to this work. They deserve my warmest thanks. David Hamel performed most of the measures used in Chap. 2. This was a lengthy and tremendous process. Alain Vuille patiently engineered and assembled the mechanical simulator used in Chap. 6. These contributions were key pieces towards the achievements described in this thesis. **Norbert Crot** machined most of the mechanical parts of this project. He was essential to the success of the final experiment. **Daniel Burnier** was always helpful when dealing with electronics and assisted me in making several boards.

I address a big “thank you” to my fellow labmates. They created a friendly working atmosphere and challenged my mind many times. **Philippe Rétornaz** is an amazing engineer. I learned a lot by working with you, and our numerous discussions at Sat helped me to overcome the difficult moments. This was an honour to finish my PhD thesis with you! I also survived thanks to the cakes and various pastries of **Fanny Riedo**. You always brought some kindness inside this harsh world. **Alain Berthoud** initiated me to the art of wild jacuzzi on frozen lakes. You showed me how a PhD thesis is easy compared to pulling sledges loaded with tons of exotic material. I have a great debt towards **Frédéric Rochat**. You took on your free time to carefully read my chapters as I was writing them. Your tireless encouragements helped me to not feel alone during the very last months. You are all great friends.

I also met a lot of nice people, both inside and outside the laboratory. I want to thank all my past and present colleagues: **Mariza Freire**, **Anne Remillet**, **Julia Fink**, **Josep Soldevila**, **Patrick Schoeneich**, **Stéphane Magnenat**, **Alexey Gribovskiy**, **Valentin Longchamp**, **François Rey**, **Hadi Ardiny**, **Christophe Barraud**, and **Frank Bonnet**. You all contributed to my quality of life during these four years. I had great times with you and wish that we can keep in touch

Acknowledgements

in the future. The coffee breaks with **Ayumi Koishi** were always good moments to forget the problems. I wish you all the best with your own thesis, regardless of my warnings!

Finally and most importantly, all of this was possible thanks to the support of my family and friends. These four years were an exhausting adventure, so the key to survival is to enjoy the good moments when they happen. **Juliette** is now my wife, but first of all she is my soulmate. You supported me each and every day, you smiled at me when I was sad and you cheered me up when life was harder than usual. This work would not have been possible without you. I also thank my parents **Guy** and **Christine** for supporting me all along my studies. The moments spent with my sister **Véréna** were likewise important for the success of my work. Your gift for a wellness weekend was an awesome idea!

Many thanks to all of you!

Abstract

Service robotics is a fast expanding market. Inside households, domestic robots can now accomplish numerous tasks, such as floor cleaning, surveillance, or remote presence. Their sales have considerably increased over the past years. Whereas 1.05 million domestic service robots were reportedly sold in 2009, at least 2.7 million units were sold in 2013. Consequently, this growth gives rise to an increase of the energy needs to power such a large and growing fleet of robots. However, the unique properties of mobile robots open some new fields of research. We must find technologies that are suitable for decreasing the energy requirements and thus further advance towards a sustainable development.

This thesis tackles two fundamental goals based on a holistic approach of the global problem. The first goal is to reduce the energy needs by identifying key technologies in making energy-efficient robots. The second goal is to leverage innovative indoor energy sources to increase the ratio of renewable energies scavenged from the environment.

To achieve our first goal, new energy-wise metrics are applied to real robotic hardware. This gives us the means to assess the impact of some technologies on the overall energy balance. First, we analysed seven robotic vacuum cleaners from a representative sample of the market that encompasses a wide variety of technologies. Simultaneous Localisation and Mapping (SLAM) was identified as a key technology to reduce energy needs when carrying out such tasks. Even if the instantaneous power is slightly increased, the completion time of the task is greatly reduced. We also analysed the needed sensors to achieve SLAM, as they are largely diversified. This work tested three sensors using three different technologies. We identified several important metrics.

As of our second goal, potential energy sources are compared to the needs of an indoor robot. The sunshine coming through a building's apertures is identified as a promising source of renewable power. Numerical simulations showed how a mobile robot is mandatory to take full advantage of this previously unseen situation, as well as the influence of the geometric parameters on the yearly energy income under ideal sunny conditions. When considering a real system, the major difficulty to overcome is the tracking of the sunbeam along the day. The proposed algorithm uses a hybrid method. A high-level cognitive approach is responsible for the initial placement. Following realignments during the day are performed by a low-level reactive behaviour. A solar harvesting module was developed for our research robot. The tests

Abstract

conducted inside a controlled environment demonstrate the feasibility of this concept and the good performances of the aforementioned algorithm. Based on a realistic scenario and weather conditions, we computed that between 1 and 14 days of recharge could be necessary for a single cleaning task.

In the future, our innovative technology could greatly lower the energy needs of service robots. However, it is not completely possible to abandon the recharge station due to occasional bad weather. The acceptance of this technology inside the user's home ecosystem remains to be studied.

Keywords: Domestic service robots, Energy-efficient robotics, Indoor energy scavenging, Solar energy harvesting

Résumé

Le marché de la robotique de service est en pleine expansion. Les robots domestiques sont désormais capables d'accomplir un grand nombre de tâches, telles que le nettoyage des sols, la surveillance ou la téléprésence. Leur vente s'est considérablement accrue au fil des ans. Passant d'un volume de 1.05 millions d'exemplaires vendus en 2009, les robots domestiques ont atteints 2.7 millions d'unités vendues en 2013. Cette croissance a engendré un surcroît d'énergie nécessaire à leur fonctionnement. Néanmoins, les robots mobiles ont des propriétés uniques qui permettent d'ouvrir de nouveaux champs de recherche. Il nous incombe de trouver les technologies permettant de diminuer les besoins énergétiques, et ainsi de progresser en direction d'un développement plus durable.

En nous basant sur une approche holistique du problème, cette thèse a deux principaux objectifs. Il s'agit en premier lieu de réduire les besoins énergétiques, et ce grâce à l'identification de technologies clés permettant une robotique plus efficiente. En deuxième lieu, il nous faut exploiter des sources d'énergies innovantes, afin d'accroître le ratio d'utilisation des énergies renouvelables récupérées dans un environnement intérieur.

Concernant le premier objectif, nous avons défini de nouvelles métriques tenant compte de l'énergie. L'application de ces métriques aux composants robotiques actuels nous ont permis d'évaluer l'impact de certaines technologies dans la balance énergétique totale. Nous avons tout d'abord analysé sept robots aspirateur. Ceux-ci sont issus d'un panel représentatif du marché, et ils englobent un vaste ensemble de technologies. Ainsi, la Cartographie et Localisation Simultanées (CLS) a été identifiée comme une technologie clé dans la réduction des besoins énergétiques durant de telles tâches. En effet, le temps nécessaire à l'accomplissement du travail est grandement réduit, malgré une légère augmentation de la consommation instantanée. Nous avons également analysé les capteurs nécessaires à la CLS, une grande diversité de technologies existant dans ce domaine. Trois capteurs, issus de trois technologies différentes, ont été examinés au cours de ce travail. Nous avons ainsi pu identifier plusieurs métriques importantes.

Par la suite, nous avons comparé de potentielles sources d'énergie avec les besoins d'un robot d'intérieur, et ce afin de tendre vers notre second objectif. La lumière du soleil passant à travers les ouvertures d'un immeuble a été identifiée comme une source d'énergie renouvelable prometteuse. Les simulations numériques ont montré l'importance d'une base mobile

Résumé

afin de tirer pleinement parti de cette disposition inédite, ainsi que l'influence des paramètres géométriques sur les revenus énergétiques annuels, et ce dans des conditions idéalement ensoleillées. La difficulté majeure de mise en pratique vient de la poursuite du faisceau lumineux au cours de la journée. Pour ce faire, l'algorithme proposé utilise une méthode hybride. L'approche cognitive haut-niveau effectue le placement initial. Les réalignements suivants au cours de la journée sont réalisés par un comportement réactif bas-niveau. Un module de collecte d'énergie solaire a été développé pour notre robot de recherche. Les tests conduits à l'intérieur d'un environnement contrôlé démontrent la faisabilité de ce concept et les bonnes performances de l'algorithme susmentionné. Toutefois, nous avons calculé qu'au sein d'un scénario utilisant des conditions météorologiques réalistes, il serait nécessaire d'attendre entre 1 et 14 jours afin d'obtenir la recharge nécessaire à l'accomplissement d'une tâche de nettoyage.

Dans le futur, cette innovation technologique pourrait permettre de grandement diminuer les besoins énergétiques des robots de service. Cependant, la suppression complète de la station de recharge n'est pas envisageable en raison des conditions météorologiques ponctuellement défavorables. L'acceptation de cette technologie dans l'écosystème de l'utilisateur reste encore à étudier.

Mots clés : Robot domestique, Robotique efficiente en énergie, Récupération d'énergie en intérieur, Collecte d'énergie solaire

Contents

Acknowledgements	i
Abstract (English/Français)	iii
List of Figures	xi
List of Tables	xv
List of Acronyms	xvii
List of Symbols	xix
1 Introduction	1
1.1 Main Motivations	2
1.2 Goals of Our Study	3
1.3 Outline of Our Work	4
2 Energy Efficiency of Domestic Robots	5
2.1 Related Work	6
2.2 Methodology	7
2.3 Experimental Setup	8
2.4 Results	12
2.4.1 Power Analysis	12
2.4.2 Navigation Analysis	16
2.4.3 Specific Energy	19
2.4.4 Cleaning Efficiency	20
2.5 Discussion	22
2.6 People Who Contributed to This Work	22
3 Robotic Platform	23
3.1 The Standard MarXbot Robot	23
3.1.1 Mobility	24
3.1.2 Short-Range Sensing	25
3.1.3 Low-Level Design of the Modules	25
3.1.4 High-Level Control	26

Contents

3.1.5	Energy Management	28
3.2	Extension Modules for Mapping	31
3.2.1	Infrared Rotating Scanner	32
3.2.2	Neato Laser Range Scanner	33
3.2.3	Hokuyo Laser Range Scanner	34
3.3	People Who Contributed to This Work	35
4	Characterisation of Sensors for Energy-Efficient Mapping	37
4.1	Related Work	38
4.2	Methodology	39
4.3	Characterisation	40
4.3.1	Thermal Drift	40
4.3.2	Gaussian Noise Hypothesis	42
4.3.3	Impact of the Target's Distance	44
4.3.4	Calibration Model	46
4.3.5	Sensitivity to the Target's Angle	52
4.4	Power Analysis	54
4.4.1	Infrared Rotating Scanner	55
4.4.2	Neato Scanner	55
4.4.3	Hokuyo Scanner	58
4.5	Concluding Analysis	58
4.6	Discussion	61
5	Theoretical Framework for an Indoor Harvesting Robot	63
5.1	Indoor Energy Harvesting for Service Robots	64
5.1.1	Heat	64
5.1.2	Light	65
5.1.3	Ambient Radio Frequencies	66
5.1.4	Pressure and Temperature Variations	67
5.1.5	Gas and Liquid Flows	67
5.1.6	Mechanical	67
5.2	Scenario	68
5.2.1	Related Work	71
5.2.2	Research Hypotheses	74
5.3	Solar Mechanics	75
5.4	Theoretical Framework	77
5.4.1	Case Studies	77
5.4.2	Results	81
5.5	Concluding Analysis	89

6	Experimental Validation of an Indoor Harvesting Robot	93
6.1	The Solar Arena	94
6.1.1	Design	94
6.1.2	Experimental Validation	99
6.2	The Solar Harvesting Module	101
6.2.1	Design	101
6.2.2	Experimental Validation	105
6.3	Methodology and Algorithms	106
6.4	Experimental Results	110
6.4.1	Phase 1: Map Building	111
6.4.2	Phase 2: Systematic Coverage	113
6.4.3	Phase 3: Sun Tracking and Recharge	116
6.5	Concluding Analysis	123
6.5.1	Summary of the Results	123
6.5.2	From Simulation to Experiment	125
6.5.3	From Experiment to Reality	125
6.5.4	A Concrete Case	127
6.6	Discussion	128
6.7	People Who Contributed to This Work	129
7	Conclusion	131
7.1	Contributions	133
7.2	Outlook	133
A	Four-Channel Power Datalogger	135
A.1	Hardware	135
A.2	Calibration	137
	Bibliography	139
	Curriculum Vitae	151

List of Figures

1.1	Estimated number of service and entertainment robots for personal use sold between 2009 and 2013.	2
2.1	Robots used during our experiments, sorted by their localisation technology.	9
2.2	The experimental setups.	11
2.3	Two plots from the <i>in situ</i> power measures.	14
2.4	Box plots of the <i>in situ</i> power for each robot under several workloads.	15
2.5	Sample of the trajectories for each robot, grouped by localisation strategy.	17
2.6	Analysis of navigation performance for the seven robots.	19
2.7	The specific energy, in J/m^2 , for each robot.	20
2.8	Analysis of the cleaning task.	21
3.1	The marXbot robot.	24
3.2	Overview of the processing units.	26
3.3	Overview of the battery pack.	28
3.4	Simplified diagram of the power distribution network for a marXbot robot.	29
3.5	The infrared rotating scanner module.	32
3.6	The Neato scanner module.	34
3.7	The Hokuyo scanner module.	35
4.1	Picture of the simple setup used to characterise the mapping sensors.	39
4.2	Results for the thermal drift test.	41
4.3	Thermal images showing the self-heating process of the Hokuyo scanner.	42
4.4	Normality test for the three sensors.	43
4.5	Absolute error Δd as a function of the nominal distance d_n	45
4.6	Calibration steps for the infrared rotating scanner.	47
4.7	Calibration steps for the Neato scanner.	49
4.8	Calibration steps for the Hokuyo scanner.	50
4.9	Polar plot of the data accumulated during 50 scans by the Hokuyo scanner.	51
4.10	Sensitivity to the target's angle for each sensor.	53
4.11	Instantaneous power consumed by the processor and the infrared rotating scanner during the boot sequence.	54
4.12	Power analysis of the infrared rotating scanner.	56
4.13	Power analysis of the Neato scanner.	57

List of Figures

4.14	Power analysis of the Hokuyo scanner.	59
5.1	Schematic overview of the proposed scenario.	70
5.2	Pictures of some state-of-the-art photovoltaic robots.	72
5.3	Schematic views of the apparent solar motion.	75
5.4	Reference frame of the solar panel inside the global reference frame (x, y, z) of the room.	78
5.5	Results of the simulated sun trajectories: winter solstice, spring/autumn equinox, and summer solstice.	81
5.6	Simulated illumination for three select days during the year.	82
5.7	Results for case study 1.	84
5.8	Results for case studies 2, 3, 4 and 5.	86
5.9	Results for the case studies 6 and 7.	88
5.10	Amplitude between the two most distant positions reached by the robot, along the x axis and for a whole year.	91
6.1	Comparison of the trajectory achieved by the proposed solution and the actual solar trajectory.	95
6.2	Mechanical overview of the solar arena.	96
6.3	Comparison of the spectral power density between ASTM E490, ASTM G173, and the black body emission of a halogen lamp.	98
6.4	Experimental validation of the mechanical design of the solar simulator.	100
6.5	Characterisation of the beam's power density.	101
6.6	Pictures of the solar harvesting module.	102
6.7	Diagram of the electronics inside the solar harvesting module.	103
6.8	Mechanical concepts envisioned for the solar harvester.	105
6.9	Simple model of the harvesting module.	106
6.10	Measured performances of the photovoltaic harvesting module. Data credits: Master thesis of M. Liniger [1].	107
6.11	Picture of the solar arena during the final experiment.	108
6.12	Low-level state machine for the third phase of the experiment.	109
6.13	Qualitative results for the first phase (map building).	111
6.14	Quantitative results for the first phase (map building).	112
6.15	Qualitative results for the second phase (systematic coverage).	114
6.16	Quantitative results for the second phase (systematic coverage).	115
6.17	Measurement of the photovoltaic power, with respect to the (mis)alignment of the panel with the beam.	117
6.18	Proceedings of the second run.	119
6.19	Zoom on the power usage during key moments of the second run.	120
6.20	Box plots of the energy-related metrics for each subsystem, when considering the totality or parts of the third phase (solar recharge).	122
6.21	Summary of the energy consumed by several subsystems, during the three phases of the experiment.	124

6.22 Comparison of the statistical yearly relative sunshine duration in Switzerland, when considering the worst and the best year of the last decade.	126
6.23 Comparison of the statistical monthly relative sunshine duration in Switzerland during February.	126
6.24 Simulated daily energy gained by the solar robot for two distinctive cases of yearly amount of sunshine.	127
A.1 Overview of the electronics.	136
A.2 Results of the calibration.	138

List of Tables

2.1	Summary of the seven robotic vacuum cleaners used during our experiment. . .	10
2.2	Detailed results for the <i>in situ</i> power analysis for each robot and under several workloads.	13
2.3	Detailed results for the global power analysis for each robot.	13
2.4	Detailed results of the navigation analysis, for each robot.	18
3.1	Specifications of the two Computers-on-Module used during this work.	27
4.1	Summary table for the final power analysis of the three sensors.	60
5.1	Summary of the commonly admitted power densities in photovoltaic applications, based on the type of light source and the atmospheric conditions.	66
5.2	Summary of the considered scavenging technologies for indoor service robots.	69
5.3	Calculation of the feasibility boundary for several scavenging technologies, given a fictive service robot consuming 60 Wh per week.	69
5.4	Summary table for the results of case studies 1 to 7.	90
A.1	Characteristics of the power datalogger.	137

List of Acronyms

Notation	Description	Page List
ACML	Adaptive Monte Carlo Localisation	108
ADC	Analog Digital Converter	104, 135
ASTM	American Society for Testing and Materials	74, 97
BT 2.0+EDR	Bluetooth 2.0 Enhanced Data Rate	27
CAN	Controller Area Network	25, 27, 30, 31, 102, 113, 135
CMOS	Complementary Metal–Oxide–Semiconductor	33
COM	Computer-on-Module	27, 31
CPU	Central Processing Unit	102, 109, 110, 113
DoF	Degree of Freedom	77–79
DSP	Digital Signal Processor	25, 27, 33
DVFS	Dynamic Voltage and Frequency Scaling	26
EPFL	École Polytechnique Fédérale de Lausanne	23
FFT	Fast Fourier Transform	55
GPIO	General Purpose Input Output	30, 33
I ² C	Inter-Integrated Circuit	25, 30, 31, 102
IC	Integrated Chip	30
IDE	Integrated Development Environment	25
IFR	International Federation of Robotics	1
IMU	Inertial Measurement Unit	25
IoT	Internet of Things	2
irDA	infrared Data Association	32, 55
LDO	Low-DropOut	30, 31, 33, 54, 55, 103

List of Acronyms

Notation	Description	Page List
LSRO	Laboratoire de Systèmes Robotiques	23, 35, 136
MDF	Medium Density Fibreboard	39, 44
MMU	Memory Management Unit	31
MPPT	Maximum Power Point Tracking	104, 118
MPU	Microprocessor Unit	31
PHY	PHysical Interface	31
PID	Proportional-Integral-Derivative	25
PMIC	Power Management Integrated Chip	31
PWM	Pulse-Width Modulation	25, 104
RAM	Random Access Memory	25, 26
RF	Radio Frequency	66
RFID	Radio-Frequency Identification	25, 66
RTG	Radioisotope Thermoelectric Generator	66
SLAM	Simultaneous Localisation and Mapping	6–9, 12, 14–17, 19, 21, 22, 33, 42, 44, 108, 111
SoC	System on Chip	27, 31
STC	Standard Test Conditions	74
TEG	Thermoelectric Generator	65
TOE	Tonne of Oil Equivalent	2
TRM	Technical Reference Manual	31
UART	Universal Asynchronous Receiver/Transmitter	135
USB	Universal Serial Bus	26, 27, 34, 136

List of Symbols

Notation	Description
A	Surface area
α	Solar altitude
α_{ab}	Seebeck coefficient
α_n	Solar altitude at noon
$\alpha_{n,ss}$	Solar altitude at noon during the summer solstice
$\alpha_{n,ws}$	Solar altitude at noon during the winter solstice
β	Photovoltaic panel inclination (tilt)
δ	Sun's declination
ϵ	Earth tilt axis $\approx 23.45^\circ$
η	Ratio of efficiency
γ	Solar azimuth
γ_p	Photovoltaic panel azimuth
λ	Wavelength
ω	Hour angle from solar noon
Φ	Current latitude
θ	Angle of incidence of direct beam relative to array surface normal
d	Day of year / Distance
d_m	Measured distance
d_n	Nominal distance
E	Energy
f_s	Sampling frequency
G	Irradiance (power per unit area on a surface)
G_n	Irradiance under standard conditions (1000 W/m ²)
h	Solar hour
I	Average current
i	Instantaneous current
N	Size of the statistical sample
P	Average power
p	Instantaneous power

List of Symbols

Notation	Description
T	Temperature / Period of time
t	Time
T_s	Sampling period
V	Average voltage
v	Instantaneous voltage
\hat{x}	Estimator of the value
\tilde{x}	Median of the population

1 Introduction

Robots are increasingly used by our society. Industrial robotic arms have been used for several decades inside factories worldwide to build cars, assemble electronic appliances, or work in harsh and difficult environments [2]. The demand for such robots has experienced a rapid growth, as their sales increased by 12% to 178,132 units sold in 2013 [3]. This expansion is linked to the perceived benefits that they offer: reduced labour costs, the suppression of exhausting or repetitive labour, or increased flexibility [4].

More recently, the market of service robots began to grow due to progresses in several enabling technologies, such as vision systems, navigation algorithms, and versatile mobile platforms. Such robots can be classified into two antagonistic categories. On one hand, service robots for professional use are a specialised market, with few units sold but a high unit value. This includes defence applications, medical robots, and logistic chains. According to the statistics given by the International Federation of Robotics (IFR), approximately 21,000 units were sold in 2013 for a total value of \$3.57 billion [5]. The average price is thus valued at about \$170,000 per unit. On the other hand, domestic robots are mass-marketed with a low unit value. Still according to the IFR, more than 4 million units were sold in 2013 for a total value of \$1.7 billion. This places the average price at only \$425 per unit.

These mass-marketed domestic robots can be further divided between entertainment systems (toy robots, educational devices, etc.) and other domestic service robots, such as floor cleaners and lawn-mowing robots. Figure 1.1 shows the evolution of sold units per year for both these categories between 2009 and 2013. These robots have encountered a strong and steady increase of their sales over the years. The number of sold domestic service robots has more than doubled in five years from 1.05 million to 2.7 million units in 2013. It is hard to know the population currently in use, however, as such robots may be short-lived or used for several years depending on the users [6].

On the millions of domestic service robots sold each year, around 95% are robotic vacuum cleaners [10]. This is, by far, the main application inside households. But in the years to come, other applications have the potential to make a breakthrough inside this market and

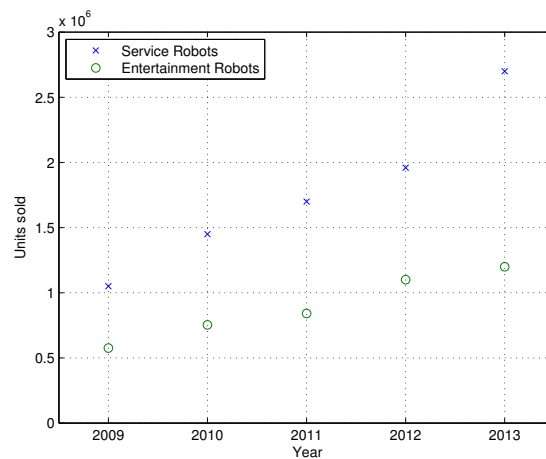


Figure 1.1 – Estimated number of service and entertainment robots for personal use sold between 2009 and 2013. Data source: International Federation of Robotics [5, 7, 8, 9, 10].

consequently further increase sales. Along with telepresence and surveillance robots, there are also numerous possible applications linked to the Internet of Things (IoT). This raises a number of anthropological questions, as this is usually the first encounter with a real robot for most people. A number of interaction paradigms are changed, and new ones are created with a mix of expectations and disillusionments [6]. This also raises a number of important technical questions. The sustained growth of the number of robots operating inside our households challenges the sustainable development of our society. This work aims at shedding light on some of the energy-related topics introduced by these new devices.

1.1 Main Motivations

Let us begin with a simple and concrete scenario. If we imagine that the 4.66 million domestic service robots sold in 2012 and 2013 are still regularly operated, what are their energy requirements? We have to make a number of educated guesses first. We saw that most of these robots are robotic vacuum cleaners. Usage patterns can be very different from one user to the other, but, on average, such a robot is probably used a few hours per week. Let us settle for two hours per week, which is a conservative estimate. This workforce will provide a total of 480 million hours of yearly labour. With an average power consumption of 20 watts¹, we would need at least 9.7 GWh or 830 Tonnes of Oil Equivalent (TOEs)². This does not take into account other energy losses, such as idle power consumption. Even with a small idle power of 1 W, this figure can be multiplied several times.

¹ This hypothesis will be validated in Chap. 2.

² One TOE is the energy released by burning one tonne of crude oil. It equals 41.868 GJ or 11.63 MWh. This does not take into account the efficiency of the power plant. As a consequence, the necessary quantity would be higher in the case of electricity generation.

The sustainment of this fleet represents a lot of energy, and this trend will only get worse in the coming years. Notwithstanding, a lot of other home appliances draw more energy inside households, so why should we focus on robotic vacuum cleaners? Compared to a television or a refrigerator, these robots have some unique and exciting properties that have to be studied. Their ability to move by themselves inside the environment brings some unparalleled opportunities to scavenge for sources of renewable energy. Even if the harvesting of energy is already used by a number of outdoor mobile robots, the indoor environment is a novel challenge.

Energy is also a daily problem for the users of domestic robots. The robot's energy must be temporarily stored inside a battery, but only a limited amount can be taken each time due to the small size and weight that can be embedded on-board. If the robot does not make an efficient use of its energy, it will need to prematurely abort its task to go back to its recharge station. This will cause an increased annoyance to the user, who will probably have to start the operation again after the robot is recharged. It is time to consider the energy as a quantity to be optimised in robotics. The losses should be minimised to the greatest extent possible, while trying to gain energy by using unconventional sources.

1.2 Goals of Our Study

The first goal of our study is to **investigate the energy usage inside existing robots**. This is a necessary step to grasp a deep understanding of the current issues and find technologies to make these robots more energy-efficient in the future. The robots have to carry out more and better work with fewer resources. Towards the accomplishment of this objective, we need to define energy-aware metrics to find the best trade-off among raw performances, functionalities, and a rational use of the resources at hand. We also need to develop simple tools and experimental methodologies, as existing robots often lack the on-board monitoring tools that are available inside research robots.

Our second goal is to **increase the ratio of renewable energies used by domestic robots**. Currently, robots operating inside households exclusively use energy from the main electricity source. With the foreseen increase of their number, it is time to search for new and unconventional sources for their supply. We have to systematically assess the adequacy between potential energies indoors and a realistic user scenario. Because these renewable energies are often intermittent and subject to some spatial and temporal patterns, it is also necessary to develop a theoretical framework regarding their harvesting by such a mobile robot. It is likewise necessary to perform an experimental validation with a real robot to account for unmodeled factors.

This work focuses on domestic robots because they are by number the most important population of service robots. But the scope of this work is in essence not limited to domestic robots, as findings may have an impact for other robots operating indoors, such as those used for professional tasks. An important difference, however, is the usage pattern within such

applications. The time dedicated to the recharge may be considerably lower, compared to a robotic vacuum cleaner used a few times per week. On the contrary, they can embed more costly technologies due to their higher retail price. For mass-marketed robots, we must come up with inexpensive solutions and leverage existing components as much as possible.

1.3 Outline of Our Work

We propose to take a holistic approach within this work, by beginning with a systemic analysis of seven robotic vacuum cleaners available on the market. This study, conducted in Chap. 2, is fulfilled according to our first goal: by performing the partition of the energy used inside actual domestic robots, we aim at creating energy-wise metrics and highlight useful technologies to reduce the energy requirements. In particular, the role of localisation technologies for such applications is studied. Chapter 3 introduces the marXbot robot, which is the experimental platform used throughout this work. Chapter 4 provides an in-depth analysis of the extension modules used for mapping purposes. Chapter 5 reviews the ways to achieve our second goal, with a survey of the possible energy sources and their suitability within our scenario. A theoretical framework is built around the best candidate, and Chap. 6 reports the results of the experimental validation. Finally, Chap. 7 concludes this work.

2 Energy Efficiency of Domestic Robots

Domestic service and entertainment robots have encountered a constant growth in their sales, as discussed in Chap. 1. Between 2009 and 2013, the sales have progressed by more than 150%, from 1.05 million to 2.7 million units. The iRobot Corporation, one of the main players in this market, claims to have sold more than 10 million units of its flagship Roomba robot between 2002 (its first release) and 2013 [11]. The forecast for the period 2014 – 2017 calls for 23.9 million units to be sold worldwide [5]. In 2012, on the 1.96 million service robots that were reportedly sold, 1.9 million were vacuum and floor cleaners [10]. De facto, robotic vacuum cleaners are the main application in most households, as they represent more than 95% of the domestic service robots, and more than 60% of all domestic robots sold in 2012. As sales will continue to progress, there is a clear need to study the energetic impact of such robots¹.

The study of current domestic robots is also crucial for the global objectives of our work. The first interesting question to answer is the following. How much energy does a regular domestic service robot consume? The answer will directly help us to size the energy harvester in Chaps. 5 and 6. Another important aspect is the design of energy-efficient robots. Current mobile robots have a battery to store the energy required to perform the task. This inherently limits the autonomy, due to the limited size and weight. For a vacuum cleaner, this translates into a limited surface that can be covered at one time. Energy is the prime resource when designing a mobile robot, and we should not waste it. There is a clear need to identify technologies that can help us design efficient domestic robots.

This work analyses the current state of the art and level of achievement in domestic robotics, with a focus on robotic vacuum cleaners and energy-related topics. Some research results and design choices for the various functions will impact the energy consumption of the mobile system and thus affect its autonomy. With our work, we aim to highlight the influence of these choices on the energy consumption in order to design energy-wise agents that are compatible with a sustainable growth of the number of robots. For example, with the evolution of technologies, domestic robots shifted from the simple “random walk” approach towards

¹ Parts of this chapter were previously published in [12] and [13].

more evolved navigation schemes, involving a localisation technology at an affordable price. Up to now, no scientific study has analysed the potential impact of these newer robots in terms of energy consumption.

In this chapter, we present an analysis of the performance of several existing robots, assessing the impact of the embedded technologies on the system's fulfilments. We present the results from a three-month experiment performed on a sample of seven robots. The related work is first discussed in Sec. 2.1. The methodology and the main hypothesis are described in Sec. 2.2. We present the robots, the tools, and the experimental setup in Sec. 2.3. The results are carefully analysed in Sec. 2.4, with a focus on several metrics used to compare the vacuum cleaners and assess the impact of several parameters on power and energy consumptions. Finally, a concluding discussion is provided in Sec. 2.5.

2.1 Related Work

Currently in the domestic environment, only a few types of mobile robots have been mass-produced. The first successful product, and now the most widespread, is the robotic vacuum cleaner. The first research related to the design of a robotic vacuum cleaner dates back to the 1980s [14], while the first prototype for domestic use was created in 1991 [15]. In this field, researchers mainly focused on technical aspects. Major research topics include the path planning algorithms used to efficiently cover a given surface [16, 17, 18], the higher level human-aware task planning [19], or an enhanced interactivity to easily control the task to be achieved [20, 21]. As domestic robots share the physical space with people, numerous studies have also addressed the various interactions that can take place, as well as how domestic robots get adopted or rejected [6, 22, 23].

Some comparative studies were also conducted. But up to now, they have addressed mobile domestic robots only from a historical or purely technical point of view [24, 25]. They do not take into account the most recent trends, such as the use of low-cost mapping technologies, nor do they address the energetic aspects.

Other commercial applications of robotics to date have included lawn-mowing, telepresence, pool and gutter cleaning [26]. In the literature, other examples such as assistive [27] or rehabilitation robotics [28] can also be found. Most of the research has focused on key aspects such as the navigation in dynamic environments [29], or more broadly, the Simultaneous Localisation and Mapping (SLAM) problem [30, 31, 32]. Some researchers have studied performance metrics, such as the coverage of several domestic mobile robots performing a random walk [33]. Again, this does not reflect the capabilities of the latest technologies currently available. The question of energy efficiency for these kinds of appliances was only considered recently and only to point out the lack of regulations and standards compared to other home appliances [34].

Our work proposes to fill in the current gap in the state of the art by studying a sample of the latest domestic robots, with a special focus on the energy efficiency of the overall system.

2.2 Methodology

First, we formulate a supporting equation to help in our reasoning around the analysis of domestic service robots. For a specific robot, let $p_{\text{robot}}(t)$ be the instantaneous power drawn from the battery, and T_{task} be the time needed to complete the considered task. The total energy consumed to achieve the task is $E_{\text{total}} = \int_0^{T_{\text{task}}} p_{\text{robot}}(t) dt$. To represent the set of possible configurations, let $\vec{\alpha}$ be a set of generalized design parameters, such as the type of localisation algorithm. These parameters will influence both $p_{\text{robot}}(t)$ and T_{task} . In addition, if we take into account the efficiency of the charging electronics, η_{charger} , the parametric total energy E_{total} becomes

$$E_{\text{total}}(\vec{\alpha}) = \frac{1}{\eta_{\text{charger}}} \int_0^{T_{\text{task}}(\vec{\alpha})} p_{\text{robot}}(t, \vec{\alpha}) dt. \quad (2.1)$$

One goal of our work is to help design more energy-aware devices; that is, we explore the design space $\vec{\alpha}$ in order to minimise $E_{\text{total}}(\vec{\alpha})$. In Eq. (2.1), two functions can be minimised by varying the design parameters: p_{robot} and T_{task} . As we will show, neither is independent, which makes the analysis of the problem non-orthogonal.

The instantaneous power $p_{\text{robot}}(t)$ comes from the “useful” power on one hand and from the losses on the other hand. The required power is minimised by removing useless functions or fusing together several functions, leading, for instance, to a decrease in the number of motors used. Losses are minimised by increasing the robot’s efficiency, for example, by reducing the numerous electrical and magnetic losses inside the motors, as well as by reducing the Joule and switching losses inside the electronics. For a mobile robot, the energy lost when braking also accounts for a part of the total losses, and it can be partially recovered by the addition of appropriate electronics [35]. The overall control, such as obstacle avoidance, is equally important, in order for the robot to follow a smooth trajectory and avoid unnecessary braking. A modified trajectory will, in most cases, influence the completion time.

The other function to be minimised is the completion time T_{task} . In this case, increasing the robot’s speed is often useless because it will increase the instantaneous power accordingly. Better planning and navigation are the keys for this strategy to succeed. When systematic coverage is desired or required, as in cleaning, patrolling, or lawn-mowing tasks, a path planning coupled with a localisation strategy will cut down the coverage time compared to a random walk approach. Recent developments in the semiconductor industry for mobile applications, coupled with algorithmic and mechatronics advances such as low-cost laser scanners [36], have made the SLAM affordable for the mass market. This benefit comes at the

price of extra sensors and computational power needed to achieve an efficient localisation, which conflicts with the reduction of the instantaneous power.

The concern is thus to choose the best technology to reach the better trade-off. Our work will assess, among other factors, the effect of the navigation strategy on the total energy based on measures performed with real mobile domestic robots.

2.3 Experimental Setup

We took a representative sample of the robotic vacuum cleaners currently available on the market. The sample consists of seven robot models, ranging from the low-cost derivatives of the Roomba robot to recent products embedding more complex sensors and algorithms. A summary of some of the technical specifications is shown in Table 2.1. These products target the mass market with an affordable price (between \$250 and \$600). We classified them according to their navigation strategy:

- Robots 1 to 3 (Figs. 2.1a – 2.1c) follow a random walk using some predefined behaviours (wall following, spirals, and obstacle avoidance for the majority).
- Robots 4 to 6 (Figs. 2.1d – 2.1f) perform Ceiling Visual SLAM (CV-SLAM), implementing an algorithm similar to the one described by [37].
- Robot 7 (Fig. 2.1g) is fitted with a low-cost laser range scanner performing 2D Laser SLAM [36].

During this experiment, we explore several performance metrics related to power consumption, navigation strategy, and cleaning efficiency of the robots.

Concerning the figures for the energy consumption used throughout this work, both the global and the intrinsic instantaneous powers were measured. For the instantaneous power, a wireless power datalogger working at 1 kHz was hooked to the battery and used to measure the robot's *in situ* power $p_{\text{robot}}(t)$ during operation (Fig. 2.2a). This module records the power consumed during the entire experiment. It gives a good, but somewhat indirect, insight into the system. This method is simple to put into practice, as the battery is easily removable most of the time. A power analysis bench (Fig. 2.2b) was used to measure the overall energy drawn by the charging station during a complete recharge of the battery. This enables us to take into account the losses of the complete system, including the recharge process.

The navigation is tested in the setup depicted in Fig. 2.2c. It reproduces a two-room apartment with a total internal surface area of 15.5 m^2 made of a concrete floor. Robots are started from their base station located in the upper-right corner. The evolution of the coverage as a function of time is measured using the overhead camera (Fig. 2.2d) and custom tracking software. The

Random Navigation



(a) Robot 1
(Trisa Robo Clean)



(b) Robot 2
(Primotecq Mambo)



(c) Robot 3
(iRobot Roomba 770)

Ceiling Visual SLAM



(d) Robot 4
(Samsung Navibot SR8855)



(e) Robot 5
(LG Hom-Bot VR5902)



(f) Robot 6
(Philips HomeRun FC9910)

Laser SLAM




(g) Robot 7
(Neato XV-11)

Figure 2.1 – Robots used during our experiments, sorted by their localisation technology. Picture credits: (a) – (f) courtesy of RTS/ABE; (g) ©Neato Robotics.

Table 2.1 – Summary of the seven robotic vacuum cleaners used during our experiment. From a global point of view, they differ mainly by their localisation technology (or the absence thereof) and by the technology of their battery. Some other minor differences are omitted here, such as the technology of their short-range sensors.

	Robot 1	Robot 2	Robot 3	Robot 4	Robot 5	Robot 6	Robot 7
Maker	Trisa	Primotecq	iRobot	Samsung	LG	Philips	Neato
Model	RoboClean	Mambo	Roomba 770	Navibot SR8855	Hom-Bot VR5902	HomeRun FC9910	XV-11
Mapping technology	None	None	None	CV-SLAM	CV-SLAM	CV-SLAM	Laser SLAM
Battery technology	Ni-MH	Ni-MH	Ni-MH	Ni-MH	Li-poly	Li-ion	Ni-MH

2.3. Experimental Setup



(a) Robot 7 during the *in situ* power analysis. The visible electronics is the custom power datalogger, connected between the battery and the robot.




(b) Power bench used during the global power analysis (Alciom PowerSpy).



(c) Arena reproducing an apartment, side view.



(d) Arena reproducing an apartment, from the overhead camera. The white ellipse marks an area of difficult access.



(e) Simple arena for cleaning tests.



(f) Robot 4 during the crack cleaning test.

Figure 2.2 – The experimental setups.

experiment is stopped when the robot returns to its base station or when it runs out of battery power.

A simple setup (Fig. 2.2e) is used to measure the cleaning capabilities. It consists of a square surface ($2 \times 2 \text{ m}^2$). Several surfaces were used: 1) smooth concrete; 2) short-pile carpet; and 3) a crack with a cross-section measuring $14 \times 6 \text{ mm}^2$ (Fig. 2.2f). Dust was simulated using a mixture exhibiting a broad granularity range: 5 g of wheat flour, 5 g of fine sand, and 5 g of wood shavings. This was randomly spread by hand over the central square (1 m^2) to avoid interference with the edges. The experiment was stopped when the robot returned to its base station or when the elapsed time reached 7 minutes 30 seconds. The collected material was weighed using a laboratory scale with a resolution of 10^{-3} g . The cleaning efficiency is defined as the ratio between the collected material and the spread material (15 g). Three trials were done for each combination.

2.4 Results

Our results cover several distinct topics but with the same final goal: designing better and more efficient service robots. We will first present the low-level details and show the power patterns that one can encounter in such applications. We will then go to a higher level and analyse the navigation algorithms and the coverage metrics. We will link both aspects together, thanks to a new metric that we call the specific energy. The cleaning efficiency, along with its relation to the power consumed by the robot, is finally analysed on a number of surfaces.

2.4.1 Power Analysis

Efficient analysis of the energy issues related to a robotic system requires that one first knows the relative impact of each subsystem inside the total budget. We want to understand where the power is being used from a systemic point of view. Simple yet informative analysis methods are used to reveal which subsystems correlate with specific power usage.

In Situ Analysis

We can measure the instantaneous power by placing the power datalogger between the robots and their battery. The detailed statistics are given in Table 2.2. Two informative plots are drawn in Fig. 2.3. In Fig. 2.3a, the start-up sequence of the cleaning process for Robot 7 (Laser SLAM) is clearly visible. Starting from the idle state, the following phases can be identified: 1) the laser's spinning motor starts and stabilises; 2) the powerful suction fan starts; 3) the main brush starts to rotate (no side brush); and 4) finally, the robot starts the driving motors and begins to clean. It can be deduced that the Laser SLAM itself consumes only 1.9 W (6.3 % of the total cleaning power) compared to the cleaning subsystem, which takes 23.8 W (78.8 %). The

Table 2.2 – Detailed results for the *in situ* power analysis for each robot and under several workloads. These values are measured by the datalogger placed between the robot and its battery. The “off” state corresponds to the storing state, with the switch turned off when available. The “idle” state happens when the robot is turned on, but motionless. For each piece of data, the mean and the standard deviation are given.

	Robot 1	Robot 2	Robot 3	Robot 4	Robot 5	Robot 6	Robot 7
Off	[W] 0.0068 ±0.10	0.0087 ±0.11	0.064 ±0.18	-0.015±0.13	-0.0019±0.23	0.0075 ±0.32	1.47 ±0.10
Idle	[W] 1.09 ±0.14	2.40 ±0.13	2.93 ±0.30	3.90 ±0.31	2.99 ±0.27	3.85 ±0.39	1.97 ±0.09
Cleaning concrete	[W] 15.6 ±2.78	20.5 ±0.73	13.03 ±1.17	19.98 ±1.74	12.9 ±1.04	23.2 ±2.58	29.95 ±0.64
Cleaning carpet	[W] 16.6 ±3.60	24.5 ±1.64	15.25 ±2.13	22.90 ±2.10	13.7 ±1.29	27.8 ±3.47	30.19 ±0.93

Table 2.3 – Detailed results for the global power analysis for each robot. The power and the energy are measured at the plug of the recharge station by using the power analysis bench.

	Robot 1	Robot 2	Robot 3	Robot 4	Robot 5	Robot 6	Robot 7
Idle power							
– Base station only ^a	[W] 1.20	3.51	1.23	1.94	0.94	0.66	0.40
– Station + robot ^b	[W] 6.13	5.95	4.32	8.06	3.19	3.61	4.63
Recharge energy	[J] 2.30·10 ⁵	2.34·10 ⁵	2.41·10 ⁵	2.19·10 ⁵	1.16·10 ⁵	1.70·10 ⁵	2.35·10 ⁵
Autonomy	[sec] 9464	3807	12099	6211	6399	6144	2870 ^c
Recharge efficiency	0.64	0.33	0.65	0.57	0.71	0.84	0.37 ^c

^a Averaged over 60 seconds

^b Averaged over 24 hours

^c It is impossible to completely turn off this robot, thus measures can be slightly biased due to the self-discharge current.

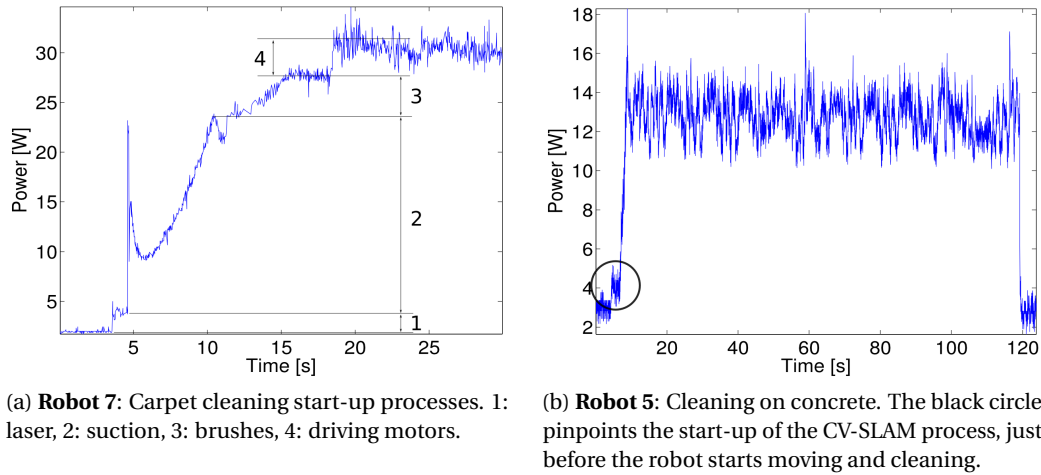


Figure 2.3 – Two plots from the *in situ* power measures.

mobility accounts for 2.5 W (8.3 %). For this specific robot, the power used for the navigation functions is a marginal addition compared to the cleaning device.

Such a clear breakdown, however, is not always possible with our simple method. For example, in the case of Robot 5 (Fig. 2.3b), the first small increase in power (black circle) is due to the Visual SLAM subsystem (camera and algorithm). It takes about 1.1 W of extra power, or 8.4 %, of the total power. This is consistent with the power typically consumed by an embedded processor. All the motors start together and thus cannot be evaluated separately.

The influence of the navigation subsystem on the power budget can be further studied. Figure 2.4a plots the distribution of the *in situ* power consumption of each robot in two working cases. Let us first consider the idle case, when the robot is turned on but not moving. The three robots performing CV-SLAM are, not surprisingly, among the top consumers, as the additional embedded processor will need between 0.5 to 1 W of extra power, even when not processing any data.

When considering the cleaning case, things are completely different. The previous increase, due to the extra processing power, is largely overwhelmed by the difference due to the driving and cleaning motors. Thus, the addition of the SLAM represents only a small part of the total consumption when compared to the energy required for moving and cleaning. As we will see during the analysis of the navigation, SLAM-enabled robots benefit from the acceleration of coverage, saving energy during the overall process.

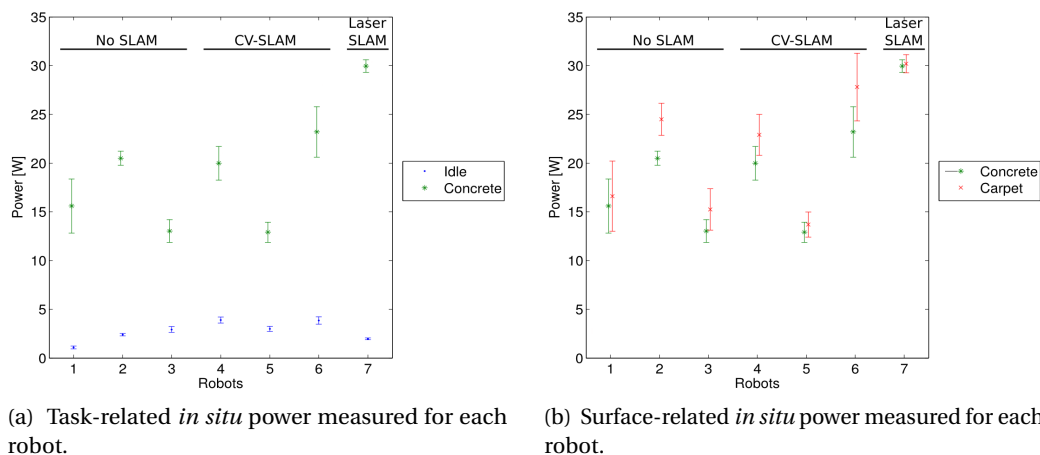


Figure 2.4 – Box plots of the *in situ* power for each robot under several workloads ($N = 1000$).

Figure 2.4b shows the averaged *in situ* power measured for the robots when cleaning two types of surfaces, namely concrete and short-pile carpet². The power distribution clearly shifts upward for all the robots when cleaning the carpet. This is explained by the increased current due to the additional frictional resistance on the cleaning brushes. Cleaning a rough surface requires more power.

Global Analysis

A power analysis was also performed directly at the plug of the recharge station. Detailed results are given in Table 2.3 (p. 13). One initial observation is the high idle power of the sole base station, with the worst result noted in the case of Robot 2 (up to 3.5 W). When the charged robot is left connected to its base station, the result further deteriorates, as power consumption increases to between 3.2 and 8.1 W depending on the robot. Unfortunately, this type of appliance is not, at present, bound by any regulations similar to the European regulation 1275/2008 [38], which limits the standby mode to 2 W. This represents a serious concern for these types of mass-produced electrical appliances.

The efficiency of the recharge station η_{charger} was computed as the ratio between the energy entirely consumed by the robot, and the amount of energy injected into the system during a recharge cycle. For a given robot (fixed set $\vec{\alpha}$ of parameters), Eq. (2.1) can be reorganized to

$$\eta_{\text{charger}} = \frac{1}{E_{\text{total}}} \int_0^{T_{\text{task}}} p_{\text{robot}}(t) dt, \quad (2.2)$$

² All robots were unable to clean a long-pile carpet due to the small distance between the ground and the robot's frame.

where E_{total} is the energy consumed at the plug to recharge the robot, and $p_{\text{robot}}(t)$ is the *in situ* power measured on the robot during the whole process.

This efficiency figure varies between 0.33 and 0.84 in our study and includes the intrinsic quality of the charger as well as the storage efficiency of the battery subsystem. The best two robots are the ones that use Li-ion batteries, while the others use the Ni-MH technology. In the case of the two worst robots, more than 50 % of the recharge power is already lost at the plug. An efficient design must take into account the whole system, and not only the domestic robot.


2.4.2 Navigation Analysis

The trajectories were recorded using an overhead video tracking system when the robot was engaged in cleaning the apartment space shown in Fig. 2.2c (p. 11). The samples of each trajectory in Fig. 2.5 clearly show distinct strategies for each robot. For example, Robot 1 is performing a random walk by using a mixture of straight lines and spirals (Fig. 2.5a), and it stops the spiral when it encounters an obstacle. This strategy is efficient to cover large empty spaces, but it fails in more complicated environments, and some areas can be left uncovered. The superiority of systematic navigation becomes clear by looking at Figs. 2.5d to 2.5g, even if the SLAM-enabled robots all take very different approaches from each other.


To estimate the surface covered, the image analysis integrates over time the surface hidden by the robot's shape. For this purpose, the Gaussian mixture-based background segmentation of [39] is applied to the calibrated pictures. This estimation does not take into account the side brushes used by most robots or the width of the main brush under the robot. In the results, hidden areas (such as under the sofa) are not taken into account. Detailed statistics are given in Table 2.4.

The evolution of the coverage, as a function of time and averaged between all the trials, is plotted in Fig. 2.6a (p. 19). The SLAM-enabled robots seem much faster than the others. This is confirmed by the computed completion times shown in Fig. 2.6b (p. 19). Localisation-less robots have no robust way to compute the achieved coverage. Consequently, they do not return to their base stations when the coverage reaches a steady state, and most of the time, they dock to their stations after an extended period of time. On the contrary, Robot 6, which is the slowest among the robots performing SLAM, is three times faster on average than the random walk robots.


Regarding the achieved coverage, random walk robots take time but achieve a robust coverage. On the other hand, Robot 4 and, to a lesser degree, Robot 5 (both relying on vision), underperform compared to the others. Looking back at the image analysis, it appears that some places are harder to reach for them. One of these places is between the sofa, the intermediate wall, and the bin (white ellipse in Fig. 2.2d on p. 11). In half of the runs, Robots 4 and 5 were unable to reach this place, thereby losing part of the coverage. On the contrary, Robot 6 was successful on all of its 11 runs because its path planning uses thinner bands, as shown in Fig. 2.5f. While

Random Navigation


(a) Robot 1




(b) Robot 2




(c) Robot 3

Ceiling Visual SLAM


(d) Robot 4



(e) Robot 5



(f) Robot 6

Laser SLAM

(g) Robot 7

Figure 2.5 – Sample of the trajectories for each robot, grouped by localisation strategy.

Table 2.4 – Detailed results for the navigation analysis for each robot. The coverage ratio is based on the image analysis given by the overhead camera. For each piece of data, the mean and the standard deviation are given.

	Robot 1	Robot 2	Robot 3	Robot 4	Robot 5	Robot 6	Robot 7
Coverage time [sec]	3838 ±1467	3737 ±428	3597 ±174	881 ±125	551 ±63	1104 ±40.0	991 ±38
Coverage reached	0.962 ±0.043	0.973 ±0.003	0.972 ±0.009	0.850 ±0.061	0.882 ±0.065	0.976 ±0.005	0.979 ±0.001
Specific energy [J/m ²]	4273 ±3250	6070 ±1169	3642 ±749	1531 ±527	553 ±175	2029 ±349	1970 ±140

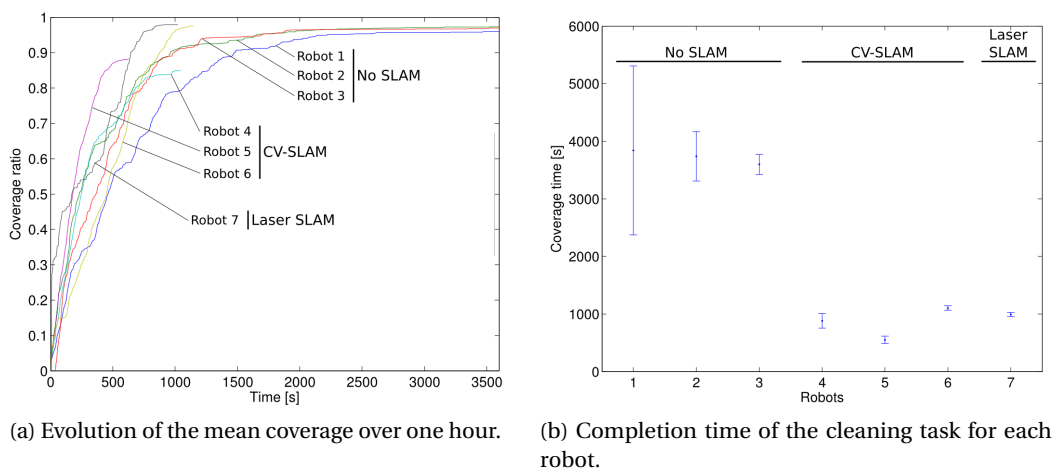


Figure 2.6 – Analysis of navigation performance for the seven robots.

some time is lost by this strategy, it gains greatly in robustness and coverage. There is clearly a trade-off to be found here, as trying to lower the required energy by all means could result in poor adaptability and robustness with respect to complex environments.

2.4.3 Specific Energy

One of our key questions is the influence of design parameters on the energy consumption, and especially the navigation strategy. To answer this question, we now compare the coverage strategy with respect to the energy. For this, we define the specific energy, which is the energy needed to cover 1 m^2 of floor. This metric makes sense for robots where a systematic coverage is required, and it can be used to make a rough estimate on the energy required for a given environment. Of course, in addition to the area, other elements can play a role. Obstacles, transitions between rooms, or dead ends will potentially require more energy to be overcome. But the covered area remains the driving factor for vacuum cleaning robots and other similar service robots.

From Eq. (2.1), we know that the energy totally consumed by the robot can be expressed as $\int_0^{T_{\text{task}}} p_{\text{robot}}(t) dt$, therefore, we have

$$E_{\text{specific}} = \frac{1}{A_{\text{effective}}} \int_0^{T_{\text{task}}} p_{\text{robot}}(t) dt, \quad (2.3)$$

where $A_{\text{effective}}$ is the surface effectively covered, as deduced from the previous coverage analysis.

Figure 2.7 shows the specific energy for the robots of our sample. Detailed statistics are also given in Table 2.4. With this metric, we can demonstrate the effectiveness of the SLAM-enabled




Figure 2.7 – The specific energy, in J/m², for each robot.

robots over those using random-walk methods, counterbalancing the increase of power by a drastically reduced T_{task} . However, no clear conclusion can be drawn between CV-SLAM and Laser SLAM robots, even if Robot 5 (CV-SLAM) is consistently two times more energy efficient than the other robots.

2.4.4 Cleaning Efficiency

From the user's perspective, cleaning efficiency is one of the most important factors that determines the usefulness of a robotic vacuum cleaner. Efficiency figures were measured on three different surfaces and averaged on three trials. These figures were then compared to the averaged *in situ* power previously measured. Plots are given as a function of the surface's type in Figs. 2.8a to 2.8c.

For the concrete surface (Fig. 2.8a), no relationship between the cleaning efficiency and the robot's power can be seen (using a linear fit, goodness of fit R^2 is only 0.018). Most robots scored about 90 % with respect to the amount of material collected during cleaning. For comparison, the same test conducted with a manual vacuum cleaner (Dyson DC05) showed an efficiency above 98 %.

In the case of the carpet (Fig. 2.8b), cleaning efficiency $\eta_{\text{cleaning,carpet}}$ exhibits a moderate dependency on the robot's power, as shown by the linear regression

$$\eta_{\text{cleaning,carpet}} = 0.007 \cdot \bar{p}_{\text{robot}} - 0.092 \quad (R^2 = 0.56) ,$$

where \bar{p}_{robot} is the averaged *in situ* power. The overall efficiency remains poor (below 35 %) in all cases.

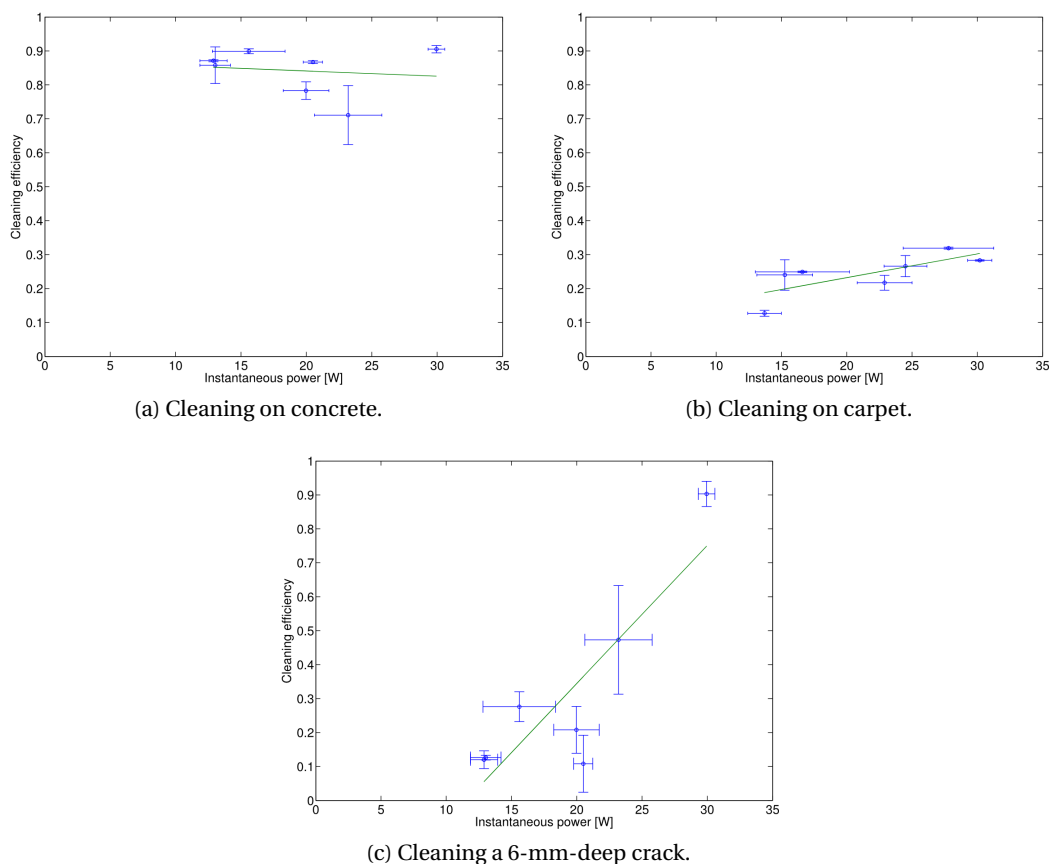


Figure 2.8 – Analysis of the cleaning task. The plots show the cleaning efficiency on a variety of surfaces as a function of the robot’s averaged *in situ* power. The green line shows the linear regression performed on the data ($N = 3$).

Finally, for the cleaning of the $14 \times 6 \text{ mm}^2$ crack (Fig. 2.8c), data are moderately explained by a linear regression

$$\eta_{\text{cleaning,crack}} = 0.041 \cdot \bar{p}_{\text{robot}} - 0.47 \quad (R^2 = 0.74) .$$

Robot 7 had a good score, but it is also the one with the highest instantaneous power.

In conclusion, the suction power does not really help when dealing with a flat and smooth surface. The design of the brushes is the primary concern in this case. However, suction power does become the main tool on a hard and uneven surface. Current robots do not adapt themselves with respect to the surface they are on, but it would be conceivable to modulate the suction power as a function of the actual ruggedness. This would spare energy on the vast majority of the cases, while keeping the possibility to clean rough surfaces with a decent efficiency.

2.5 Discussion

Based on *in situ* and global power measurements of existing products, the influence of some technologies could be determined. From this analysis, it is established that robots required to fully cover an area have a clear advantage, from the energy point of view, if they can rely on an embedded SLAM system. Ceiling Visual SLAM and Laser SLAM are clearly advantageous for energy consumption in that they reduce the coverage time, and they have a relatively small impact on the instantaneous power. However, no conclusion could be drawn between the two SLAM systems.

If we take a systemic approach, the entire energy chain has to be considered, starting from the charging up to its final usage. As noticed during the experiments, the efficiency of commercial charging stations is rather low, with important losses in the electronics. Li-ion batteries, when properly used, also benefit from the efficiency of the energy storage. However, because they are more expensive compared to Ni-MH technology and need dedicated recharge electronics, their adoption is for now limited in such devices.

Regarding the total coverage, some of the SLAM-enabled robots underperformed compared to the others. This results from the path planning strategy, as some locations can be missed if only a coarse displacement is performed. It is believed that such problems are a matter of fine-tuning the system. To add a Visual SLAM-based navigation into a product, the price is mainly driven by the cost of the processor and camera boards. As the smartphone market expands, the cost of such components rapidly decreased during the past years.

Finally, we have the necessary elements to answer this question: Compared to a manual vacuum cleaner, does a robot perform better in terms of energy? We will focus on the setup of Fig. 2.2c (p. 11). With a robot consuming roughly 20 W (Robots 1 to 5 on a concrete floor), it will take around 15 minutes for the fastest one, equaling an amount of 18 kJ. On the contrary, a human will take about 5 minutes with a vacuum cleaner, consuming usually 1000 W or more, equaling an amount of at least 300 kJ. The robot consumes roughly 15 times less energy per cleaning. If the robot completely replaced the traditional vacuum cleaner, its usage every day of the week would spare about half of the energy consumed by a weekly manual vacuuming.

In reality, the robot will usually never completely replace the manual vacuum cleaner, as it cleans less efficiently. For most users, the robot will engender a decreased usage of the manual vacuum cleaner for spot cleanings, and it will help to keep an overall cleaned state by doing frequent cleanings. However, actual figures depend on the considered group of users.

2.6 People Who Contributed to This Work

This work was made possible thanks to David Hamel, who performed all the trials and the measures. Part of the robots were lent by Télévision Suisse Romande (TSR).

3 Robotic Platform

In the previous chapter, we showed the significance of the energy considerations for the efficient design of domestic robots, as well as how the use of a localisation technology can help reduce energy requirements. However, to push our analysis further, we need a more modular and versatile robotic platform to experiment with various concepts, ideas, and algorithms. This chapter presents the robotic platform that will be used throughout this work. We introduce the existing modules, explain the major features and summarize the main characteristics.

3.1 The Standard MarXbot Robot

The marXbot is a small form factor, modular robot developed between 2007 and 2010 by the Laboratoire de Systèmes Robotiques (LSRO) at the École Polytechnique Fédérale de Lausanne (EPFL). During the development, the main focus was on collective robotics, but it can be also used in a variety of research fields, thanks to the standardized interface between the modules. For example, modules were developed for the collective attachment of swarm robots [40], the handling of magnetic building blocks [41, 42], the structured deposition of thermoplastic, and even imprinting experiments between a robot and young chicken [43].

The standard version of the marXbot is shown in Fig. 3.1b. It includes, from bottom to top: the mobile base, the attachment module, the range and bearing module, the infrared rotating scanner, and the main computer. We will not use the attachment module, nor the range and bearing module. A complete overview of these modules, and the robot itself, can be found in [44, 45].

In this section, we will summarize the main design decisions concerning the subsystems that play a key role in the remaining chapters of this thesis. We will focus on the mobile base, the low-level and high-level controls, and the power management. The long-range sensors will be covered in Sec. 3.2.




Figure 3.1 – The marXbot robot.

3.1.1 Mobility

The marXbot is propelled by its mobile base, shown in Fig. 3.1a, which is a self-contained module. It has an outer diameter of 17 centimetres, which is rather small compared to other robots with similar capabilities. The motion is produced by two geared DC motors, arranged in a differential configuration. The output of each gearbox drives a treel, which is a combination of a wheel and a track. The wheel offers good precision during the dead reckoning localisation, while the track offers versatile mobility in a wide variety of terrains. In reality, the contact point between the treel and the ground is not always well defined, leading to an increased error on the odometry compared to a wheeled robot.

Each motor is driven separately by its own microcontroller. The control is achieved at the software level by three nested Proportional-Integral-Derivative (PID) controllers controlling the current, the velocity, and the position¹. The controller runs at 1 kHz, and the motor's state is computed based on the measured current, as well as the digital position encoder. The control at the electrical level is achieved by generating a Pulse-Width Modulation (PWM) to drive the power electronics.

3.1.2 Short-Range Sensing

When using the standard configuration of the marXbot, the environment close to the robot is sensed by 24 short-range sensors mounted on the mobile base. They are working based on the reflected intensity of a pulsed infrared beam. They have a usable range of 5 centimetres, which is enough to detect and avoid obstacles but is too short-sighted for mapping applications.

Other sensors are also available, even if they are not used in the scope of this work. We have, for example, ground infrared sensors, an Inertial Measurement Unit (IMU), a Radio-Frequency Identification (RFID) tag reader, or a set of microphones. The interested reader can refer to [44] for more details.


3.1.3 Low-Level Design of the Modules

Each module is completely independent from the others inside the robot. The control at the level of the module is done by one or more microcontroller(s). All the microcontrollers used on the marXbot are from the Microchip dsPIC33 family. They are versatile 16-bit microcontrollers, with an integrated Digital Signal Processor (DSP) core to perform advanced computations.


These microcontrollers are networked together using a Controller Area Network (CAN) bus. This allows the easy exchange of messages between modules, as well as with the high-level control if present. The CAN bus is a multi-master self-arbitrated differential bus [46], enabling each module to communicate with its neighbours without a centralized authority. The microcontrollers are also networked using an Inter-Integrated Circuit (I²C) bus, which is a master-slave synchronous bus [47]. This bus is used for some specific purposes, such as waking up microcontrollers from their hibernation state.

A new technology was necessary to ease the development of the embedded programs inside such a network of microcontrollers. Aseba was developed with such a goal in mind. It allows the modular development of complex behaviours by relying on the execution of bytecode inside a lightweight virtual machine executing on each microcontroller [48]. The code is easily developed with the provided Integrated Development Environment (IDE) and loaded inside the microcontrollers' Random Access Memory (RAM) with a one-click step.

¹ More details can be found inside "Rétornaz P., Magnenat S., Vaussard E.: MarXbot User Manual, (2012)", available online (last accessed: 5th September 2014): <http://mobots.epfl.ch/data/robots/marxbot-user-manual.pdf>.



(a) LSRO original computer board, based on a Freescale i.MX31 processor.



(b) LSRO new computer board, based on a TI OMAP3 processor.



(c) Gumstix Overo Computers-on-Module.

Figure 3.2 – Overview of the processing units.

3.1.4 High-Level Control

The high-level control is achieved by an embedded ARM processor located on a dedicated module. The first version of the marXbot is using a Freescale i.MX31 (ARM1136v6k) running at 532 MHz. This processor is assembled on a small processor board, alongside 128 MB of RAM and 32 MB of NOR Flash memory to store the bootloader and the Linux kernel. This processor board is then connected on a motherboard, which is providing the necessary expansion connectors such as the Universal Serial Bus (USB) and electrical connections with the rest of the robot. The electronics is shown in Fig. 3.2a.

A new version of the processor board was designed for this work in order to expand the capabilities by using newer hardware. A special focus was also placed on better low-power performances by means of revised electronics and enhanced Linux support for low-power states, Dynamic Voltage and Frequency Scaling (DVFS), and similar features.

3.1. The Standard MarXbot Robot

Table 3.1 – Specifications of the two Computers-on-Module used during this work. Both are manufactured by Gumstix, and they are compatible with the motherboard designed during this work.

Model	Processor		Memory		Wireless	Price (US \$)
	Name	Frequency (MHz)	RAM (MB)	NAND (MB)		
EarthSTORM	TI AM3703	800 MHz	512	512	–	109.00
FireSTORM	TI DM3730	800 MHz	512	512	WiFi 802.11b/g BT 2.0+EDR	189.00

The core of the new version is architected around a commercially available Computer-on-Module (COM) called Overo, which is sold by Gumstix Incorporated². The advantage of this solution is the flexibility offered by the large choice of compatible modules (17 as of this writing), offering various processors' variants, memory options, or wireless communications. For the present work, two variants were selected. They are shown in Fig. 3.2c, and the principal characteristics are detailed in Table 3.1. The main difference is the availability of the wireless communication electronics, enabling the IEEE 802.11 WiFi and Bluetooth 2.0 Enhanced Data Rate (BT 2.0+EDR) wireless channels in the case of the Overo FireSTORM. At the level of the System on Chip (SoC), both are very similar and based on an ARM Cortex-A8 core with the following on-chip cache architecture: 32 kB L1i, 32 kB L1d, and 256 kB L2. The AM3703 found in the Overo EarthSTORM is lacking the TMS320C64x+ DSP core, as well as the POWERVR SGX hardware graphics accelerator, but both are unused in the scope of this work. These on-chip hardware blocks are not powered, and the corresponding clocks are disabled. As a result, the power drawn by both processors' variants is expected to be almost identical.

This COM is embedded on a new motherboard (Fig. 3.2b), which is similar in appearance to the first design, as most of the external peripherals were kept. A major difference is the addition of an on-board USB hub in order to provide up to four USB ports to interface common robotics hardware (vision camera, Microsoft Kinect, etc.). Other differences include a reworked CAN interface for better performance, and a reworked voltage regulators hierarchy to further decrease the consumption when entering low-power modes.

The processor is running with a mainline Linux kernel³ (version 3.16 as of this writing), and a customized user space root filesystem based on an ARM port of Ubuntu 14.04⁴.

² More information: <http://www.gumstix.com>.

³ Most of the patches developed during this work were successfully merged upstream.

⁴ More information: <https://wiki.linaro.org/Platform/DevPlatform/Rootfs>.




Figure 3.3 – Overview of the battery pack.

3.1.5 Energy Management

The energy is at the core of our work, and it cannot be wasted. But within such a complex system, with microcontrollers and power consumers distributed all around the robot, the energy must be intelligently managed at all of the robot's levels.

This endeavour begins with the battery pack, which embeds a high-precision Lithium battery monitoring electronics, based on the Maxim DS2764 integrated circuit⁵. The current measurement has a resolution of 3.125 mA, with a dynamic range of ± 12.8 A and a sampling frequency of 1456 Hz. The voltage measurement has a resolution of 4.88 mV. The available energy with a new battery pack is 37 Wh (or $135 \cdot 10^3$ J), with a rated current of 10.4 A. The assembled battery pack is shown in Fig. 3.3a, and the inner details are shown in Fig. 3.3b.

The form factor of the battery, coupled with sliding contacts, allows the battery to be easily exchanged. It was shown that a continuous operation is possible, 24 hours per day, when used in conjunction with a custom automatic battery charger [49]. This can be advantageous for certain type of robotic experiments, but in the present work, we are trying to achieve zero-replacement through the use of harvested energies. This exchange feature is thus unused.

Now let us consider the global picture. The relation between the power providers and power consumers can be summarized inside a power tree to show the hierarchy and the dependencies between the modules and the components. Figure 3.4 shows such a power tree for a marXbot robot constituted by the mobile base, some generic modules, and finally the processor board on top. These three parts have specific behaviours with respect to low-power operations.

⁵ We describe here the voltage and current measurement capabilities of this integrated circuit, but this feature is unused within the scope of our work. Indeed, a measurement at high frequency could cause itself an increased power consumption, thus biasing the results. An external datalogger is used instead, as described in Appendix A.

3.1. The Standard MarXbot Robot




Figure 3.4 – Simplified diagram of the power distribution network for a marXbot robot constituted by the mobile base (lower part), a generic kind of module (middle part), and the processor board (upper part).

Mobile Base

The mobile base hosts the battery pack, which is providing the power at a voltage varying between 3.6 V (nominal voltage) and 4.2 V (maximum rated voltage). The power flows from the battery by going through the first motor module, then the proximity sensors module, and finally the second motor module.

The numerous Low-DropOut (LDOs) regulators along the way are powered by the battery, but they may not be enabled (thus in low-power state). The LDO powering the microcontroller on each module is disabled by default and can be enabled only by an I²C access to the I²C General Purpose Input Output (GPIO) expander located on the proximity sensors module. The user has the choice to make such an access from the main processor, from any microcontroller on another expansion module, or from both. For practical reasons, the I²C GPIO expander is always powered on, but the MCP23008 Integrated Chip (IC) has a maximum supply current of 1 μ A in idle state. The remaining LDOs are directly controlled by the corresponding microcontrollers, and they will be kept in low-power state as long as the controlled functionality remains unneeded. The LDOs used in the mobile base (TI LP2992) have a typical idle consumption of 50 nA.

Apart from the GPIO expander, another component is always powered on by design: the power electronics of each motor. This component (SI9986) has an idle consumption of 55 μ A at 4 V, resulting in a power loss of 230 μ W. With a total of 6 ICs, this gives a total of 1.38 mW for the typical losses.

To summarize, when the mobile base is in low-power mode, it will at least consume 1.4 mW, the driving factor being the power electronics of the motors. As a comparison, this is what is dissipated by a 10 k Ω resistor connected to a 4 V supply.

Intermediate Modules

Figure 3.4 depicts a generic extension module on top of the mobile base. As can be seen, there is also a generic mechanism to turn off the LDO powering the microcontroller. The designer of the module is free to use it or not. In both cases, it is also possible to use software to ask the microcontroller to enter into a deep sleep mode by sending it a special CAN packet. Once in this state, it can be woken up only by sending an I²C command to the microcontroller⁶. On the marXbot's modules, the most common family of microcontrollers is the Microchip dsPIC33FJxxxMCxxx, with a typical sleep current of 211 μ A for the core. One should also take into account the consumption of the I²C peripheral, as well as the power dissipated by the LDO. In overall, the expected power consumption should be less than 4 mW in deep sleep.

⁶ Another implementation-specific interrupt would be also possible if no I²C bus is available.

If the module has other functionalities, the corresponding LDO or DC/DC converters should be controlled by the on-board microcontroller to turn off the features when unused. In any case, the regulator should be turned off when the microcontroller enters the sleep state.

Processor Board

The processor board is constituted by the Gumstix Overo COM mounted on top of the motherboard. The processor in itself is powered by its own set of regulators, known as the Power Management Integrated Chip (PMIC). This is a complex piece of hardware that provides, among other features, one boost converter, three buck converters, and ten LDOs⁷. It interacts with the processor through a dedicated high-speed I²C bus and a few control signals to provide the correct power sequence when switching power states.

The processor is even more complex. The ARM Cortex-A8 core shares the silicon die with dozens of other peripherals inside the SoC. The Microprocessor Unit (MPU) inside the TI AM/DM37xx SoC has five power domains with associated clocks⁸. Considering that each power domain has four possible states (active, inactive, retention, and off), the MPU alone has 14 valid operating modes⁹. This does not even take into account the other major peripherals, such as the interconnect buses, the memory subsystem, or the Memory Management Unit (MMU). The SoC has a total of 16 voltage domains, 18 power domains, and numerous clock sources. The kernel of the operating system — Linux in our case — is in charge of enabling and disabling the resources when it deems necessary. The detailed analysis of such a system is beyond the scope of this work, but we will strive to extract power figures for well-defined operating conditions.

The motherboard has five LDOs and one boost converter, as shown in Fig. 3.4 (p. 29). Most of them are always turned on, as the processor relies on them to communicate with the rest of the robot, as well as with the outside world. However, the three LDOs controlling the camera power domain can be controlled by the processor. The processor also controls two individual power-hungry components: the CAN Physical Interface (PHY) and the sound amplifier.


3.2 Extension Modules for Mapping

The basic version of the marXbot robot is severely short-sighted. Its proximity sensors can only sense within 5 centimetres. To overcome this limitation, several scanning modules were developed by numerous people within the scope of several other projects. In this section, we will describe three extension modules targeting mapping applications. They all have radically different designs, and they all embody the current state of the art in their respective category.


⁷ More information on the TI TPS65950 PMIC can be found inside the Technical Reference Manual (TRM) at <http://www.ti.com/product/tps65950>.

⁸ They are ARM Core Logic, ARM L2 RAM, NEON Engine, MPU Interrupt Controller, and Debug Logic.

⁹ More details can be found inside the AM/DM37xx TRM at <http://www.ti.com/product/dm3730>, and inside the ARM Cortex-A8 TRM at <http://infocenter.arm.com/help/topic/com.arm.doc.subset.cortexa.a8>.



(a) Simplified diagram of the power usage inside the module.



(b) Picture of the infrared rotating scanner.

Figure 3.5 – The infrared rotating scanner module.

3.2.1 Infrared Rotating Scanner

The infrared rotating scanner was initially developed for the Swarmanoid project [40] and it is described in details in [44, 50]. Figure 3.5a presents a simplified diagram of the module, focusing on the power distribution chain inside the module. A picture of the module is shown in Fig. 3.5b.

This module is split in two parts. The primary assembly is fixed to the robot and hosts the main microcontroller. The secondary assembly is supported by the first one, and a motor actuates its rotation along a vertical axis. A slip ring is usually used in such cases to transmit power and data between the stationary and the rotating parts. But in order to reduce the cost and increase the lifetime of the module, an air-core transformer is used instead to transfer energy from the primary part to the secondary part. It uses a resonant inductive coupling at 228 kHz and achieves an overall efficiency of 69%. The bidirectional half-duplex communication is achieved with infrared Data Association (irDA) transceivers on both the fixed and the rotating parts.

The infrared range sensors are embedded into the rotating assembly. Two sets of sensors are used: short-range sensors (40 – 300 mm) and long-range sensors (200 – 1500 mm). The two sensors inside the same set are placed back to back, while the two sets are shifted by 90° . It is thus possible to cover the whole range between 40 and 1500 mm by combining the readings from both sets. This only works if the robot moves at low speed, however, due to the time shift caused by the angle of 90° between the two sets of sensors. Each sensor is sampled at 60 Hz, leading to a total bandwidth of 240 Hz. The rotation speed can be adjusted, therefore changing the angular resolution.


The air-core transformer is an interesting idea. But its low efficiency, combined with the multiplicity of infrared sensors, makes a potentially high power usage. This will have to be validated in practice. We can finally note that this is not the only rotating module with infrared sensors, but this is the most advanced one. The work of [51] achieved a low-cost module (under \$100), but with a bandwidth of only 30 samples per second.

3.2.2 Neato Laser Range Scanner

Neato Robotics is an American company that developed a family of robotic vacuum cleaners. Neato's robots feature an innovative laser range scanner, coupled with a particle-based Simultaneous Localisation and Mapping (SLAM) algorithm [36, 52, 53]. The Neato XV-11 has already been tested in Chap. 2, and it was among the fastest robotic vacuum cleaners. An extension module for the marXbot was developed based on the sensor extracted from one of the Neato robotic vacuum cleaners and a reverse-engineered communication protocol. Figure 3.6a presents a simplified diagram of the module, and Fig. 3.6b shows a picture of the module mounted on top of the processor board. Compared to other laser sensors with similar characteristics, this sensor is cheaper by an order of magnitude: its fabrication costs \$30 [36], and the robot is sold for less than \$400.

This laser sensor computes the distance based on the triangulation of a 650 nm-infrared laser beam. The emitting laser, the Complementary Metal–Oxide–Semiconductor (CMOS) camera and the DSP are all embedded into the rotating part, while the power and data are transmitted through a slip ring. The sensor is able to take scans with a field of view of 360° and an angular resolution of 1° .

The extension module is an interface between the marXbot and the sensor. The data coming from the sensor are read by the microcontroller through a serial link, and the rotation speed is controlled by the microcontroller. The power is shared between the microcontroller and the sensor. Like other extension modules, the LDO supply can be disabled. But unlike other modules, it is not designed to be stacked onto the mobile base but to be mounted on top of the processor board. As a consequence, the enable signal is directly provided by one of the GPIOs of the processor.



(a) Simplified diagram of the power usage inside the module.



(b) Picture of the Neato module, mounted on top of the processor board.


Figure 3.6 – The Neato scanner module.

3.2.3 Hokuyo Laser Range Scanner

The laser range sensors manufactured by the Japanese company Hokuyo Automatic Co. have been successfully used for several years by the robotics community. They are compact, lightweight, and offer a good precision with a high bandwidth. A three-dimensional scanning module was made for the marXbot with a Hokuyo URG-04LX mounted on a horizontal rotation axis. Figure 3.7a presents a simplified diagram of the module, and Fig. 3.7b shows a picture of the module mounted on top of the mobile base.

This laser sensor computes the distance based on the phase difference between the emitted and received 785 nm-infrared laser beam. It uses two waves, modulated at 46.55 and 53.2 MHz [54]. The scanning device uses a lightweight rotating mirror, while the rest of the optics is fixed inside the housing. This makes scanning at a high frequency (10 Hz) possible. The sensor provides scans with a field of view limited to 240° but a high angular resolution of 0.36° .

The developed module acts as a simple interface between the marXbot and the Hokuyo sensor. Due to the high bandwidth of the sensor, the data are not read by the microcontroller but sent directly over a USB link to the processor. Like other extension modules, the power can be disabled in order to save energy. This also cuts the power of the sensor by shutting down



(a) Simplified diagram of the power usage inside the module.



(b) Picture of the Hokuyo module, mounted on top of the mobile base.

Figure 3.7 – The Hokuyo scanner module.

the DC/DC supply. The horizontal rotation axis is actuated by a DC motor, which is directly controlled by the microcontroller of the module.

3.3 People Who Contributed to This Work

The marXbot was developed by a team of people at the LSRO laboratory. Michal Bonani designed the mechanics, in collaboration with Tarek Baaboura. He also designed some of the electronics, in collaboration with Daniel Burnier, Philippe Rétornaz and Simon Fivat. The embedded software is the work of Philippe Rétornaz, with some help from Stéphane Magnenat for the control framework. I designed the first version of the mobile base's electronics, as well

Chapter 3. Robotic Platform

as the second version of the processor board. The first version of the processor board was designed by Daniel Burnier and Valentin Longchamp. I also performed the Linux support for the second version.

The infrared rotating scanner was developed by Michael Bonani, Philippe Rétornaz, and Tarek Baaboura. The air-core transformer was done with the help of Paolo Germano. The Neato module was engineered by Michael Bonani, and the embedded software by Philippe Rétornaz. The Hokuyo was designed by Daniel Burnier (electronics) and Philippe Rétornaz (embedded software), in collaboration with ETH Zürich (mechanics).

4 Characterisation of Sensors for Energy-Efficient Mapping

In Chap. 2, we performed a coarse power analysis of robotic vacuum cleaners with the aid of a datalogger between the battery and the robot. We were able to gain a better understanding of the situation with the analysis of energy-related metrics. We also proved how the addition of a localisation technology can reduce the energy needs in the long run. This chapter will specifically focus on the sensing capabilities that are required for mapping applications.

Chapter 3 presented the robotic platform that will be used throughout this work. Three mapping modules were developed for the marXbot robot. All three modules use a rotating scanner, but each one has some unique specificities and performance. In this chapter, we will characterise these sensors and determine the influence of several external factors on the quality of the measures.

We will also perform a detailed power analysis of each sensor to advance towards energy-efficient robots. This analysis will be useful to compare the sensors between them and find their advantages. A good sensor should consume the least energy possible, however, a low-power sensor may have poor performance. Thus we need to define metrics that take into account not only the raw performance but also the energy required to reach this result. A tradeoff must be found.

We will first discuss the related work in Sec. 4.1, and describe our methodology in Sec. 4.2. Two sets of experiments are performed. First, the characterisation of the sensors is conducted in Sec. 4.3. We will of course analyse important attributes, such as the accuracy, but the impact of several external factors is also studied. A power analysis is then done for each sensor in Sec. 4.4. These two experiments are linked together in Sec. 4.5, principally by defining new metrics to take into account both the performance and the power consumption. We will finish with the discussion in Sec. 4.6.

4.1 Related Work

Our first concern is the characterisation of mapping sensors. A number of metrics can be used to assess the quality and usefulness of a sensor for such applications. The range and the accuracy are of course two important aspects to evaluate. It is important to be aware that a number of environmental parameters have an influence on the accuracy. For example, the angle with respect to the object, the object's material, or the sensor's thermal drift can all play important roles.

Not all the sensors under consideration received the same amount of attention from the robotics community. As the infrared rotating scanner has a limited user base, it is not surprising that the literature is very limited [50].

Even if the Neato laser scanner is mass-produced and more widely available, it has received comparatively no attention. Its unavailability as a standalone component is a possible explanation. Up to now, a full robot must be bought to get access to this sensor. The only available characterisation data are from its seminal article [36], which describes the calibration procedure used by the integrated electronics.

The Hokuyo URG-04LX laser scanner is currently the one with the most scientific coverage. Like some other scanners before it (the Sick LSM200, for example), it has been widely adopted by the robotics community and was used in numerous applications. Several scholars characterised this sensor and considered its suitability for a variety of domain. For a good summary of these publications, refer to [55]. If we focus on mapping applications, the characterisations performed in [56, 57] are of interest, but the most complete and up to date one can be found in [58]. The authors investigate the effects of a broad range of parameters on the accuracy.

An informative comparison of the Sick LMS200, the Neato laser scanner, and the Hokuyo URG-04LX is performed by Rockey [59], targeting autonomous wheelchair applications. The authors provide mainly qualitative data, unfortunately, as they are primarily interested in the detection range of various objects. No rigorous comparisons are made regarding the other metrics.

When we considered the corpus of all published work on mapping sensors, we were unable to find publications related to power considerations. Up to now, this topic has been of little interest for the robotics community, as researchers using these sensors were principally focused on the high-level algorithms and applications. Likewise, we were unable to find some metrics relating the performance of the sensor with the energy required to get these data.




Figure 4.1 – Picture of the simple setup used to characterise the mapping sensors. The nominal distance d_n between the target and the sensor is measured with respect to the robot's main axle tree. The target makes a right angle with the robot.

4.2 Methodology

The core of Sec. 4.3 uses the measure of the distance between each sensor and a fixed target under a variety of conditions. These measurements are performed inside our arena (see Sec. 2.3) under controlled conditions. The target is a 50 x 50 cm² square made of Medium Density Fibreboard (MDF) wood. While most authors use a white target to perform characterisation, this is not very representative of a real environment, especially when the sensor operates near the limits of its sensing range. Note that some authors have compared measures while using several types of targets [58].

The measures of Secs. 4.3.1 to 4.3.3 are performed with the setup pictured in Fig. 4.1. The target is placed at a known distance in front of the robot, at a right angle with respect to the centre line. The nominal distance d_n between the target and the sensor is measured with respect to the robot's main axle tree. The nominal distance was independently measured with a laser rangefinder (Bosch DLE40), which has a typical measurement accuracy of ± 1.5 mm.

The measures of Sec. 4.3.4 are performed with the target placed at a known angle with respect to the centre line. The angle is measured with a steel protractor that has a resolution of 1°. The overall angular accuracy with such a setup is not expected to be optimal, probably between 2° (at best) and 5° (at worst).

We wrote a simple acquisition software in C++ for these experiments. The principal interest was to have customised software triggers at our disposal. With such a tool, it is possible to

reduce the amount of useless captured data and consequently reduce the post-processing time.

For the power analysis of Sec. 4.4, a custom four-channel power datalogger (see Appendix A) is used. With this tool, it is possible to record the power consumed by various parts of the system at the same time. The placement of the channels depends on the sensor under consideration, which is further explained in the course of Sec. 4.4.

4.3 Characterisation


The complete characterisation of a range sensor is a broad topic. In this work, we focus on the external factors that have an impact for an all-purpose mapping application. We will first analyse the possible thermal drift that may affect the sensors and explore the downsides of such an effect. The distribution of the noise on the data is also considered and compared with a Gaussian model. Then, the principal analysis focuses on the relationship between the distance to the target and the sensor's output. A calibration model is also researched to improve the accuracy of the sensors. Finally, the sensitivity of the sensor's output with respect to the angle made with the target is inspected.

4.3.1 Thermal Drift

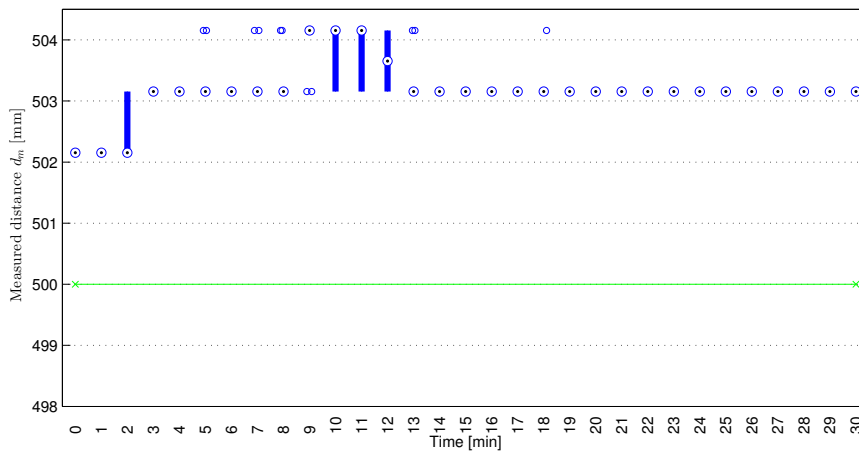
Several authors reported a thermal drift with the Hokuyo scanner, due to the self-heating process [56, 57, 58]. This thermal drift is best seen when cold-starting the sensor. A waiting time of up to 90 minutes has been reported before observing stable data. Such behaviour was also previously reported for the older Sick LMS200 [60], with a waiting time up to three hours. We tested our mapping sensors to see if they are subject to this effect. The test is conducted by placing the target at a fixed distance from the sensor (inside the known working range) and taking measurements at a constant time interval since the sensor's start-up.

Figure 4.2 shows the results for the three sensors. The infrared rotating scanner (Fig. 4.2a) does not exhibit any correlation between the time and the measured distance. The same can be said for the Neato scanner (Fig. 4.2b). The measured distance oscillates between three values during the first 15 minutes, but the amplitude of this oscillation is only 0.2% of the total value. This is interpreted as a small noise on the discretized measures.

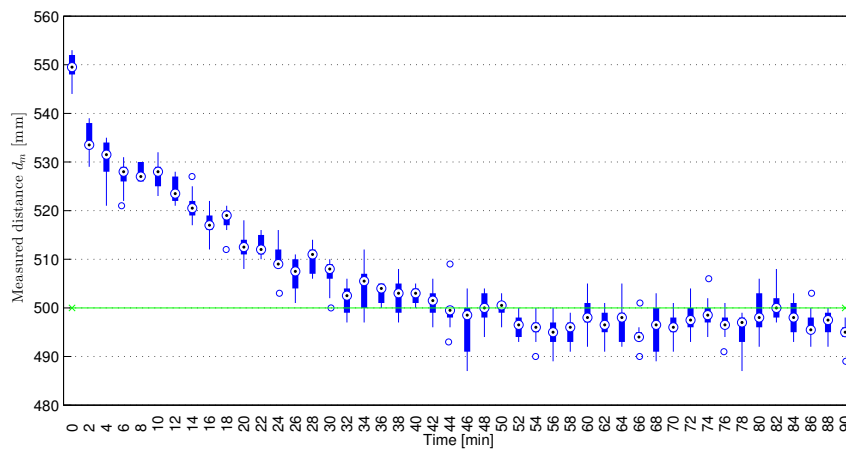
As expected, the Hokuyo scanner has a rather strong dependency between the measured distance and the run time (Fig. 4.2c). The measures stabilise after about 45 minutes, which is faster than reported by some authors [57] but comparable to others [56, 58]. This difference in the repeatability may be due to a variation of the sensor's thermal conductivity, for example, caused by a different mounting frame on the test bench. Thermal images of Fig. 4.3 show the self-heating process. After 45 minutes, the sensor has heated by about 10 degrees. This heat is



(a) Infrared rotating scanner — $d_n = 150$ mm, $\Delta t = 60$ s.



(b) Neato scanner — $d_n = 500$ mm, $\Delta t = 60$ s.



(c) Hokuyo scanner — $d_n = 500$ mm, $\Delta t = 120$ s.

Figure 4.2 – Results for the thermal drift test. The measured distance is plotted as a function of the time. The horizontal axis is not identical between plots ($N = 10$ for each time bin).




Figure 4.3 – Thermal images showing the self-heating process of the Hokuyo scanner.

mainly concentrated in the lower part, where the spindle motor is located [54]. Afterwards, we always waited for at least 60 minutes before starting the measures.


This thermal drift has two major consequences. First, it is a big constraint imposed to the final user, who will have to wait for a long time before getting reliable measures. This also means a bigger sensitivity to changes in the room's temperature. Finally, from an energy point of view, this leads to a waste of energy during the warm-up time. The effect of this loss will be comparatively bigger if the user wants to use the sensor only for a short period of time. Both the infrared rotating scanner and the Neato scanner have a clear advantage here.

4.3.2 Gaussian Noise Hypothesis


All the sensors have noise on the measures. The overwhelming majority of Simultaneous Localisation and Mapping (SLAM) algorithms have explicit or implicit hypotheses regarding the probability distribution of both the robot's motion model and the sensor's observation model. Models with a zero mean uncorrelated Gaussian noise are used in most of the cases, which result in algorithms that are computationally less expensive [30]. Only a few authors have investigated the use of non-Gaussian noise for the motion model [61] or the observation model [62, 63]. It is of interest to examine the noise distribution of the sensors at hand and validate if this noise can be the result of a Gaussian process. Normal probability plots are drawn in Fig. 4.4. On each plot, the dashed line shows the ideal distribution of the measures according to a normal distribution. We repeated the measurement for each sensor with a target placed at two different distances inside the working range. For each measurement, 1000 samples were collected.

The infrared rotating scanner consists of two short-range sensors and two long-range sensors (Sec. 3.2.1). We tested both ranges by using a close target (Fig. 4.4a) and a distant one (Fig. 4.4b).


4.3. Characterisation




(a) Infrared rotating scanner — $d_n = 200$ mm.




(b) Infrared rotating scanner — $d_n = 500$ mm.




(c) Neato scanner — $d_n = 500$ mm.



(d) Neato scanner — $d_n = 2000$ mm.



(e) Hokuyo scanner — $d_n = 500$ mm.



(f) Hokuyo scanner — $d_n = 2000$ mm.

Figure 4.4 – Normal probability plots for the three sensors. For each sensor, two different nominal distances d_n are tested ($N = 1000$ for each distance).

We can observe that this scanner is the one with the highest deviation from a normal distribution, compared to the two others. The noise distribution is not even symmetrical with respect to the median.

The Neato scanner has a low level of noise when dealing with close targets (Fig. 4.4c). It is consequently hard to analyse the distribution of such noise. With more distant targets (Fig. 4.4d), the noise around the central tendency is close to a normal distribution, but values are then underrepresented when moving away. If we draw a qualitative representation, the noise follows a bell-shaped curve, but it is more densely packed around the central tendency compared to a normal distribution.


Finally, when considering the short range, the Hokuyo scanner has a noise distribution that is closest to a Gaussian process (Fig. 4.4e). The tail of its distribution slightly deviates from normality though. When considering the long range (Fig. 4.4f), things are getting worse. The tendency is inverted compared to the Neato scanner, and the curve is more spread out on both sides of the central tendency compared to a normal distribution.

To confirm our visual conclusions, we can apply a statistical normality test. In all the cases, the goodness of fit estimated using the χ^2 test rejects our null hypothesis (the data comes from a normal distribution) with the standard 5% significance level. In conclusion, statistically no sensor has a normal distribution of its noise. Some are far from it, like the infrared rotating scanner. The Hokuyo scanner is the closest one to a Gaussian process, but it also has more noise compared to the Neato.


But at the end, does strict normality really matter? Both the Neato and Hokuyo scanners are close realisations of a normally distributed noise, even if they statistically deviate from this ideal model. In practice, mapping algorithms will be probably more affected if the distribution is not zero mean, which relates to the accuracy of the sensor (see Sec. 4.3.4). Due to the non-normal distribution of the noise, some statistical estimators, such as the standard deviation, will be slightly biased. But all these sensors were already used in the past within SLAM frameworks with some success, as these algorithms can cope with some imperfections on the model's parameters. However, it would be interesting to compare standard approaches against the aforementioned methods proposed to handle non-Gaussian noise [62, 63].

4.3.3 Impact of the Target's Distance


The principal characteristic of any distance sensor is its output with respect to its distance with an object. Several metrics are of interest, such as the accuracy, the linearity, and the noise on the readings. We will first qualitatively investigate the absolute error $\Delta d = d_m - d_n$ on the measured distance for each sensor. The next section will investigate the feasibility of a calibration model and its consequence on the accuracy. We placed the MDF target at various distances with an increasing interval and then took 50 readings for each target to have a representative amount of noise. As we are dealing with rotating scanners, we consider only



(a) Infrared rotating scanner.



(b) Neato scanner.



(c) Hokuyo scanner.

Figure 4.5 – Absolute error Δd as a function of the nominal distance d_n . The horizontal axis uses a logarithmic scale, and it is not identical between plots ($N = 50$ for each distance).

the readings that are on the centre line. A calibration procedure was first carried out before any measures to guarantee the sensor's alignment with the target. The results are shown in Fig. 4.5.

The infrared rotating scanner is a bit exotic, as it embeds two short-range (40 – 300 mm) and two long-range sensors (200 – 1500 mm). We merged together the data from the short-range sensors and likewise for the long-range sensors. This gives two series of data (Fig. 4.5a). The short-range series outputs meaningful data between 100 and 300 mm. The sensors can take closer measures, but we are limited by the volume of the robot. Beyond 300 mm, the error starts to grow linearly with the distance as the sensor saturates. Regarding the long-range series, no reliable data can be taken below the distance of 450 mm, and the error starts to be significant after 1000 mm. If we superimpose the two series, we have a dead zone, as it is almost impossible to get reliable data between 300 and 450 mm.

The Neato scanner has a linear response between 200 and 1000 mm (Fig. 4.5b). We can discern a constant bias, which is possibly due to a discrepancy between the origin of the sensor and the robot's centre. Regardless, this can be easily corrected afterwards with a calibration model. Between 1000 and 4000 mm, we notice an increase of the error but with a rate inferior to the saturation level. A calibration model may be possible, which is investigated in the next section.


The Hokuyo scanner is the sensor with the smallest error on the working range out of the box (Fig. 4.5c). The error below 300 mm is a bit unsteady, with some high values and no clear tendency. Above 300 mm, the error apparently follows a third-order model as previously noticed by [58]. This will be investigated in the next section.

4.3.4 Calibration Model


As noted in the previous section, some more work is necessary to get accurate measures with these sensors. We have to fuse together both data series for the infrared rotating scanner. The Neato scanner has a growing error for ranges beyond 1 metre. And the Hokuyo scanner, while the best among the others, is not yet perfect. We will attempt to devise good calibration models and assess the resulting accuracy.

Infrared Rotating Scanner


The major challenge with the infrared rotating scanner is the fusion of the two ranges into a consistent output. Figure 4.6a shows a close-up on the two series, with the saturated values removed. The dead zone between 300 mm and 450 mm becomes evident. A simple algorithm can be used to fuse both series together: we use the short-range sensor if the long-range sensor gives no readings. Otherwise, we use the long-range reading. The resulting relative error is plotted in Fig. 4.6b. The accuracy and the noise are both not very satisfactory.




(a) Absolute error as a function of the nominal distance d_n ($N = 50$ for each distance).



(b) Relative error on the filtered sensor as a function of the nominal distance d_n ($N = 50$ for each distance).



(c) Schema of the test environment. The sensor is placed at the origin.



(d) Polar plot of the test environment with the filtered data ($d_n = 1000$ mm, $N = 50$).

Figure 4.6 – Calibration steps for the infrared rotating scanner.

To better show the problems, the test environment of Fig. 4.6c is used. The sensor is placed on the central axis of the arena with walls on both sides at 2 metres apart. The wood target is at 1 metre on the front and another wall behind at 33 centimetres. The corresponding scans are plotted in Fig. 4.6d. We can notice the saturation of the sensor at around 1 metre when there is no obstacle. We could get rid of these spurious data, but this would imply an aggressive filtering and shorten the usable range even more.

There are some phantom points around the edges of the target. This behaviour is not symmetric with respect to the target. This asymmetry is probably caused by the direction of rotation (clockwise) and some spurious reflections on the side of the target. Overall, the quality of this scanner is low, and it is difficult to make it usable inside a spacious environment.

Neato Scanner

We noted earlier two distinct error models. Before 1000 mm, the sensor experiences a constant bias. By performing a linear least squares analysis, this bias is estimated to -14 mm. The resulting linear error model is shown in Fig. 4.7a.

Above 1000 mm, a quadratic error model is fitted onto the data, as shown in Fig. 4.7a. The estimated error $\widehat{\Delta d}$, as a function of the measured distance d_m , is

$$\begin{aligned} \widehat{\Delta d} &= \\ &= d_m - \hat{d}_n = 2.35 \cdot 10^{-5} \cdot d_m^2 - 0.0196 \cdot d_m - 21 . \end{aligned} \quad (4.1)$$


As a consequence, the model for the estimated distance \hat{d}_n , as a function of the measured distance d_m , is

$$\hat{d}_n = \begin{cases} d_m + 14, & \text{if } d_m \leq 1000 \\ -2.35 \cdot 10^{-5} \cdot d_m^2 + 1.02 \cdot d_m + 21, & \text{otherwise} \end{cases} . \quad (4.2)$$


The absolute error between the estimated distance \hat{d}_n and the nominal distance d_n is shown in Fig. 4.7b, and the relative error is given in Fig. 4.7c. With our model, we can bound the error to $\pm 1\%$ of the nominal distance on the whole range (200 – 4000 mm). The precision and the noise are especially good below 3000 mm.

Hokuyo Scanner


With the Hokuyo scanner, we noticed a trend on the error for larger ranges. As with the Neato scanner, we can fit a third-order model by using least squares minimisation. We used data above 600 mm, as the error appears less predictive below. The fitted model is shown in Fig. 4.8a.



(a) Absolute error on the raw data as a function of the measured distance and the corresponding error models.




(b) Absolute error $\hat{d}_n - d_n$ on the filtered data as a function of the nominal distance.




(c) Relative error $(\hat{d}_n - d_n) / d_n$ on the filtered data as a function of the nominal distance.

Figure 4.7 – Calibration steps for the Neato scanner ($N = 50$ for each distance).


Chapter 4. Characterisation of Sensors for Energy-Efficient Mapping



(a) Absolute error on the raw data as a function of the nominal distance and the corresponding error model (above 600 mm).



(b) Absolute error $\hat{d}_n - d_n$ on the filtered data as a function of the nominal distance.



(c) Relative error $(\hat{d}_n - d_n) / d_n$ on the filtered data as a function of the nominal distance.

Figure 4.8 – Calibration steps for the Hokuyo scanner ($N = 50$ for each distance).




Figure 4.9 – Polar plot of the data accumulated during 50 scans by the Hokuyo scanner with the target at a nominal distance of 150 mm. We can see how the plane of the target is distorted by the high error occurring between 125 mm and 150 mm.

The estimated error $\widehat{\Delta d}$, as a function of the measured distance d_m , is

$$\begin{aligned} \widehat{\Delta d} &= \\ &= d_m - \hat{d}_n = -3.23 \cdot 10^{-9} \cdot d_m^3 + 2.66 \cdot 10^{-5} \cdot d_m^2 - 0.0563 \cdot d_m + 24 \quad . \end{aligned} \quad (4.3)$$

The first polynomial root is at $d_m = 592$ mm. As a consequence, the model for the estimated distance \hat{d}_n , as a function of the measured distance d_m , is

$$\hat{d}_n = \begin{cases} d_m, & \text{if } d_m \leq 592 \\ 3.23 \cdot 10^{-9} \cdot d_m^3 - 2.66 \cdot 10^{-5} \cdot d_m^2 + 1.06 \cdot d_m - 24, & \text{otherwise} \end{cases} \quad . \quad (4.4)$$

The absolute error between the estimated distance \hat{d}_n and the nominal distance d_n is shown in Fig. 4.8b, and the relative error is given in Fig. 4.8c. The error on the filtered data is better compared to the raw data. The absolute error above 400 mm is bounded to ± 10 mm, and the relative error above 800 mm is better than 1%. The error on the short-range data (below 300 mm) is quite important with the Hokuyo scanner. We have an area between 125 mm and 150 mm where the error is even unexpectedly high (up to 20%). This can be visualised in Fig. 4.9. The plane of the target is strongly distorted. Such behaviour was not observed by other authors, but they used a coarser resolution for their measures at low range [57, 58] or did not performed measures at low range [56]. The cause of this behaviour should be investigated further.

In summary, this scanner has a slightly better range compared to the Neato scanner, but data between 125 mm and 150 mm are unusable. The behaviour of this scanner's error is opposite to the Neato scanner, as the Neato has good accuracy below 2500 mm, but is getting worse above. In addition, data have more noise on most of the range compared to the Neato scanner.

4.3.5 Sensitivity to the Target's Angle


Finally, we analyse the sensitivity of the sensor's output with respect to the incident angle with the target. Ideally, the sensor should be insensitive to the target's angle; otherwise, the perceived distance would be biased. As the angle of the target is usually unknown a priori, it is difficult to correct this effect with a post-processing step. Results are shown in Fig. 4.10 with the calibration models of the previous section applied on all the data.

The infrared rotating scanner has two series of data (Fig. 4.10a). The median of the long-range sensors is little influenced by the angle, but the dispersion of the data significantly increases with the angle. The dispersion of the short-range sensors is less impacted by the incident angle, but the median suffers from a deviation. This deviation may be due to a small misalignment of these sensors with respect to the central line. The higher increase of the dispersion for the long-range series might be due to the higher distance from the target (500 mm instead of 200 mm for the short-range series), resulting in a bigger cross-section of the beam. These sensors are indeed less focused compared to a laser beam. In both cases, the readings are no longer reliable above 70° .


The Neato scanner is also visibly impacted by the incident angle (Fig. 4.10b). As the incident angle increases, the median of the error shifts from -10 mm to +10 mm. Thanks to the low noise of this sensor (at $d_n = 2000$ mm), the absolute error stays bounded between -15 mm and +15 mm. Just like the infrared rotating scanner, the readings are no longer reliable above 70° .

The Hokuyo scanner seems less perturbed by the incident angle (Fig. 4.10c). The median of the error is almost constant between 0° and 60° . But due to the higher amount of noise on the readings, the absolute error is bounded only between -25 mm and +10 mm. This is, in any case, not better than the Neato scanner.


In summary, triangulation sensors (infrared rotating scanner and Neato scanner) are likely more impacted by the target's angle, compared to sensors using a measure of the phase shift like the Hokuyo scanner. But the Neato scanner stays competitive with the Hokuyo scanner when comparing the bounds of the absolute error, thanks to its lower noise.



(a) Infrared rotating scanner — $d_n = 200$ mm (short range) and $d_n = 500$ mm (long range).



(b) Neato scanner — $d_n = 2000$ mm.



(c) Hokuyo scanner — $d_n = 2000$ mm.

Figure 4.10 – Sensitivity to the target’s angle for each sensor ($N = 50$ for each angle).




Figure 4.11 – Instantaneous power (in milliwatts) consumed by the processor and the infrared rotating scanner during the boot sequence. The Low-DropOut (LDO) supply of the sensor is then manually turned off (at $t = 72$ s), before being turned on again (at $t = 82$ s). Finally, a power off sequence is initiated at $t = 92$ s.

4.4 Power Analysis

A detailed analysis of the power consumption is now performed for each sensor. As discussed in Sec. 4.2, we use the four-channel power datalogger described in Appendix A. As an example, Fig. 4.11 shows the instantaneous power consumed by the processor and the infrared rotating scanner during a complete power on/power off cycle.

For our calculations, however, let us first define a few conventions. The datalogger samples the signals with a sampling frequency f_s or, conversely, a sampling period T_s . The measured power is a discrete-time signal $p[k]$. Computing the energy under the discrete curve between time t_1 and t_2 is easily achieved by using the piecewise constant interpolation

$$E = \sum_{k=k_1}^{k_2} p[k] \cdot T_s = T_s \cdot \sum_{k=k_1}^{k_2} p[k] , \quad (4.5)$$

with $k_1 = t_1/T_s$ and $k_2 = t_2/T_s$. It is often more convenient to compute the average power, as it can be used to predict the energy required during a given amount of time. Likewise, the average power between time t_1 and t_2 is computed as

$$P_{avg} = \frac{1}{k_2 - k_1} \sum_{k=k_1}^{k_2} p[k] . \quad (4.6)$$

We also used this average power to perform a filtering of the signals plotted in Figs. 4.12 to 4.14. In this case, the average is performed on a moving window, where the size of the averaging window is defined by $\Delta t_w = t_2 - t_1 = T_s(k_2 - k_1)$.

4.4.1 Infrared Rotating Scanner

The power analysis of the infrared rotating scanner is shown in Fig. 4.12. A few interesting points are worthy of interest. When looking at the overall view (Fig. 4.12a), the life cycle of the sensor can be divided into four consecutive states.

At the beginning, the LDO supply of the sensor is turned off. During this state, the sensor consumes around 20 mW. The supply is then turned on. During a short period of time, the microcontroller is booting and consumes around 350 mW (Fig. 4.12b). The energy required during this transitional state equals 0.23 J. The sensor is then ready to work and consumes 215 mW on average. This idle state requires about ten times more energy compared to the off state.


The most interesting state happens when the sensor is performing a scan. On average, the power needed for this phase is 2.25 W. If we look at the details (Fig. 4.12c), the instantaneous power experiences short pulses at 7 W, followed by moments of low consumption. During the pulses, the energy is transferred to the secondary stage, while in-between data are exchanged using the infrared Data Association (irDA) transceivers (Sec. 3.2.1). This inductive coupling generates a lot of noise on the power supply of the robot, especially if this supply has a high internal resistance.

4.4.2 Neato Scanner


From an overall view, the power analysis for the Neato scanner (Fig. 4.13) is very similar to the one of the infrared rotating scanner. The powered off Neato scanner consumes around 21 mW and 314 mW when in idle state (Fig. 4.13a). During the scanning process, and after the stabilisation of the system, the sensor consumes around 1.07 W. This is about two times less compared to the previous sensor. Likewise, this sensor also needs some time and energy to initialise (Fig. 4.13b).

With some soldering work, we were able to record separately the power used by the motor from the power used by the rest of the logics. A two-second close-up on the scanning phase is shown in Fig. 4.13c. The brushed DC motor is not the principal consumer inside the sensor. Indeed, with 350 mW on average, it consumes just two times less compared to the rest of the electronics (710 mW). The power of the motor strongly oscillates, with a peak-to-peak variation of 55% compared to the average value. A Fast Fourier Transform (FFT) analysis shows a principal frequency at 8.21 Hz, which is exactly two times the measured rotating speed of the scanner (4.17 Hz). This is logical because the power inside a load is twice its supply's frequency¹.


¹ Proof: if the supply voltage is $v(t) = V_p \sin(\omega t)$, then the current inside the load is $i(t) = \frac{v(t)}{Z_L}$. Consequently, the power is $p(t) = v(t) \cdot i(t) = \frac{V_p^2}{Z_L} \sin^2(\omega t) = \frac{V_p^2}{Z_L} \frac{1}{2} (1 - \cos(2\omega t))$.



(a) Overview of the whole sequence: power off, idle, scanning, idle, power off.




(b) Zoom-in on the start-up sequence with the computation of the start-up latency and energy.




(c) Zoom-in of one second during the scanning sequence.


Figure 4.12 – Power analysis of the infrared rotating scanner (averaging window: 0.2 s).



(a) Overview of the whole sequence: power off, idle, scanning, idle, power off.



(b) Zoom-in on the start-up sequence with the computation of the start-up latency and energy.



(c) Zoom-in of two seconds during the scanning sequence.

Figure 4.13 – Power analysis of the Neato scanner (averaging window: 0.2 s).

4.4.3 Hokuyo Scanner

Similarly to the other two sensors, we can measure the power required by the different states of the Hokuyo scanner (Fig. 4.14). The powered off Hokuyo scanner consumes around 11 mW and 206 mW when in idle state (Fig. 4.14a). During the first few seconds of the scanning state, the power is momentarily higher, before reaching an average value around 2.52 W. When looking at the details (Fig. 4.14c), the power is much steadier, with a peak-to-peak variation of only 4% compared to the mean value. Indeed, this device is using a brushless DC motor with regulation electronics.


From an energy point of view, one of the main concern of this scanner is the amount of energy required for the start-up and the warm-up sequences. The initialisation of the electronics does not require a lot of energy (Fig. 4.14b). On the contrary, there is a delay of 22 s between the start of the scanning sequence and the first data, which represents 56 J of lost energy. But worse, we saw in Sec. 4.3.1 that this sensor needs a consequent warm-up time. If we have to wait 45 minutes before getting reliable measures, we need around 6800 J just for doing nothing. If we plan to use this sensor during several consecutive days, this expense might not be too tremendous. Otherwise, this is a serious handicap for any energy-wise robot.

4.5 Concluding Analysis


We have characterised the performance of the three sensors when measuring the distance to a target (Sec. 4.3), including the negative influence of several external factors. We have also performed a power analysis during different activation states (Sec. 4.4). Would it be possible to define a metrics to link both fields together? In such a case, which sensor is the most energy-efficient?

As we saw, our sensors can be set in three different operational modes: powered off, idle, or scanning. We have summarized the power required by each mode in Table 4.1. Each mode consumes roughly ten times more power compared to the preceding one. The weight of each mode inside the total energy budget depends on the exact user scenario, and it cannot be anticipated. The idle mode might be negligible, for instance, or used 99.9% of the time depending on the activation profile. So for the efficiency metrics, let us focus on the scanning mode, where most of the power is consumed.


While scanning, the sensors have a number of interesting characteristics. We already discussed the accuracy and noise. We also tested the maximum range. Another important characteristic is the data bandwidth B_W , which is the number of points measured per second. We have experimentally estimated this bandwidth, and the results are also shown in Table 4.1. Concerning the infrared rotating scanner, the raw bandwidth is actually twice the reported value (four sensors at 60 Hz), but we have to fuse the short-range and long-range sensors together. This effectively divides the practical bandwidth by two.



(a) Overview of the whole sequence: power off, idle, scanning, idle, power off.



(b) Zoom-in on the start-up sequence with the computation of the start-up latency and energy.



(c) Zoom-in of two seconds during the scanning sequence.

Figure 4.14 – Power analysis of the Hokuyo scanner (averaging window: 0.2 s).

Chapter 4. Characterisation of Sensors for Energy-Efficient Mapping

Table 4.1 – Summary table for the power analysis of the three sensors. The upper part of the table reports the average power measured for the three operation modes. The middle part reports data related to the performance of each sensor, such as its data bandwidth. Finally, the lower part presents a number of computed metrics to assess the energy efficiency of each sensor.

		Infrared Scanner	Neato Scanner	Hokuyo Scanner
Average power				
– Powered off	[W]	$23 \cdot 10^{-3}$	$21 \cdot 10^{-3}$	$11 \cdot 10^{-3}$
– Idle	[W]	0.22	0.31	0.21
– Scanning	[W]	2.27	1.07	2.52
Data bandwidth				
	$[\frac{1}{s}]$	120	1500	6820
Range				
– Maximum range	[mm]	1000	4000	5000
– Bits per point	[bit]	10	12	13
Efficiency metrics				
– Energy per point	[J]	$18.9 \cdot 10^{-3}$	$713 \cdot 10^{-6}$	$370 \cdot 10^{-6}$
– Energy per bit	$[\frac{J}{\text{bit}}]$	$1.89 \cdot 10^{-3}$	$59.4 \cdot 10^{-6}$	$28.4 \cdot 10^{-6}$
Energy losses				
– Start-up energy	[J]	$228 \cdot 10^{-3}$	$38 \cdot 10^{-3}$	$211 \cdot 10^{-3}$
– Warm-up energy	[J]	0	0	6800

We can already define the first energy-efficiency metrics. Dividing the scanning power by the data bandwidth informs us about the energy necessary to acquire one data point. This metrics does not take into account the quality of each measure, but it is used to scale the power with the amount of measures produced each second. Results are shown in Table 4.1. Even if the power consumed by the three sensors is inside the same order of magnitude, we have a big disparity on the data bandwidth. At the end, the energy needed to acquire a measure is highly uneven. The infrared rotating scanner is the worst of all, as it requires 26 times more energy for a single point compared to the Neato scanner. And the Neato scanner requires itself two times more energy compared to the Hokuyo scanner. Even if the Neato scanner is the one with the smallest power requirement, it only comes second due to its lower bandwidth.

This first metrics does not take into account the intrinsic nature of each data point, but only their number per second. For the second metrics, we can easily integrate the range of each sensor inside the calculation. Taking into account that all the sensors provide a measure x in millimetres, we can estimate the number of bits necessary to encode the information with $\lceil \log(\max(x)) / \log(2) \rceil$ if all the measures inside the range are equiprobable and uniformly distributed. If we divide the first metrics by this number, we get the necessary energy per bit of information (see Table 4.1). In this case, the difference previously noticed with the first metrics is even bigger, due to the small range of the infrared rotating scanner and the longest

range hold by the Hokuyo scanner. Now there is a difference of more than 60 between both ends of the scale.

This second metrics is still not perfect. It does not take into account a number of important pieces of information, such as the noise or the accuracy of each sensor. But integrating such information inside the metrics is harder, partly due to the correlation between these quantities and the nominal distance of the object. To alleviate some of these problems, we could use the entropy. With such a measure, we can take into account a more complex probability distribution of the data points inside the range of the sensor. But this is still a very synthetic metrics, as it will not be possible to take into account all of the possible effects on the sensor.

A good solution is to define the metrics inside a real, well-defined environment. By placing the sensor at a known fixed location, we can build a static map of the environment as seen by the sensor. If we can compute the disparity between this map and the ground truth, for example, by using the Hausdorff distance [64, 65] or the Mahalanobis distance [66, 67], we can get an overall assessment of the sensor's quality. This figure can then be used to define a better energy-efficiency metrics. Some work is of course necessary to design an informative environment around the sensor. This was not carried out in the present work because we judged the previous energy-efficiency metrics to be satisfactory enough as a first approximation.

A last question remains, however. We have not taken into account the transition energy between the operational modes. As we saw in Sec. 4.4, some energy is lost when powering on the sensors due to the start-up latency. In addition, we saw in Sec. 4.3.1 that the Hokuyo scanner has a substantial thermal drift, which forces the user to wait for at least 45 minutes before getting stable data. To account for these effects, the start-up and warm-up energies are computed and shown in Table 4.1. The start-up energy is probably insignificant for most users. On the contrary, the warm-up energy required by the Hokuyo scanner can be a major problem for most service robotic applications.

4.6 Discussion

The three sensors considered in this chapter are all very different, based on the working principle, the level of performance, and the power necessary during the entire life cycle. To this end, the choice of a particular sensor depends on the application. Some users might focus solely on the performance. In such a case, if a low level of noise is required, the Neato scanner is a good candidate. If the accuracy above 3 m or a high bandwidth are important criteria, the Hokuyo scanner should be used.

But for a mobile robot, the energy required for the operation of the sensor is also an important piece of data. It is not easy to give a one-size-fits-all metrics, as several considerations should be taken into account inside the global energy balance. This heavily depends on the usage profile over time. In our case, we are planning to use the sensor for maybe one hour per day and then shut it down for the rest of the time. The warm-up cost of the Hokuyo scanner, both

Chapter 4. Characterisation of Sensors for Energy-Efficient Mapping

in time and energy, becomes prohibitive under such circumstances. This is why we have decided to use the Neato scanner during the final experiments of Chap. 6.

5 Theoretical Framework for an Indoor Harvesting Robot

In Chap. 2, we introduced the importance of the energy for indoor service robots. Energy is a scarce resource, and a mobile robot has to store it inside a battery. A frequent recharge is needed due to the limited capacity of current technologies. Chapters 3 and 4 presented the features of the marXbot robot and summarized the strategies used to spare energy, while providing a highly modular robotic platform. The present chapter will dig into the topic of energy harvesting, and we will try to find a suitable energy source for indoor service robots.

Our overall goal is to find innovative energy supplies in indoor environments with enough power to fulfil the needs of common mobile robots. We saw examples of such robots in Chap. 2, but our work is not limited to robotic vacuum cleaners. For example, we can think about applications such as floor scrubbing, patrolling, or remote presence robots. Robotic vacuum cleaners are nice examples, because they set a high bound on the required energy. Our study is split in two halves. The present chapter will explore the available energy sources, set up a reasonable scenario, and provide a framework to compute the theoretical upper bound on the available energy. The next chapter will use a physical robot to evaluate the feasibility and get a practical upper bound.

We have divided the work into several steps. A survey of the possible sources of indoor energy is conducted in Sec. 5.1, with a focus on the renewable energies and on the achievable performances regarding our application¹. The envisioned scenario is then detailed in Sec. 5.2, including a look at the related work from other scholars and the definition of the hypotheses used throughout this work. Our theoretical framework is developed in Sec. 5.3 and Sec. 5.4, and a concluding analysis is performed in Sec. 5.5.

¹ Parts of Secs. 5.1 and 5.2 were previously published in [68].

5.1 Indoor Energy Harvesting for Service Robots

Wasted energy is environmentally ubiquitous and represents a potentially cheap source of power. We will present a short survey of available power technologies with an emphasis on renewable energies in indoor environments. Comprehensive books and state-of-the-art reviews already exist [69, 70, 71, 72] but usually focus on energies that are unavailable in our indoor setting.

In previous works, authors have often considered the use of scavenged energy in order to move robots around, but at present, only a few people have actually managed to do so, especially indoors. This is mostly due to either the small power density of the source or the difficulty in predicting the spatial and temporal distribution in a reliable way. For example, mobile harvesting is addressed in theory with outdoor unmanned vehicles by [73]. They conclude with the possibility of using solar and kinetic flow energies to power devices in the range of 1 ~ 10 W. Solar power is also commonly used in outer space exploration robots with some success [74]. The EcoBot-II is quite original, in the sense that it embeds a microbial fuel cell converting unrefined insect biomass into electricity [75]. However, this gives just enough energy to allow it to travel a few centimetres.

Most of the service robots will have to move around for their mission, implying higher energetic needs. But self-mobility is not always mandatory. By exploiting the interactions already in place between the humans and those specific objects, the humans can be used as a transportation vector, thus lowering the requirements. This concept is called parasitic mobility [76]. For example, the robotic glasses designed by [77] are used like normal glasses most of the time, except when they are empty and autonomously go back to get refilled. But unfortunately, this concept is hard to apply to most kinds of service robots, as the human is expecting it to perform its task autonomously, without any human support in the loop.

5.1.1 Heat

Heat is a universal source of energy around us, which is part of the wasted power generated by thermal engines and exhausts, by poorly insulated buildings, and even by the human body. Thermal energy scavenging devices are primarily ruled by the laws of thermodynamics. Their efficiency η is inherently limited by the *Carnot cycle*

$$\eta_C = \frac{T_H - T_C}{T_H} = \frac{\Delta T}{T_H}, \quad (5.1)$$

where T_H and T_C are the temperatures of the hot and cold sides, respectively [78]. The temperature gradient ΔT should be maximized, favouring applications such as in the automotive industry [79] or power dissipation in electronic devices [80]. Making use of thermal gradients is, however, not well suited for obtaining energy under ambient conditions. For example, when applied to scavenging energy between the human body ($T_H=309$ K) and the ambient air ($T_L=293$ K), the maximum achievable efficiency is a mere 5.2%.

5.1. Indoor Energy Harvesting for Service Robots

The total efficiency will moreover be lowered by the non-ideal nature of the transducer. The conversion of the heat into electricity is the most widely used solid-state technology at low temperature. As presented by Min in [69, Chap. 5], Thermoelectric Generators (TEGs) are mainly based on the *Seebeck effect*. A temperature gradient ΔT applied to a junction will generate a voltage V_{ab} , following

$$V_{ab} = \alpha_{ab} \Delta T , \quad (5.2)$$

where α_{ab} is the Seebeck coefficient of the considered junction. Optimizing a TEG is first a matter of material engineering. For ambient temperatures, with Bi_2Te_3 used in the case of the human body, the expected overall efficiency falls around 0.91%.

Seiko was the first company to release a watch powered by human body heat [81]. Using this strategy, it is also feasible to power wearable sensor nodes, providing about $100 \mu\text{W}$ when placed on the wrist [82]. A TEG half-buried in the soil can provide up-to 1 W/m^2 (peak) to an outdoor sensor node, by exploiting the temperature difference between the buried side and the side exposed to the ambient air [78]. Low-cost materials are also foreseen [83], but no major progress regarding efficiency is expected.

5.1.2 Light

The sun is a huge source of energy, which can be easily exploited by solar cells, at least during daylight. Indoor, room lighting is another ubiquitous source of energy². The first calculators to rely solely on this source were already designed in the 1980s [84]. There are some discrepancies in the literature regarding the peak power levels available from the visible electromagnetic spectrum. The most widely accepted values are presented in Table 5.1 based on [69, 72, 85].

The efficiency of photovoltaic cells is rather low, compared to solar heating systems. Most industrial cells will have at most 20% efficiency, and this figure heavily depends on a number of factors. For example, one should take care to note the angle between the panel and the incoming beam, the current illumination level [86], or the spectral sensitivity of the cell with respect to the light source. The optimal technology thus depends on the considered light source. Single crystal silicon is usually used for outdoor conditions, with an efficiency between 15 and 20% [87]. Thin film amorphous silicon or cadmium telluride cells are proposed for indoor conditions [72]. In this case, the efficiency often drop down to a mere 10% due to the low-lighting conditions [88]. The expected power densities at the output of the cells is computed in Table 5.1 based on these values.

Other technologies exist, such as the dye sensitized cells [89] and the multi-gap cells [90]. This latter technology is promising, but currently very costly because efficiency is increased by the use of multiple band gaps inside the semiconductor material. This allows the capture of

²Strictly speaking, this is not a renewable source of energy like the sun, as it depends on human-supplied electricity. But this is a universal source of wasted energy that one can take advantage of.

Chapter 5. Theoretical Framework for an Indoor Harvesting Robot

Table 5.1 – Summary of the commonly admitted power densities in photovoltaic applications, based on the type of light source and the atmospheric conditions. Based on the simplified efficiency of most photovoltaic cells, the conceivable power production is computed. One has to keep in mind that this view remains simplistic. For example, the solar power depends on a number of parameters, such as the latitude, time of the year, atmospheric pollution level, and so on. The indoor lighting also has a wide range of specific light types, each one with its unique spectral density.

Light Source	Conditions	Raw Power $\frac{W}{m^2}$	Cell Efficiency	Scavenged Power $\frac{W}{m^2}$
Sun	Clear Sky	500 – 1000	20%	100 – 200
Sun	Cloudy	100	15%	15
Artificial	Indoors	1 – 5	10%	0.1 – 0.5

photons from a wider range of wavelengths, consequently improving the harvesting of the entire solar spectrum. For example, most of the spatial exploration robots use the solar power by embedding multiple band gaps panels³.

5.1.3 Ambient Radio Frequencies

Nowadays, wireless transmissions and electromagnetic waves belonging to the Radio Frequency (RF) band are everywhere: television, radio, and cell phones among others. It would be tempting to exploit part of this radiated power.

The *Friis equation* can be used to compute the received power P_r in an antenna with gain G_r , assuming a wave of wavelength λ emitted at a distance d by an isotropic source of gain G_t and propagating in free space, far enough from the emitter (far field condition)

$$P_r = G_r G_t P_0 \left(\frac{\lambda}{4\pi d} \right)^2 = G_r \frac{(E\lambda)^2}{4\pi Z_0}, \quad (5.3)$$

where the radiation impedance of free space, Z_0 , equals 377 ohms [91]. P_0 is the emitted power, while E is the field strength. When considering indoor conditions, the received power is closer to $P_r \sim d^{-4}$. Assuming a unity gain and a field strength of 1 V/m at 2.4 GHz, the scavenged power is around 3.3 μ W.

While the scavenged power is far too low for our use case, this is enough for wireless sensor nodes [92] or Radio-Frequency Identification (RFID) tags, even at a distance of a few metres [93].

³Radioisotope Thermoelectric Generator (RTG) is another common power source, used, for example, on the Viking 1, Viking 2, and Mars Science Laboratory missions. The principle is based on a thermoelectric generator powered with the heat produced by the decay of a radioactive element.

5.1.4 Pressure and Temperature Variations

The daily atmospheric changes could also power devices. This technology has been in use since the early 1900s, providing power to clocks [94], and more recently to watches [95], based on a closed dilating volume. However, apart from the watch industry, no other device has been powered by this kind of energy.

In fact, according to [96], a variation of 1 K or 400 Pa provides $3.5 \cdot 10^{-3}$ J of energy. This would be enough to power a small device with a consumption of $50 \mu\text{W}$ for only 70 seconds. This is, however, enough to power the “Atmos” clock for 2 days [97], which implies that this device has an impressive power consumption of only 20 nW.

5.1.5 Gas and Liquid Flows

The fluid mechanics has long been exploited to scavenge energy in windmills and watermills. This energy is also available inside houses, hidden in small air breezes generated by air-conditioning or heating systems, as well as inside water and sewage pipes. This principle is, for example, exploited by the self-powered shower handle [98], but at the expense of a reduced pressure at the output.

According to the *Betz law*, for a flow of speed v considered over a surface A , only 59.3% of the total power can be scavenged

$$P_{\max} = C_p \cdot P_{\text{flow}} = C_p \cdot \frac{1}{2} \rho A v^3 \quad \text{where } C_p = \frac{16}{27} \approx 0.59 \quad . \quad (5.4)$$

The power is proportional to the flow density ρ . Air and water have respective densities of 1.2 and 1000 kg/m^3 . This makes a theoretical power of 0.35 and 290 W/m^2 , respectively, when considering a cross-section perpendicular to the current and for a flow speed of 1 m/s. Water is undoubtedly more powerful.

Even if big mills tend to have an efficiency close to the Betz limit, small devices operating at low speeds suffer from the increased mechanical and viscous frictions, reducing their efficiency to about 1/6 of the Betz limit [99]. To scavenge energy from the air flow, most researchers are using a custom rotor coupled with a standard DC motor operating in generator mode [100, 101]. This can be enough to power sensor nodes [102, 103], but the available power remains low under normal living conditions.

5.1.6 Mechanical

Mechanical scavenging is a broad topic, with numerous applications [104, 105]. However, to the best of our knowledge, no real robot has been powered solely based on the mechanical

work. Low-power sensor nodes are better suited, for example, to monitor the stress of life-critical materials [106].

Several technologies are currently being studied [70]. Electromagnetic generators are simple and widely used in the industry [107, 108]. Piezoelectric scavengers convert a stress into a voltage, often using a lead zirconium titanate (PZT) ceramic [109, 110, 111]. Finally, electrostatic devices are based on a variable capacitor, operated either at fixed voltage or fixed charge [112, 113, 114].

In the home environment, possible mechanical energy sources include machine vibrations such as clothes dryers and microwave ovens, as well as low-amplitude vibrations such as structural vibrations of the windows and walls [115]. The available power is low, however, when compared to some industrial applications (engines, machine tools), something below 1 mW/cm^3 .

For some kind of service robots, scavenging human power could be an interesting solution. Scavenging the power of the human gait could produce up to a few watts [116], but scavenging too much energy would result in a disturbance for the user. The first research focused on shoes, producing up to 230 mW [117]. Piezoelectric scavengers implemented in floors and pavements are another option. A prototype [118] is said to produce 3.8 W with a walking person.

5.2 Scenario

Now we have to determine the best scavenging technology for our indoor service robot, in light of the scavenging options discussed in the previous section. Table 5.2 summarizes, in its middle column, the theoretic upper bound power density that can be estimated for each technology. This calculation takes into account the expected operating conditions, as well as ideal energy transducers. The column on the right of Table 5.2 shows the achievable power densities, given the current state of the art.

As an illustration, let us assume a 30 W service robot operating for 2 hours per week, resulting in a weekly energy need of 60 Wh. Table 5.3 shows the computed feasibility boundary, based on the figures of Table 5.2 and the estimated charging time slot per day. Several technologies, such as the energy scavenged from the variations in air temperature, were already discarded beforehand due to the extremely small attainable power density.

It can be noted how water and solar are powerful energy sources. Watermills have been used for more than two millenniums⁴ to harvest the power of the rivers. Scavenging a water flow could be really interesting, but a free flow is not a common phenomenon inside a house. The sewage system would be a good candidate, but a service robot could be able to access this

⁴The earliest known written reference to the engineering of watermills dates back at least to the 3rd century BC [119].

Table 5.2 – Summary of the considered scavenging technologies for indoor service robots. The first column shows the theoretic upper bound for the power density, under the given conditions. The second column gives the achievable power density, based on the current state of the art. Base units are mW and cm for an easier comparison at the robotic scale.

Source	Upper Bound Power			Practical Power		
Solar Light	100	$\frac{\text{mW}}{\text{cm}^2}$	Outdoors; Clear sky	20	$\frac{\text{mW}}{\text{cm}^2}$	[87]
Artificial Light	0.1	$\frac{\text{mW}}{\text{cm}^2}$	Indoors	0.01	$\frac{\text{mW}}{\text{cm}^2}$	[88]
Radio Frequencies	0.003	mW	$f = 2.4 \text{ GHz}; E = 1 \text{ V m}^{-1}; G_r = 1$			
Acoustic	$0.96 \cdot 10^{-3}$	$\frac{\text{mW}}{\text{cm}^3}$	100dB			
Thermoelectric	0.1	$\frac{\text{mW}}{\text{cm}^2}$	On the body	0.02	$\frac{\text{mW}}{\text{cm}^2}$	[82] ^a
Pressure Variation	$7.8 \cdot 10^{-6}$	$\frac{\text{mW}}{\text{cm}^3}$	[72] $\Delta p = 677 \text{ Pa}$			
Temperature Variation	0.017	$\frac{\text{mW}}{\text{cm}^3}$	[72] $\Delta T = 10 \text{ K}$	$20 \cdot 10^{-6}$	mW	[96]
Shoe Impact	8400	mW	[116]	230	mW	[117]
Vibrations	0.3	$\frac{\text{mW}}{\text{cm}^3}$	[72]	0.10	$\frac{\text{mW}}{\text{cm}^3}$	[105]
Air Flow	0.58	$\frac{\text{mW}}{\text{cm}^2}$	2.54 m/s at Betz limit	0.10	$\frac{\text{mW}}{\text{cm}^2}$	[100] ^b
Water Flow	480	$\frac{\text{mW}}{\text{cm}^2}$	2.54 m/s at Betz limit	80.7	$\frac{\text{mW}}{\text{cm}^2}$	Hyp. ^c


^aOn the wrist

^b2.54 m/s


^c2.54 m/s; 1/6 of Betz limit.

Table 5.3 – Calculation of the feasibility boundary for several scavenging technologies, given a fictive service robot consuming 60 Wh per week. Power densities are taken from Table 5.2. The charge duration is estimated on a daily basis under optimum conditions. The estimated boundary shows the minimum size for the considered transducer in order to reach the weekly energy goal.

Source	Power Density		Charge Duration per Day	Scavenged Energy per Week		Feasibility Boundary
Solar Light	20	$\frac{\text{mW}}{\text{cm}^2}$	4 h	560	$\frac{\text{mWh}}{\text{cm}^2}$	1 dm ²
Artificial Light	0.010	$\frac{\text{mW}}{\text{cm}^2}$	4 h	0.28	$\frac{\text{mWh}}{\text{cm}^2}$	> 20 m ²
Thermoelectric	0.020	$\frac{\text{mW}}{\text{cm}^2}$	24 h	3.4	$\frac{\text{mWh}}{\text{cm}^2}$	2 m ²
Shoe Impact	230	mW	2 h	3.2	Wh	Several humans
Vibrations	0.10	$\frac{\text{mW}}{\text{cm}^3}$	24 h	17	$\frac{\text{mWh}}{\text{cm}^3}$	0.35 m ³
Air Flow	0.10	$\frac{\text{mW}}{\text{cm}^2}$	24 h	17	$\frac{\text{mWh}}{\text{cm}^2}$	35 m ²
Water Flow	80	$\frac{\text{mW}}{\text{cm}^2}$	1 h	560	$\frac{\text{mWh}}{\text{cm}^2}$	1 dm ²



(a) Mock-up view of a bedroom, with a robot waiting for the child to come back from school. The window could maybe offer enough daylight to recharge its batteries?



(b) The solar harvesting robot, searching for the best location. The intensity of the sun is not uniform, due to the aperture in the wall, and its distribution varies along the day.

Figure 5.1 – Schematic overview of the proposed scenario.

source of energy only with great difficulties. This would also increase the risk of flooding due to the improper flow management.

Scavenging shoe impacts is interesting only in crowded locations, such as what was already done on dance floors [120]. Other gathering places could be equally interesting, such as fitness clubs or shopping malls. But no lightweight solution exists up to now, like a magic piezoelectric carpet, and this requires an entire infrastructure to be installed into the floor of the buildings. In such a case, photovoltaic cells on the rooftop could be a better and cheaper solution.

Let us finally come to the various lighting sources. Indoor lighting is common but would require a huge photovoltaic panel to reach the weekly goal of 60 Wh. This leaves us with solar radiations. It is true that the buildings protect people from most of the direct sunlight to maintain comfortable living conditions. But sunlight can enter through apertures, such as windows and dormers. Today, architects even tend to maximize the natural lighting in our living spaces. This represents both a huge amount of free energy lurking on the floor of the rooms. and also a considerable challenge.

Figure 5.1a shows a simplified view of a small bedroom, where an entertaining robotic box is waiting for the children to come back from school. This application is a good candidate for our scenario because this robot is expected to operate only a few hours per week. For the rest of the time, it can wander around in search of a source of energy to recharge its batteries. And most of the bedrooms have at least one window. If this window is oriented toward the south, odds are that we can perform a significant part of the recharge process based on the incoming sun during the day. But as depicted in Fig. 5.1b, the bright spot's location on the ground varies along the day and across the year. The robot will probably need to move in order to stay inside the beam, but moving will cost some energy. A trade-off will have to be found.

5.2.1 Related Work

No previous state of the art exists regarding the objective of scavenging the solar energy coming inside a room through apertures made into the walls. There is, however, a good track of sun-powered robots and related developments. We will first cover the few documented examples of indoor photovoltaic robots, before covering the relevant state of the art for solar robotics, with a focus on mobile robots.

Indoor Photovoltaic Robots

The robot of [121] seems to be the first attempt to power a micro-robot using a photovoltaic cell (Fig. 5.2a). It is made of a miniature legged inchworm, with an integrated solar panel and electronics, enclosed in an area of $3.6 \times 1.8 \text{ mm}^2$. One prototype was able to move 3 mm in 30 minutes, under the light coming from a fibre optic.

Other researchers have fitted the miniature robot Alice with solar panels [122] (Fig. 5.2b). They used a 3000 ANSI lumens beamer as a power source, and achieved power densities of 7.1 and $27.8 \mu\text{W}/\text{mm}^2$ with crystalline silicon and thin film amorphous silicon photovoltaic cells, respectively. The achieved power balance is unknown.

The work of [126] has achieved to increase the autonomy of the Boe-Bot robot by 15%, under a light bulb of 20 W. They used four crystalline silicon cells, for a total area of 4.8 cm^2 . No more details are known.

Another recent trend is the use of “solar engines”, mainly among the hobbyist community [127, 128, 129]. The principle is very simple: a capacitor is used to store the energy from the photovoltaic panel and restores it when the voltage reaches a certain threshold. Albeit very cheap, this system can only drive a motor by burst, when exposed directly to the sun. It is used in solar-enabled toys.

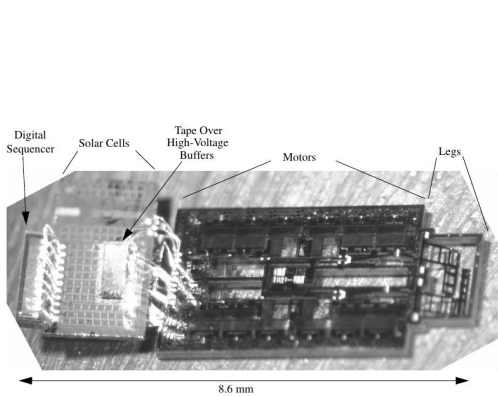
Outdoor Robots

We can list several reasons that motivate the installation of a solar cell onto an outdoor robot: increase the longevity, access remote places far from the power network (deserts, other planets), or maintain operations even during power outages.


Regarding exploration robotics, the “Cool Robot” is one of the most achieved designs [123, 130, 131] (Fig. 5.2c). It is a 160 W rover aimed at autonomously exploring the Antarctic. The robot has been carefully designed, taking into account the light reflected on the snow, which provides up to 40% of the total power. Although several tests were done in real conditions, no tests were performed on the long run, as initially planned.

The study of [125] tries to assess the theoretic feasibility of powering a robot (Tribot) using solar light in different configurations (Fig. 5.2e). With conditions of 4 to 5 kWh per day, it


Chapter 5. Theoretical Framework for an Indoor Harvesting Robot



(a) Solar-powered, milligram robot. Picture credits: S. Hollar [121]. All rights reserved.




(b) The solar "Alice" microrobot. Picture credits: A. Boletis, et al. [122] ©2006 IEEE.




(c) The Cool Robot. Picture credits: J.H. Lever, et al. [123] ©2007 Elsevier B.V.



(d) Indoor surveillance security robot using a solar panel. Picture credits: H.T. Lee, et al. [124].



(e) Different solar designs to supply a mobile robot. Picture credits: T. Marco, et al. [125] ©2010 IEEE.



(f) The Sojourner Mars rover. Picture credits: NASA/JPL.

Figure 5.2 – Pictures of some state-of-the-art photovoltaic robots. This includes indoor robots (a)–(b), outdoor robots (c)–(e), and space robots (f).

concludes to the possibility of powering only the electronics, not the mechatronics. This robot is, however, not designed to be energy-efficient, with a heavy locomotion and an electronics drawing 2 W in idle state.

Regarding the possible applications, Lee [124] is aiming at building a patrolling robot (Fig. 5.2d). Other researchers are investigating the use of solar energy for a power line monitoring robot [132]. In the case of both preliminary works, no details are given related to the solar harvester, nor the achieved performances. On the market, one can find a solar-enabled robotic lawnmower [133]. It is not designed to be powered solely using the sun, however, as its main power source is a recharge outlet.

The integration of a solar-enabled mobile node has also been studied for wireless sensor networks. The Robomote of [134] has a solar panel producing up to 54 mW of useful power, which is not enough to sustain the platform but does enhance its lifespan. Rahimi [135] studies the feasibility of cooperation between mobile and still nodes to power the whole network. Using a small test bed of static nodes and Robomotes, it concludes the feasibility of using 40% of mobile nodes, but no details are given on the experimental conditions.

Up to now, only a few people have investigated the sun tracking problem on a mobile robot. The first work seems to be [136], where a tracked robot is fitted with a tilted solar panel and four small photodiodes at each corner. The primary focus was the survival strategy, using a set of neuronal networks to switch between several behaviours. The robot has randomly navigated for 33 hours with success, putting itself in idle state when necessary. The robot of [137] uses two light-dependent resistors to track the sun following the horizontal plane. Results are inconclusive, however.

Outdoor Space Robots

Two of the three NASA robotic missions targeting Mars have used solar panels to power the rovers. The design of space robots tends to be very conservative, using well-known components to guarantee safe operation. The cost is not the main concern, however.

The Sojourner rover [138] (Pathfinder mission) was fitted with 0.22 m² of GaAs/Ge solar panels, providing up to 16 W of peak power at noon (Fig. 5.2f). The normal driving consumption was about 10 W. Solar panels were designed to be the primary source of energy, as the primary backup battery was used only to provide peak power. Regarding the power management, the available power is roughly estimated using a temperature sensor, a short-circuited cell, and an open cell. Actions are planned using a priori information. It has successfully operated for three months.

Spirit and Opportunity robots (Mars Exploration Rover mission) use triple-junction cells, arranged into one fixed and five deployable panels [139]. Rechargeable lithium batteries are used to absorb the excess energy [140]. The predicted output power was 480 W, but the output is more around 190 W (150 W/m²) when operating under good conditions [74]. The robots

harvested up to 900 Wh per Martian day during the first days, a value which decreased down to 600 Wh due to the dust covering the panels. The robots achieved an unexpected lifetime of 6 and 10 years respectively, and the Opportunity rover is still operating as of 2014.

5.2.2 Research Hypotheses

We will first set out a number of hypotheses that will be used throughout this work. This is necessary because the research field is extensive and needs to be restricted first in order to reach concluding results. It would be desirable to relax most of these constraints in a future work.


South-oriented wall We will begin our considerations by using the ideal case of a south-oriented wall without any obstacles on the outside (trees, etc.) that may cast a shadow on our aperture and lower the incoming light.

Fixed latitude The latitude for most of the simulations, and later the experiments, is fixed to $\Phi = 46.5^\circ$, which corresponds to the latitude of the city of Lausanne in Switzerland.

Direct radiations Only direct radiations are considered, and thus diffuse illumination is neglected. The validity of this hypothesis varies a lot depending on the room's material, but even with white walls, the amount of reflected energy will be at least an order of magnitude lower compared to direct beams. Inside the experimental arena, the measured diffuse illumination when both the artificial sun and the artificial lighting were switched on never exceeded 10 W/m^2 . The same observation was repeated inside several real rooms experiencing sunny conditions and reached the same conclusion.

Standard Test Conditions We are trying to stick to the Standard Test Conditions (STC), as defined by IEC 61853-1, which are used to test the performances of solar panels worldwide. The standardized input irradiance is fixed at $G_n = 1000 \text{ W/m}^2$ with a light spectrum as defined by the American Society for Testing and Materials (ASTM) G173 standard [141]. This simplification has an important consequence. The variation of the incoming irradiance, due to a longer or shorter path through the upper layers of the atmosphere, is not taken into account. In place, the standardized irradiance is used, which means that, in reality, the average incoming irradiance may be higher during the summer and lower during the winter. Several models exist that predict the transmission ratio of the atmosphere as a function of the incident angle. For a comparison example, see [142].

Clear atmosphere For this exploratory study, we are interested in the upper bound that one can get with such a scenario. As a consequence, we are not introducing any atmospheric disturbances, such as clouds or pollution. In reality, this would reduce the peak power and thus the energy along the day. Knowing that a cloudy day will get only a fraction of the usual illumination, this can be a serious limitation in reality. We will relax this hypothesis in Sec. 6.5 by using statistics for the yearly sunshine duration.



(a) Apparent motion of the sun with respect to a fixed observer on Earth⁶. The position of the solar plane depends on the day of year.

(b) Reference frame used to describe the motion of the sun (S) with respect to an observer placed at the origin. The position of the sun is given by the altitude α and the azimuth γ .

Figure 5.3 – Schematic views of the apparent solar motion.

Ideal window For this work, the window's material is set to be ideal. Thus, the light transmittance factor, as well as the refraction index, are neglected. In reality, only a part of the incoming radiations will be transmitted, the rest being reflected or dissipated by the glass material. The behaviour will depend on the rays' incident angle and even on the wavelength of the incoming light. Observed light transmittance ratios for standard double glazing vary from about 0.5 to 0.8. Like the atmospheric conditions, this can greatly impact the efficiency of the overall system.

Mobility on the ground We restrict the movement to the ground plane ($z = 0$), whereas the aperture's centre is located at $z = H_w$ above the ground. This considerably simplifies the simulations, without loss of generality.

5.3 Solar Mechanics

The trajectory of the sun in our sky has been studied since ancient times and used, for example, in sundials. In our scenario, the shadow of the sundial's needle is substituted by the bright spot of the window on the ground. Let us first formulate the basic equations that describe the movement of the sun in the sky with respect to an observer placed at a location with latitude Φ on the Earth⁵. The reference frame is established as shown in Fig. 5.3b.

⁵Most of the equations in this section are from [143, Appendix B], excepted Eqs. (5.8) and (5.9).

⁶Adapted from <http://physics.weber.edu/schroeder/ua/default.html>. Original author: Daniel V. Schroeder. Licensed under a Creative Commons Attribution-NonCommercial 3.0 Unported License (CC BY-NC 3.0).

The apparent motion of the sun in the sky varies with the season. It travels high in the sky during the summer (for the northern hemisphere) and stays low during the winter, as shown in Fig. 5.3a. This seasonal variation is called declination and equals

$$\delta = \epsilon \cdot \sin\left(\frac{360^\circ}{365}(d - 81)\right) \approx 23.45^\circ \sin\left(\frac{360^\circ}{365}(d - 81)\right) , \quad (5.5)$$

where d is the considered day of year. If we choose the day such as $\{k \in \mathbb{N} : d = 81 + k \cdot \frac{365}{2}\}$, an equinox will happen (the length of the day will equal the length of the night). The summer solstice (longest day) happens for $d = 81 + \frac{1}{4}365$, whereas the winter solstice (shortest day) is for $d = 81 + \frac{3}{4}365$. ϵ is the axial tilt (or obliquity) of the Earth's rotation axis and is at the origin of the seasonal variations. Its value is approximately 23.45° , but the obliquity is subject to slow variations over the centuries⁷, as well as small periodic oscillations (called nutation) caused by the tidal forces from the sun and the moon principally. Considering the precision required by this work, these variations can be safely neglected.

The position of the sun on its circular trajectory, as a function of the current time, is called the hour angle ω . By convention, $\omega = 0^\circ$ when the sun is at the solar noon, which corresponds to its highest point during the day. The hour angle equals

$$\omega = 15^\circ(h - 12) , \quad (5.6)$$

with h being the current solar hour. The relation between the solar hour and the wall-clock time is complex, as it depends on the current time zone, possible use of daylight savings time, or the longitude. For simplification and generality reasons, we will stick to solar hours.

With δ and ω in hand, it is now possible to deduce the apparent position of the sun as seen from the observer. First, the solar altitude α is the angular height of the sun, relative to horizontal plane

$$\alpha = \arcsin(\sin(\delta) \sin(\Phi) + \cos(\delta) \cos(\Phi) \cos(\omega)) . \quad (5.7)$$

An important special case occurs at solar noon, when the sun reaches its higher altitude. In this case, the solar altitude at noon (hereafter noted α_n) equals

$$\alpha_n = \arcsin(\sin(\delta) \sin(\Phi) + \cos(\delta) \cos(\Phi)) = \arcsin(\cos(\Phi - \delta)) . \quad (5.8)$$

If we restrict to latitudes such that $\Phi - \delta \in [0; 180]$ (which implies $\Phi > \epsilon$ if we consider the whole year, a condition always true for any points north of the Tropic of Cancer), this can be further simplified to

$$\alpha_n = 90^\circ - (\Phi - \delta) . \quad (5.9)$$

⁷ A good approximation for the obliquity over several centuries can be found in the *Astronomical Almanac* published by the U.S. Naval Observatory [144].

Secondly, the solar azimuth γ is the angle between the geographic north and the sun's heading (north-clockwise by convention):

$$\begin{aligned} \gamma' &= \arccos\left(\frac{\cos(\Phi)\sin(\delta) - \cos(\delta)\sin(\Phi)\cos(\omega)}{\cos(\alpha)}\right) \\ \gamma &= \begin{cases} \gamma', & \text{for } \omega < 0 \\ 360^\circ - \gamma', & \text{for } \omega > 0 \end{cases} . \end{aligned} \quad (5.10)$$

An important derived quantity for our scenario is the angle of incidence θ of a direct beam relative to the normal vector for a south-facing surface, inclined by β relative to the horizontal

$$\theta = \arccos(\sin(\alpha)\cos(\beta) + \cos(\alpha)\sin(\beta)\cos(\omega)) . \quad (5.11)$$

Indeed, if the incoming beam has a power density given by G_n , the power density projected onto the tilted surface is given by $G = G_n \cos(\theta)$.

5.4 Theoretical Framework

A simulation framework was built using Matlab in order to clear out a number of important questions regarding our novel scenario. For example, what is the expected energy gain between a motionless solar panel and a mobile object tracking the sun (for now, we will neglect the energy required to move)? When can we expect a maximum energy income, and what is the ratio with the worst period of the year? Or what is the influence of the principal geometric parameters?

5.4.1 Case Studies

To explore the numerous possibilities and combinations offered by our scenario, we will define a number of case studies. The aim is to get an upper bound of the energy received by a unitary surface along the year. This also allows us to explore the relation between the received energy and the system's complexity. After all, a complex mobile robot may not be significantly better when compared to a motionless flower pot.

If we consider the system as a whole, we can isolate several components with their own Degrees of Freedom (DoFs). This breakdown helps to orthogonally analyse the system, and it will be always possible to combine components together at a latter stage of the design.

To begin with, the *base* can be one of the following: a) fixed for the whole year, b) fixed during the course of the day, or c) moving with the sun. The energy required to perform this kind of tracking is expected to be linked to the frequency of the displacements, as well as the amount of distance to be overcome. Thus, solution a) will be the less costly, but it will also provide




Figure 5.4 – Reference frame of the solar panel inside the global reference frame (x, y, z) of the room. The panel is located at \vec{p} , with azimuth γ_p (north-clockwise) and inclination β (with respect to the horizontal plane). The beam \vec{S} is coming from the sun, and makes an angle θ with respect to the panel's normal vector. The centre of the aperture is located at $(0, 0, H_w)$.

less energy along the year. Solution c) is the extreme opposite on the complexity scale. The position of the base inside the global reference frame is set to $\vec{p} = (x, y)$. The origin $(0, 0)$ of the frame is at the vertical of the aperture's centre, with the x axis oriented towards the north, as shown in Fig. 5.4.

Then, the *solar panel* can have up to two DoFs. Several sun tracking configurations are possible, depending on the position of the rotating axes [145]. This usually has no dramatic influence on the collected energy. For this work, we set the first axis to be vertical — i.e. collinear with the robot's axis of rotation — with a rotation angle γ_p (north-clockwise). The second axis is horizontal, with an inclination angle β . Figure 5.4 shows this configuration inside the global reference frame. This configuration has the significant advantage of possibly using the mobile base to perform the tracking along the vertical axis. For our case studies, each axis can be fixed at a given angle or can actively change its angle to track the sun.

We can define seven interesting cases, based on the combination of the above-mentioned configurations. For each case study, we are interested in the yearly average energy received by the considered location. Let us first define $G(h, d, \beta, x, y)$ as the instantaneous illumination on a surface inclined by β at any given location (x, y) , at time h (expressed in solar hours) during

the day d . This function equals

$$G(h, d, \beta, x, y) = \begin{cases} G_n \cos(\theta(h, d, \beta)), & \text{if } (x, y) \text{ is in direct sunlight,} \\ 0, & \text{otherwise,} \end{cases} \quad (5.12)$$

where $\theta(h, d, \beta)$ is given by Eq. (5.11).

Case 1: The Lonely Flower Pot The object has no DoF; it stays at its fixed location for the whole year. In this case the best location (x, y) to put this flower pot inside the room and the inclination angle β of the solar panel can be optimised. For simplification reasons, we will set the vertical rotation angle to 180° in order to have a south facing solar panel. As we are expecting the optimal placement to be in the direct alignment of the window, this is a reasonable guess. Although this case has no direct implication for a mobile robot, it is useful in order to establish a lower bound on the available energy at any point inside the room.

The map of the yearly average energy for a panel with a fixed inclination angle β is expressed in $\frac{\text{J}}{\text{m}^2 \cdot \text{day}}$ and is computed as

$$E_{c1, \beta}(x, y) = \frac{1}{365} \sum_d \int_h G_\beta(h, d, x, y) dt . \quad (5.13)$$

We can then search for the optimum location by using a global maximum search on the entire map. By repeating this operation on the search space $\beta \in [0^\circ; 90^\circ]$, we can also find the optimum inclination angle. The maximum energy is thus given by

$$E_{c1, \text{opt}} = \max_\beta \left(\max_{x, y} (E_{c1, \beta}(x, y)) \right) . \quad (5.14)$$

Case 2: The Optimized Flower Pot Compared to case 1, this object is placed at the beginning of each day at location (x, y) , where it will get the best illumination for the whole day. This could be a human taking care of the flower pot or a service robot stopping its nightly labour at the right location every day. We establish the yearly average energy as

$$E_{c2, \beta} = \frac{1}{365} \sum_d \max_{x, y} \left(\int_h G_\beta(h, d, x, y) dt \right) . \quad (5.15)$$

As in case 1, we can search for the best inclination angle β by performing a global maximum search, and the resulting energy is

$$E_{c2, \text{opt}} = \max_\beta (E_{c2, \beta}) . \quad (5.16)$$

Case 3: The Sunflower Pot Compared to case 2, we add a tracking along the vertical axis. The expected effect should be an improved energy accumulation along the day. Equations (5.15) and (5.16) remain valid, but we have to adapt the expression of $\theta(h, d, \beta)$. In this case, as the panel is always vertically aligned with the sun, $\cos(\omega)$ is always equal to 1, and Eq. (5.11) simplifies to

$$\begin{aligned} \theta_{c3}(h, d, \beta) &= \arccos(\sin(\alpha) \cos(\beta) + \cos(\alpha) \sin(\beta)) = \arccos(\sin(\alpha + \beta)) \\ &= 90^\circ - \arcsin(\sin(\alpha + \beta)) . \end{aligned} \quad (5.17)$$

If we further restrict to small angles such that $\alpha + \beta \in [-90^\circ; 90^\circ]$, this can be simplified to

$$\theta_{c3}(h, d, \beta) = 90^\circ - (\alpha + \beta) . \quad (5.18)$$

In our simulation, we will use the general formulation given by Eq. (5.17).

Case 4 Compared to case 2, and opposite to case 3, we add a tracking along the horizontal axis. This should also improve the energy accumulation along the day. In this case, Eq. (5.11) simplifies to

$$\theta_{c4}(h, d, \beta) = \theta_{c4}(h) = \omega(h) . \quad (5.19)$$

Case 5: The Evolved Sunflower Pot In this case, we combine case 3 and case 4 in order to have a tracking along both axes. As the panel is always aligned with the sun, θ_{c5} is always null. Consequently, $G(h, d, \beta, x, y) = G_n$ when (x, y) is in direct sunlight.

Case 6: The Simple Robot We now have a mobile robot that is tracking the location of the sun spot on the floor. As the robot can spin on itself, the tracking along the vertical axis is a natural consequence. For this case, we fix the panel's horizontal axis to an angle β . If the global illumination map $G_\beta(h, d, x, y)$ is known, which is always the case in our simulation, the path of the robot is unequivocally defined by

$$\vec{p}(h, d) = \underset{x, y}{\operatorname{argmax}} G_\beta(h, d, x, y) . \quad (5.20)$$

The yearly average energy in this case is computed by

$$E_{c6, \beta} = \frac{1}{365} \sum_d \int_h \max_{x, y} (G_\beta(h, d, x, y)) dt , \quad (5.21)$$

and the optimized inclination angle is found by performing a global maximum search, as in Eq. (5.16).





Figure 5.5 – Results of the simulated sun trajectories (blue curves, on the left of the masking wall). The three trajectories show extreme periods of the year: the winter solstice (lowest sun), the spring/autumn equinox, and the summer solstice (highest sun). The crosses mark the position of the sun for each solar hour, with the corresponding projections inside the room through the aperture (green patches). The solar noon is aligned with the aperture.

Case 7: The Complex Robot We add a second tracking axis to the panel. Our robot now has all the tracking capabilities: both regarding the position inside the room and the angle of the panel with respect to the incident sunbeam. Regarding the simulation, this is a combination of case 5 (two-dimensional tracking panel) and case 6 (mobile robot).

5.4.2 Results

All the simulations are run with a number of predefined parameters. First, the size of the room is fixed to $4 \times 4 \text{ m}^2$, as it corresponds to our experimental arena (see Sec. 6.1). The aperture has a size of $0.3 \times 0.3 \text{ m}^2$, and its centre is 1 metre above the ground. This is again to match our experimental setup. The computed illumination and energy maps have a spacial resolution of 5 cm, the search space for the optimum β is done with a resolution of 1° , and the simulation is run with a temporal resolution of 1 minute. For readability, all energies are given in watt-hours (Wh).

The simulated environment is shown in Fig. 5.5 with the sun's trajectory drawn for three important days of the year: the winter solstice (shortest day), the spring/autumn equinox (day and night are the same length), and the summer solstice (longest day). The position of the projected aperture is shown on the ground (green patches) for each solar hour. Figure 5.6 shows the maps for the illumination time during those three days (Fig. 5.6a) and the resulting accumulated energy by an horizontal surface (Fig. 5.6b). Both maps look very similar, but they are not exactly identical.



(a) Room's map of the accumulated illumination time in minutes. (b) Room's map of the accumulated energy, considering a horizontal panel ($\beta = 0^\circ$).

Figure 5.6 – Simulated illumination for three select days during the year: the winter solstice (on the right side of each figure), the equinox, and the summer solstice (on the left side). The aperture is at (0,0), and the x axis is oriented towards the north.

We can already draw a number of interesting conclusions. First of all, two elements limit the daily illumination inside the room. In winter or before the equinox, when the sun is low on the horizon, the projected spot moves rapidly inside the room. Consequently, when the sun has an elevation near the horizon ($\alpha \approx 0$) early in the morning or before sunset, the sun illuminates exclusively the side walls. This case is rejected by the simulation, as we only consider the beam when it illuminates the ground. The size of the room will have a small influence on this boundary condition. When we are approaching the summer, the sun is higher above the horizon. In this case, around the sunrise and the sunset, the sun azimuth is behind the south wall (thus directly illuminating the outside east and west walls, respectively), which corresponds to $\gamma < 90^\circ$ or $\gamma > 270^\circ$. This can be seen in Figs. 5.6a and 5.6b, as the illuminated surface does not extend to the side walls.

It is also clear from Fig. 5.6 that the sun stays longer near the zenith during the summer. The surface on the ground near the aperture rapidly accumulates energy. But due to the high elevation, the projected shape is relatively small. The resulting hotspot, which corresponds to a long illumination time and a high accumulated energy, is consequently spatially condensed. During the winter, the sun rapidly sweeps a large area on the back of the room. The illumination time is smaller, but the covered area is wider.

Case Study 1: The Lonely Flower Pot

The results for this first case study are shown in Fig. 5.7. This case is interesting in many regards. Figures 5.7a and 5.7b show the average on a whole year for the daily illumination time $\overline{T}_{\text{light}}(x, y)$ at any point of the room. While most locations inside the room see the light less than 15 minutes per day on average, there is one place close to the aperture (at $\vec{p}_{\text{opt}} = (0.47, 0)$) that receives the light up to 30 minutes per day. But one should keep in mind that this is a value averaged on the whole year. As we will soon see, this position has no direct sunlight for most of the year.

We can remark a strong correlation between the position of \vec{p}_{opt} and the trajectory of the sunbeam during the summer solstice. This is during this period that one can accumulate the most energy by standing at the same place, as the sunbeam's velocity on the ground is the slowest (see Fig. 5.5b on p. 81). We can thus devise an analytic approximation for \vec{p}_{opt} . At the summer solstice (we will use the "ss" subscript in this case), α_n equals

$$\alpha_{n,ss} = 113.45^\circ - \Phi , \quad (5.22)$$

or $\alpha_{n,ss} = 66.95^\circ$ with our reference latitude. Given the height of the aperture's centre H_w , its projection on the ground at the summer solstice is

$$\hat{\vec{p}}_{\text{opt},ss} = \left(\frac{H_w}{\tan(\alpha_{n,ss})}, 0 \right) = (0.425, 0) . \quad (5.23)$$

In our case, this estimator gives an error of 45 mm (9%). The real \vec{p}_{opt} is located slightly more at the back of the room to accumulate energy during the biggest number of days.


By placing a south-facing panel at \vec{p}_{opt} with a varying inclination angle β , we can perform a maximization of the accumulated energy, as explained with Eq. (5.14). The plot in Fig. 5.7c shows the average received energy per day \overline{E}_d as a function of β . The maximum is found at $\beta_{\text{opt},c1} = 26^\circ$, with an average energy of $501 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$. With such a fixed panel, Fig. 5.7e and 5.7f show the map $\overline{E}_{d,\text{opt}}(x, y)$. We can also devise an estimator for the optimal inclination by placing the panel perpendicularly to the summer solstice's beam

$$\hat{\beta}_{\text{opt},c1} = 90^\circ - \alpha_{n,ss} . \quad (5.24)$$


This estimator gives $\hat{\beta}_{\text{opt},c1} = 23.05^\circ$ in our case, which corresponds to an error of 3° . Again, this is due to the optimal location being slightly more in the back of the room compared to the position at the summer solstice.

Finally, Fig. 5.7d shows how a fixed light-sensitive object (be a plant or a photovoltaic gadget) placed at \vec{p}_{opt} might behave along the year. For most of the year (280 days of 365), this spot has no direct sunlight. The first ray of light comes on the 10th of May, the energy reaches a maximum of $2300 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$ during June, and the 3rd of August marks the last sunny day. Apart


Chapter 5. Theoretical Framework for an Indoor Harvesting Robot




(a) Room's map of the average illumination time per day, considering the simulation of an entire year. The aperture is at (0,0), the x axis is oriented along a south-north axis.




(b) Same as (a), top view.




(c) Average received energy at the best location (0.47,0), as a function of the panel's inclination angle β , and considering the simulation of an entire year. The maximum is reached for $\beta_{\text{opt}} = 26^\circ$.



(d) Received energy per day at the best location (0.47,0) and for an optimally inclined panel, as a function of the day of year.



(e) Room's map of the average received energy, considering the simulation of an entire year and for an optimally inclined panel at $\beta_{\text{opt}} = 26^\circ$.



(f) Same as (e), top view.

Figure 5.7 – Results for case study 1.

from a plant, most power-hungry objects would have trouble surviving with such lighting conditions during most of the year.

Case Studies 2, 3, 4, and 5

If one takes care of the plant, or if our light-sensitive object is able to move by itself before sunrise, we can chase the sun along the year and try to find an optimum position on a daily basis. This is a form of low-cost mobility, and it is the foundation of case studies 2 through 5. They differ from each other by the combination of sun tracking axes (horizontal or/and vertical) in use. Results are shown in Fig. 5.8.

When using a panel with a fixed inclination angle β (cases 2 and 3), we first need to perform a maximization step to find out the best inclination angle, as shown in Figs. 5.8a and 5.8c, respectively. Both cases have an optimum for $\beta_{\text{opt},c23} = 39^\circ$. In this case, case 2 can collect an average of $1290 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$, which is an improvement by 2.5 times compared to case 1. Case 3 slightly improves this figure by adding a tracking along the vertical axis, with an average of $1294 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$. This is a small gain of $4 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$ (0.3%), which is probably not worth the added complexity.

We can analyse what is happening on a day-to-day basis. As with case 1, the object should remain centred with the aperture ($y = 0$). Figures 5.8b and 5.8d present the daily received energy and the optimum position along the x axis, as a function of the day of year, for cases 2 and 3, respectively. The only difference between both figures is the absolute value of the received energy. Both cases reach a minimum of $414 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$ around the 21st of December, but case 3 performs slightly better with a maximum of $2323 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$ around the 17th of June, compared to $2312 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$ for case 2.

We can see that the optimum trajectory is identical in both cases, which is not surprising, as the illumination conditions are identical. The optimum position during the summer solstice is $(0.425, 0)$, which corresponds exactly to our estimator given in Eq. (5.23). Conversely, we can formulate an estimator for the farthest point from the window, during the winter solstice (“ ws ” subscript thereafter). In this case, α_n equals


$$\alpha_{n,ws} = 66.55^\circ - \Phi , \quad (5.25)$$

or $\alpha_{n,ss} = 20.05^\circ$ with our reference latitude. Thus, we have an estimator for the winter solstice


$$\hat{\vec{p}}_{\text{opt},ws} = \left(\frac{H_w}{\tan(\alpha_{n,ws})}, 0 \right) = (2.74, 0) . \quad (5.26)$$

This is not exactly what is shown by our simulations $\vec{p}_{\text{opt}} = (2.375, 0)$. This can be explained by looking at Fig. 5.6b (p. 82). In winter, the illuminated area has a relatively uniform density


Chapter 5. Theoretical Framework for an Indoor Harvesting Robot




(a) **Case study 2:** Average on one year of the received energy, as a function of the panel's inclination angle β . The optimum angle is for $\beta_{\text{opt}} = 39^\circ$.




(b) **Case study 2:** Received energy per day (solid blue), and optimal position along the x axis (dashed green), as a function of the day of year. The panel is fixed at β_{opt} .




(c) **Case study 3:** Average on one year of the received energy, as a function of the panel's inclination angle β . The optimum angle is for $\beta_{\text{opt}} = 39^\circ$.



(d) **Case study 3:** Received energy per day (solid blue), and optimal position along the x axis (dashed green), as a function of the day of year. The panel is fixed at β_{opt} , with a tracking along the vertical axis.



(e) **Case study 4:** Received energy per day (solid blue), and optimal position along the x axis (dashed green), as a function of the day of year. The panel is south-oriented, with a tracking along the horizontal axis.



(f) **Case study 5:** Received energy per day (solid blue), and optimal position along the x axis (dashed green), as a function of the day of year. The panel performs a tracking on both axes.

Figure 5.8 – Results for case studies 2, 3, 4 and 5.

of energy. Thus a large interval of points can be promoted as the optimum position, and our simulation has a bias toward the point closest to the aperture. If we use Eq. (5.26) to compute the projected position of the aperture's lower edge (by subtracting half of the aperture's height from H_w), we get (2.33, 0), which is very close to the point chosen by the simulation (and with a spatial resolution limited to 0.05 m, the difference becomes insignificant). More generally, a good guess for the optimum position at any day of the year is

$$\hat{p}_{\text{opt}}(d) = \left(\frac{H_w}{\tan(\alpha_n(d))} \right) = \left(\frac{H_w}{\tan(\arcsin(\cos(\Phi - \delta(d))))}, 0 \right). \quad (5.27)$$

Case 4 explores the use of a tracking along the horizontal axis, and case 5 enables the tracking on both axes. Again, the optimum trajectory remains the same because the optimum location depends only on the base's mobility and not on the panel's mobility. Case 4 gets a yearly average energy of $1322 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$ (a gain of $32 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$, or 2.5%, compared to case 2), whereas case 5 gets $1335 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$ (a gain of $45 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$, or 3.5%, compared to case 2).

Case Studies 6 and 7: The Mobile Robot


The last two cases enable a full mobility of the base, for example, by using a mobile robot. Both cases have a tracking along the vertical axis because most mobile service robots can usually rotate on themselves. In addition, the last case enables a tracking along the horizontal axis. Results are shown in Fig. 5.9.

The optimum inclination of the panel in case 6 equals $\beta_{\text{opt},c6} = 53^\circ$, as shown in Fig. 5.9a. This scenario leads to a yearly average energy of $7412 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$, which is an improvement of about 5 folds compared to the previous cases. The full sun tracking solution, as brought by the last case study, even leads to an average of $7651 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$. Of course, these figures are impressive, but one has to keep in mind that the simulation does not take into account the energy required to continuously move and track the sun. This has a cost that could hinder — or even negate — the benefits offered by these solutions. This will be sorted out during the experimental study of Sec. 6.4.


Let us analyse what is happening on a day-to-day basis. Figures 5.9b and 5.9c show, for case 6 and 7, respectively, the daily energy received by the robot, as well as the number of hours where the light enters the room. Both curves are closely related because the robot is now able to make use of every ray of light inside the room. It can also be seen how the efficiency is slightly better during the summer, when the sun is high in the sky, especially for case 7 where the inclination of the panel can be adjusted in relation to the sun's altitude.

The number of workable hours of sun per day needs a small explanation. We can see in Figs. 5.9b and 5.9c that the maximum is reached two times per year, on the 1st of May and the 12th of August, with just above 10 hours of sun per day. It may seem counterintuitive that the summer solstice produces fewer hours of sun per day (around 8 hours and 43 minutes). Al-


Chapter 5. Theoretical Framework for an Indoor Harvesting Robot



(a) **Case study 6:** Average on one year of the received energy, as a function of the panel's inclination angle β . The optimum angle is for $\beta_{\text{opt}} = 53^\circ$.



(b) **Case study 6:** Received energy per day (solid blue), and total illumination time per day (dashed green), as a function of the day of year. The panel is fixed at β_{opt} , with a tracking along the vertical axis.



(c) **Case study 7:** Received energy per day (solid blue), and total illumination time per day (dashed green), as a function of the day of year. The panel performs a tracking on both axes.

Figure 5.9 – Results for the case studies 6 and 7.

though it is true that the summer solstice is the longest day in the year outside, the constraints on our indoor environment change the behaviour of the problem. Between the 1st of May and the 12th of August, the limiting factor is not the elapsed time between the sunrise and the sunset but the geometry of the sun path with respect to the south-facing wall. Indeed, when the summer reaches its peak, the light begins to enter by the aperture of the south wall about one hour after sunrise and conversely disappears about one hour before sunset. Indeed the sun's azimuth must satisfy the condition $\gamma \in [90^\circ; 270^\circ]$ to shine inside the room. Early in the morning, the sun shines only on the outside of the east wall as its azimuth is below 90° . Conversely, the sun shines only on the outside of the west wall late in the afternoon at its azimuth goes past 270° .

An analytical expression for the ideal robot's trajectory can be derived. If one knows the height of the aperture's centre H_w , the time of day h , and day of year d , the optimal robot's state $(\vec{p}, \beta, \gamma_p)$ is

$$\vec{p}_{\text{opt}}(h, d) = \left(\frac{H_w}{\tan(\alpha(h, d))} \cos(\omega(h)), \frac{H_w}{\tan(\alpha(h, d))} \sin(\omega(h)) \right), \quad (5.28)$$

$$\beta_{\text{opt}}(h, d) = 90^\circ - \alpha(h, d), \quad (5.29)$$

$$\gamma_{p,\text{opt}}(h, d) = \gamma(h, d). \quad (5.30)$$

The robot's position \vec{p} is set to follow the projection on the ground of the aperture's centre. This can be somehow relaxed by taking into account the size of the aperture, and thus the size of the beam on the ground, minus the size of the robot's panel. If the panel is small enough compared to the light beam on the ground, the robot does not need to move continuously. It can wait until it reaches the edge of the bright spot before moving again to its optimal position. This leads to a temporally discontinued displacement and a reduced use of energy.

5.5 Concluding Analysis

The cases that we studied in this section cover a large range of scenarios and produced some interesting results. To better see the big picture, a summary of the important statistics for each case study is shown in Table 5.4.

The yearly average illumination time \bar{T}_{light} is shown to depend directly on the base's mobility. With a higher mobility, we can increase the number of hours where the sun shines directly on the panel. A motionless object will not see the sun for most of the year, whereas a mobile robot could get up to seven and a half hours of sun on average. As a logical consequence, the daily accumulated energy is also directly related to the base's mobility, as it can be seen with the yearly average \bar{E}_d , as well as $\min_d(E_d)$ and $\max_d(E_d)$. This is thus the single most influential design choice and the reason why a mobile robot could outperform other simpler solutions.

On the contrary, the sun tracking performed by the panel has much less influence on the resulting gain. As an example, the gain between case 2 (no tracking) and case 5 (tracking on both axes) is only $45 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$ (3.5%) on average. This is slightly more interesting in the case of a mobile robot, with a gain of $239 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$ on average, but this represents only 3.2% compared to the total energy. The usefulness of this addition can only be evaluated with the knowledge of the energy required to move the panel.

In the case of a panel with a fixed horizontal axis, the optimal inclination angle β_{opt} can be computed. This parameter is also directly related to the mobility of the base. As it was shown with Eq. (5.24), the yearly optimum angle for a fixed object is close to the optimum angle at

Table 5.4 – Summary table for the results of case studies 1 to 7. For each case, the mobility of the base and panel are decisive. \bar{T}_{light} is the yearly averaged illumination time received by the robot, and \bar{E}_d is the corresponding received energy. $\min_d(E_d)$ is the minimum amount of energy during the year, and $\text{argmin}_d(E_d)$ is the day when it happens. Conversely, $\max_d(E_d)$ and $\text{argmax}_d(E_d)$ give the same information, but for the maximum amount of energy.

Case	Mobility	γ_{panel}	β_{opt}	\bar{T}_{light} min day	\bar{E}_d Wh m ² .day	$\min_d(E_d)$ Wh m ² .day	$\text{argmin}_d(E_d)$	$\max_d(E_d)$ Wh m ² .day	$\text{argmax}_d(E_d)$
1	Fixed	180°	26°	30.2	501	0	[04Aug, 09May]	2308	02Jun
2	Daily Fixed	180°	39°	80.1	1290	414	21Dec	2312	17Jun
3	Daily Fixed	Tracking	39°	80.1	1294	414	21Dec	2323	17Jun
4	Daily Fixed	180°	Tracking	80.1	1322	483	16Dec	2360	17Jun
5	Daily Fixed	Tracking	Tracking	80.1	1335	483	16Dec	2400	17Jun
6	Tracking	Tracking	53°	458	7412	4494	21Dec	9763	{1May, 12Aug}
7	Tracking	Tracking	Tracking	458	7651	4767	17Dec	10030	{1May, 12Aug}




Figure 5.10 – Amplitude between the two most distant positions reached by the robot, along the x axis and for a whole year, as a function of the place's latitude Φ . The red dash-dotted lines mark the position of the two polar circles at $\Phi = \pm 66.55^\circ$.

the summer solstice. When increasing the mobility, the optimization will gradually choose steeper angles to maximize the harvesting along the whole year.

A remaining question is the influence of some of the problem's geometric parameters, such as the placement and size of the aperture or the room's dimensions. Let us begin with the aperture. We already saw with Eqs. (5.23), (5.26), and (5.27) that the optimal placement at noon is along a south-north axis passing by the aperture's centre, and that the distance along the x axis from the wall is geometrically related to the **height of the aperture's centre**. The yearly amplitude Δx along the x axis can be analysed based on the difference between Eq. (5.23) and (5.26)

$$\begin{aligned} \Delta x &= H_w \left(\frac{1}{\tan(\alpha_{n,ws})} - \frac{1}{\tan(\alpha_{n,ss})} \right) = H_w \frac{2 \cdot \sin(\alpha_{n,ss} - \alpha_{n,ws})}{\cos(\alpha_{n,ws} - \alpha_{n,ss}) - \cos(\alpha_{n,ws} + \alpha_{n,ss})} \\ &= H_w \frac{2 \cdot \sin(2\epsilon)}{\cos(2\epsilon) + \cos(2\Phi)}. \end{aligned} \quad (5.31)$$

The result depends only on the height of the aperture's centre H_w , the Earth axial tilt angle ϵ , and the considered latitude Φ . Figure 5.10 shows Δx as a function of the latitude, using a unitary height H_w .

The asymptote is located at

$$\begin{aligned}\cos(2\Phi_A) &= -\cos(2\epsilon) \\ \Phi_A &= \pm \frac{1}{2} \arccos(-\cos(2\epsilon)) = \pm(90^\circ - \epsilon) = \pm 66.55^\circ .\end{aligned}\quad (5.32)$$

This value corresponds to the location of the polar circles, where no light shines for a part of the year. The values of Δx beyond this latitude are not valid anymore due to this restriction. Figure 5.10 clearly shows the feasibility area for an indoor robot that needs to operate all year round solely based on the solar power, i.e. places with a latitude $\Phi \in [-60^\circ; 60^\circ]$. The feasibility further decreases if we limit the allowable amplitude inside the room. The minimum required amplitude corresponds to locations on the Earth's equator ($\Phi = 0$). In this case,

$$\min(\Delta x(\Phi))_{\Phi=0} = 2 \cdot H_w \frac{\sin(2\epsilon)}{\cos(2\epsilon)} = 2 \cdot H_w \tan(2\epsilon) \approx 2.14 \cdot H_w . \quad (5.33)$$

The **size of the aperture** will of course change the size of the beam on the ground. This has two benefits. First, a bigger aperture will increase the feasible exposition time (the width will increase the daily exposition time, whereas the height will compensate for the change in altitude of the sun along the year). This will be of notable help for fixed objects, but a mobile robot will not gain a lot, as this will change only the boundary conditions when reaching the edge of the room. Secondly, in the case of a mobile robot, a bigger beam will induce fewer wake-up events in order to move the robot back inside the beam, thus sparing energy in the long run. The total amount of path to travel remains the same, however. On a design note, the size of the aperture will also limit the size of the usable photovoltaic panel, thus the total amount that one can harvest.

The **size of the room** will influence the boundary conditions. A larger room will slightly increase the hours of light per day, when this dimension is the limiting factor (for example, the period before the 1st of May or after the 12 of August in Fig. 5.9b on p. 88). A room that is not deep enough will also limit the possibility of solar charging during the winter, when the sun is possibly hitting exclusively the back wall during the winter's peak, leaving the robot starving on the floor.

6 Experimental Validation of an Indoor Harvesting Robot

Chapter 5 explored the possible sources of indoor energy, showed how the sun coming through the apertures of the building can be a promising supply, and set a theoretical upper bound on the available energy along the year with respect to the mobility of the system. This theoretical framework used a number of hypotheses, including a null power consumption of the system. In reality, this is of course not the case, especially when tracking the sun's location along the day. The goal of the present chapter is to get a practical upper bound on the available energy by using a real system, inside a controlled environment.

We need three important pieces of hardware to experimentally challenge the results of our simulations:

1. A controlled environment: In any scientific work, it is necessary to have repeatable conditions and reproducible results. This implies the development of a controlled environment and the ability to set representative conditions with respect to the expected conditions on a whole year. A controlled environment is also important in order to get a consistent behaviour when comparing the performance of different solutions. For this reason, an artificial mobile lighting system was developed to fit inside our experimental arena to reproduce the sun's trajectory. The main design keys and the characterisation process are presented in Sec. 6.1. Of course, further developments would need to assess the performance of various algorithms inside a real setting. But such a setup is difficult to work with and time-consuming during the early development phase, especially when fine-tuning the algorithms.
2. A solar-harvesting module for our robotic platform: This is the final piece of hardware that will enable the marXbot robot to autonomously harvest some energy from its environment. The design of this module is crucial, as it will heavily influence the overall efficiency of the harvesting process. The important design considerations are reported in Sec. 6.2¹.

¹ Parts of Sec. 6.2 were previously published in [146].

3. A versatile energy logging tool: This tool must meet a number of requirements. First, it should provide a good insight of the inner workings of the subsystems constituting the robot. Then, its precision must be good enough to allow a precise assessment of the overall energy balance. Finally, it must not influence the measured system. Such a tool was already developed for the measurements of Sec. 4.4 and is presented in Appendix A.

Once the required hardware is available, the experiment can take place. The experimental methodology and the algorithms are presented in Sec. 6.3. The experiment is divided into three phases: the building of the environment's map, the systematic coverage of the surface, and finally the autonomous tracking of the solar beam to recharge the robot. Results are described and analysed in Sec. 6.4 and Sec. 6.5. The final discussion is performed in Sec. 6.6.


6.1 The Solar Arena

Our controlled environment is made of two parts. The first piece of hardware is the experimental arena already used in Sec. 2.3. This arena has a controlled artificial lighting system, and a well-defined shape. The second piece of hardware is the solar simulation system. The aim is to faithfully simulate the sun's trajectory, with respect to the aperture made in the separation wall. The power density of the beam should also match the one of the sun as closely as possible. We will describe in this section the mechanical design and the experimental validation of the system's conformity compared to the real sun.

6.1.1 Design

The main challenge, within the solar arena, is the design of the mobile solar simulation system. Current solar simulators are motionless, and they are used to characterise the performance of solar panels. In our scenario, this system has to move along the day, according to the solar mechanics developed in Sec. 5.3. In order to provide our scenario with a realistic simulation tool, a number of characteristics must be met. Let us set out the main design requirements and the retained solutions.

Mechanical Trajectory The trajectory of the system, with respect to the aperture, should match the one of the sun as close as possible. This is the primary requirement. Exact kinematics is, however, hard to achieve with a simple system. Figure 6.1a visually shows the difference between the trajectory of the sun and the achievable trajectory obtained by spatially shifting a circle of constant radius. More specifically, Fig. 6.1b looks at the difference for the altitude angle α between both trajectories, as a function of the solar hour and the day of year. It can be observed that the ideal circle deviates from the sun's trajectory by a few degrees, especially when approaching the solstices. The difference remains small, however, and this solution is retained for its simplicity.



(a) Simulated isometric view of the trajectory achieved by the fixed radius circle (dashed blue curve) and the scaled trajectory of the sun (solid red curve). The crosses are placed at every solar hour, and the three trajectories correspond to the winter solstice, equinox, and summer solstice (from bottom to top, respectively).

(b) Comparison of the altitude angle α between both trajectories for seven different days along the year, between the summer solstice ($d = 172$) and the winter solstice ($d = 355$).

Figure 6.1 – Comparison of the trajectory achieved by the proposed solution of a circle with a fixed radius, and the actual solar trajectory (as seen by an observer located at the aperture).

Figure 6.2 shows the mechanical design of the solar simulation system and its integration inside the arena. The mechanical schema is detailed in Fig. 6.2a, while the corresponding hardware is depicted in Fig. 6.2b. The working principle is fairly simple: when the system is vertical (i.e. comprised inside the $(x - z)$ plane), the angle made by rod ② with the vertical defines the considered latitude Φ . This joint is thus fixed during the experiment, unless one wants to experiment with different latitudes. The angle made by the light emitting device ③ with the vertical defines the incident angle of rays at noon α_n , and therefore relates to the simulated day of year, following Eq. (5.8) on p. 76.

Along the day, the system rotates around the axis defined by rod ①, following the relation $\omega(h)$ defined by Eq. (5.6) on p. 76. This is achieved by using a stepper motor and dedicated electronics. As the joint between rods ① and ② is fixed for a given latitude, the radius of the circle described by this system is fixed irrespective of the day of year, as designed for the retained solution.

There is an additional geometric constraint on the system. Because the light emitting device has a limited size, comparable to the size of the aperture, it must remain aligned with the aperture during the whole experiment. This is achieved by aligning the axis of rod ① and the normal to the light emitting device ③ with the aperture's centre, as shown in Fig. 6.2a. By doing so, the light rays will be always aligned with the aperture for any angle ω . As a consequence, when changing the simulated day (by adjusting the




Figure 6.2 – Mechanical overview of the solar arena. The solar simulation system is on the left side of the wall ($x < 0$); the arena is on the right side ($x > 0$). The system is depicted when located inside the $(x - z)$ plane, i.e. when $w(h) = 0$. Along the day, the system rotates along its main axis, leading to a displacement of the projected spot inside the arena. Φ is the simulated latitude; $\alpha_n(d)$ corresponds to the incident angle at solar noon, and thus to a specific day of year.

angle α_n), the height of the support under the system must be adjusted in order to keep this geometric constraint satisfied.

Light Power Density The light emitting device is another critical piece of hardware. The light rays coming out of the device should have parallel trajectories, with a uniform power density across the beam, in order to match the conditions experienced by an observer on Earth. In our system, this device is made of two pieces: a light bulb, and a 380 mm off-the-shelf parabolic reflector fixed on the back. The chosen halogen bulb is designed for a professional lighting system, with six independent filaments in order to homogenise the power density over a large area. The parabolic reflector on the back gathers and focuses the backward rays, thus enhancing the overall efficiency of the system. This system is considerably lighter compared to an optical system made of Fresnel lenses (as used inside lighthouses), which is an advantage given the substantial lever effect caused onto the mechanical actuator.

Light Spectrum It is desirable to have a light source with a spectrum close to the sun. Indeed, solar panels have a quantum efficiency that dependants on the incoming wavelength. If we want to have results that are transposable to the real world, the spectrum of the light bulb should be close to the American Society for Testing and Materials (ASTM) G173 standard [141].

Several options exist. Static solar simulators generally use plasma or metal-halide (xenon) lamps. They have a baseline spectrum similar to the sun, usually with some additional peaks (around 900 nm for metal-halide lamps). Combined with filters, they are used to characterise solar panels in laboratories.

For our solar simulator, we have decided to use halogen lamps. Although their spectrum is shifted towards infrared wavelengths, they are more affordable. Figure 6.3 compares the reference ASTM E490 [147] (extraterrestrial irradiance) and ASTM G173 (irradiance at sea level) spectral power densities with the computed black body model of our halogen lamp² (colour temperature of 3200 K). The power density of the black body model was normalized to 1000 W/m^2 to match ASTM G173. Crystalline silicon has a bandgap of 1.11 eV at 302 K [148], and consequently the upper limit for the quantum efficiency is around 1100 nm (thick dashed line in Fig. 6.3). Light with a higher wavelength will not be converted into an electron-hole pair. The excess of infrared light principally happens beyond the working limit of silicon panels. This will probably result in an increased heating of the solar panel. Below 1000 nm, the halogen lamp has a lower spectral power density, resulting in less harvested power. The overall effect should be a lower photovoltaic efficiency with our simulation system, compared to a usage in the real sun.

²Model: Ushio HPL+ 575/230V.




Figure 6.3 – Comparison of the spectral power density of the two reference models (extraterrestrial irradiance ASTM E490, and irradiance at sea level ASTM G173), with the black body emission of a halogen lamp (colour temperature of 3200 K). The spectral bands are given according to ISO 20473. The thick dashed line shows the equivalent wavelength for the bandgap energy of crystalline silicon (1100 nm), which is the upper limit for the photovoltaic conversion.

The integration of the simulation system inside the arena is done as shown in Fig. 6.2c. The reference coordinate frame is set exactly like the one of the simulation performed in Sec. 5.4, and the free space between the four walls is likewise a square surface of $4 \times 4 \text{ m}^2$. The size of the aperture made in the separation wall is likewise $0.3 \times 0.3 \text{ m}^2$, with its centre located 1 m above ground. The physical realisation is shown in Fig. 6.2d, and an illumination test is done in Fig. 6.2e. It is important to note that the size of the aperture should be smaller than the diameter of the beam to have a satisfying occultation.

For the duration of our experiments, we have fixed the latitude angle Φ to 46.5° and the incident angle at noon α_n to 30° . This incident angle corresponds to the conditions experienced during both the 14th of February ($d = 45$) and the 26th of October ($d = 299$). If we consider the whole year, 70% of the days will experience better solar conditions (refer to Fig. 5.9c on p. 88), and we are close in amplitude to the worst conditions experienced during the winter solstice. This setting thus offers representative experimental conditions and enables us to extrapolate our results to the real world with confidence.

6.1.2 Experimental Validation

We have undertaken several validation steps, according to the requirements set in the previous section. The mechanical design is validated by comparing the resulting light spot on the ground with the software simulated counterpart for every solar hour from 9am to 3pm. The results are shown in Fig. 6.4a. They are obtained by blending together several images taken by the overhead camera. The results of the software simulation were reported by hand on the ground. The experimental solar simulator matches pretty well the results predicted by the software simulation. The spot on the ground is slightly bigger, however, by about 10 to 20% in each direction. Indeed, the rays coming from the light source are not perfectly parallel. But the reflector added on the back of the light is still effective, and the light source is far from being omnidirectional. This can be verified by comparing the software simulations for a light source with perfectly parallel rays (Fig. 6.4b) and a perfectly punctual source (Fig. 6.4c). The results of Fig. 6.4a are clearly closer to an ideal solar source than a simple light bulb.


The characterisation of the power density over a cross-section of the beam is shown in Fig. 6.5. These measures are done by systematically placing a solarimeter³ on the grid of a target. The target is inclined by $\beta = 90^\circ - \alpha_n = 60^\circ$ in order to be perpendicular to the incoming beam, and it is placed at 2.25 m from the light source, as shown in Fig. 6.5c. The grid has a fine spatial resolution of 0.05 m, resulting in 20 by 20 measurement points over an area of 1 m². For the first measurement, the power density is measured over the whole beam, without any aperture, whereas the second measurement is done with the square aperture in place. Results are found in Fig. 6.5a and 6.5b, respectively. The shape of the beam is well defined in both cases, with sharp edges. The projected square of Fig. 6.5b measures 400 mm x 350 mm (width x height), which is slightly bigger than the aperture. If we take the points inside the square beam, the box plot of the power density is shown in Fig. 6.5d. The power density is not perfectly uniform inside the beam, and it has a few outliers with a higher value. They are most likely the image of the bulb's filaments. The median of the values equals 515 W/m², which is below the standard power density of 1000 W/m². This is not a big issue because real-world conditions are usually below the standard conditions due to clouds and other atmospheric perturbations. We can always scale the results of our simulations to match the experimental conditions.

The light spectrum of the halogen source is hard to validate, as we do not have the necessary hardware to perform such measurements. It is not critical, however, as we are not interested in the detailed spectral power density but only in the total power density that can be potentially converted by our solar panel. This measure was already carried out when measuring the power density in Fig. 6.5, with a sensor using the same silicon technology. Thus the integral of the spectral power density should give similar results with our solar panel.


³ Model: Kimo SAM20. The accuracy is 5% over the measurement range [1; 1300] W/m². The sensor is made of polycrystalline silicon with a spectral response between 400 nm and 1100 nm.



(a) Fused images as seen from the overhead camera when the light source moves from one hour to the next. Green rectangles: Position given by the software simulation. White line: Room's boundaries.





(b) Simulation of the light's location on the ground every hour, in the case of an ideally far sunlight with parallel rays.




(c) Simulation of the light's location on the ground every hour, in the case of an omnidirectional light source, such as a light bulb without any focusing device.

Figure 6.4 – Experimental validation of the mechanical design of the solar simulator. The green patches are based on the software simulations. Simulation parameter: $d = 45$ ($\alpha_n = 30^\circ$), $\Phi = 46.5^\circ$.



(c) Experimental setup used for the characterisation of the power density. They are the light source (right), the aperture (middle), and the target (left).



(d) Box plot showing the distribution of the power density inside the beam, in the case of the square aperture.

Figure 6.5 – Characterisation of the beam’s power density. Distance of the target from the light source: 2.25 m.

6.2 The Solar Harvesting Module

Our scenario requires the robot to be able to harvest energy from the incoming sunlight. For this purpose, we need a versatile solar harvesting module to put on top of the marXbot robot. This section presents the design and the experimental validation of this module. The final hardware is shown in Fig. 6.6.

6.2.1 Design

We have a number of requirements for this module. It must be low-power to minimise the losses and efficiently accumulate the energy provided by the solar panel, while being able




Figure 6.6 – Pictures of the solar harvesting module. The harvesting electronics sits on top of the laser scanner, while the interface electronics sits in-between the mobile base and the processor board.

to interface with the robot through the Controller Area Network (CAN) bus. It should also have an access to the main Inter-Integrated Circuit (I²C) bus to manage the robot's resources and replace the main Central Processing Unit (CPU) if we decide to sleep or shut down the processor board. This requirement is fulfilled by dividing the module into two dedicated electronics. The first low-power electronics is in charge of controlling the harvesting process by measuring the panel's states and performing an optimisation of its operating point. This part of the module is always switched on. The second electronics serves as an interface with the robot by hooking into the CAN and I²C busses, and it can control the robot when the CPU is not available. This part can be switched off on demand, and it can also enter a low-power sleep state for a quantified amount of time.

The module must also optimise the placement of the solar panel with respect to the incoming light. We have seen in Sec. 5.4 that a mobile differential robot is able to orient itself along the vertical axis as a part of its displacement process (case study 6). Our solar module will add a tracking angle along the horizontal axis to build the system studied by case 7. We are expecting a gain of a few percentage points with such tracking, but moving the solar panel has a cost. This cost should be minimised and put in balance with the achieved gain during the experiment.

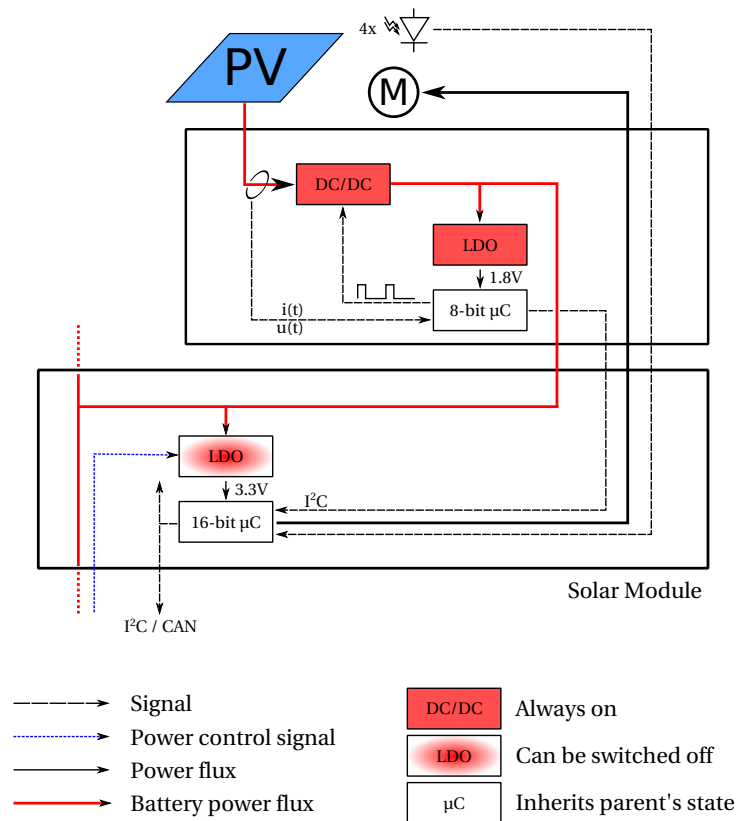


Figure 6.7 – Diagram of the electronics inside the solar harvesting module. The electronics is split in two distinct boards, connected through wires. The lower part is the interface electronics with the marXbot robot, while the upper part is the low-power harvesting electronics.

Electronics

The diagram of the module's electronics is shown in Fig. 6.7. The interface electronics is similar to most modules shown in Fig. 3.4 (p. 29). The 16-bit microcontroller can be turned off by acting on the Low-DropOut (LDO), but the firmware also allows the microcontroller to enter a deep sleep state. The electronics of the photovoltaic harvester uses a low-power 8-bit microcontroller.

The principal element of our system is the photovoltaic cell. We have characterised and compared four different photovoltaic cells. Two cells were made of monocrystalline silicon, while the two others were made of amorphous silicon. During the characterisation process, the current/voltage curve was recorded using a variable resistive load. We considered two lighting conditions: the sunlight, as received behind a south-oriented window, and an artificial halogen light similar to our solar arena. One of the monocrystalline cell⁴ then outperformed

⁴Sunways AH508480G.

the others. It has an active area of 190 cm^2 with a rated efficiency of 18.6%, which is enough to produce 3.5 W under standard conditions.

This single cell is providing a high current (8 A under standard conditions) at low voltage ($V_{\text{open circuit}} = 0.63\text{ V}$). Due to space constraints, only one cell fits on top of the marXbot, thus we cannot reach a higher voltage by chaining them in serial. A boost DC/DC converter is mandatory to increase the voltage up to the one provided by the battery. But no off-the-shelf DC/DC converter could work at such an operating point, so we had to make one ourselves using discrete components. As shown in Fig. 6.7, the 8-bit microcontroller senses the voltage and the current at the output of the solar panel, and then controls the DC/DC converter with a Pulse-Width Modulation (PWM) signal.

The elements of the harvesting electronics were optimised to reduce the power losses. This includes the Joule losses, the switching losses, and the leakage currents. It is hard to find the real optimum due to the huge number of discrete components and technologies available on the market, making a tremendous number of possible combinations. An iterative approach was used to solve this problem: at each iteration, the losses for each component are computed, the dominant losses are identified, and the corresponding components are changed for the next iteration. Results are further discussed in [1]. Despite this optimisation, losses are inevitable. When the photovoltaic panel operates at low power, the dominant losses come from the power consumed by the microcontroller and the reverse leakage current of the DC/DC's diode. When the power increases, and consequently increases the current going through the system, the Joules losses inside the transistors and the inductor prevail.

Solar panels have a power at the output following a bell-shaped curve. The maximum-power point is reached only at a specific load. The Maximum Power Point Tracking (MPPT) is done by the microcontroller. The impedance of the DC/DC converter is controlled by changing the duty cycle of the applied PWM. Numerous MPPT algorithms exist and have been studied in the literature [149]. The study in [150] compares the energetic performance of 10 widely adopted algorithms and their associated cost. Based on their findings, we chose to use the *Perturb and Observe* algorithm for its versatility, performance, and ease of use. The main disadvantage is a constant oscillation around the maximum-power point, giving rise to some ripples on the output current and voltage.

The electronics includes several other components with the goal of increasing the collected energy. A photodiode is placed at each corner of the photovoltaic panel to easily perform a gradient ascent and find the location of the sun spot on the ground. These photodiodes are connected to the Analog Digital Converter (ADC) of the main microcontroller. The stepper motor, used to actuate the mechanics, is also driven by the main controller.




Figure 6.8 – Mechanical concepts envisioned for the solar harvester. Concepts 1 to 3 use a leadscrew (black rectangle) to convert the motor’s rotary motion into a linear speed. Concept 4 uses the motor in direct drive. Images credits: Master thesis of M. Liniger [1].

Mechanics

Several mechanical concepts were studied for the rotation of the panel around the horizontal axis. Figure 6.8 shows the four retained candidates. Three concepts use a leadscrew to convert the motor’s rotary motion into a linear speed that will then actuate the panel around its rotation axis. They offer a very good transmission ratio, allowing us to use a low-power actuator. The last concept directly actuates the pivot joint with the geared motor.

We studied the energy required by each concept to move the panel between the two extreme angles. Concept 4 requires less energy. The main disadvantage of concepts 1 to 3 is the low efficiency of most leadscrews compared to gearboxes. Concept 4 is also the simplest one. Figure 6.6a (p. 102) shows the final system. It uses a geared stepper motor in direct drive with the rotation axis. The printed circuit board of the harvesting electronics serves as a mechanical support for one of the two roll bearings.

6.2.2 Experimental Validation

A simple model of the harvester is shown in Fig. 6.9. The efficiency of the harvester $\eta_{\text{harvester}}$ is limited by the efficiency of the two constitutive elements. The first element is the photovoltaic cell, which is converting the incoming illumination into electricity, with an efficiency η_{cell} . The second element is the boost DC/DC converter, which has an efficiency $\eta_{\text{DC/DC}}$. Finally, the output power of the harvesting module flows to the robot. If the consumption of the robot is lower than the harvested power, the excess goes into the battery.

The efficiency of the photovoltaic cell is experimentally measured under halogen lighting conditions with a power between 100 and 400 W/m² (Fig. 6.10a). The median of the efficiency

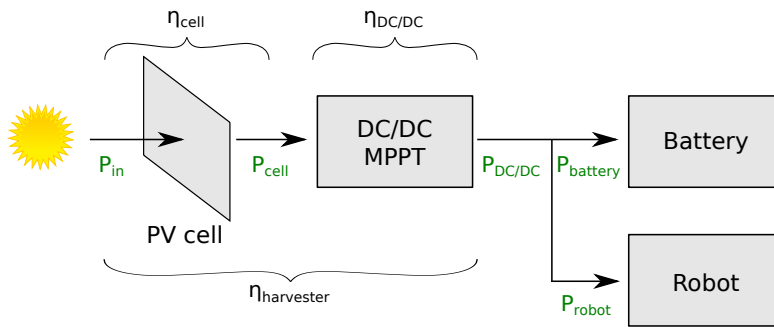


Figure 6.9 – Simple model of the harvesting module.


is at 12.6%, which is below the nominal efficiency of the cell (18.6%). Two causes are the likely reason for this low efficiency. First, the measures were done under halogen lighting. As we saw in Sec. 6.1.1, and especially in Fig. 6.3 (p. 98), the spectrum power density of such lights is shifted towards infrared compared to the natural sunlight. As our solar cell cannot harvest electromagnetic waves above 1200 nm, part of this power is lost. Second, our test conditions have a lower illumination level compared to standard test conditions (1000 W/m^2). It is known that photovoltaic cells have a lower efficiency under low light conditions [151].

The efficiency of the DC/DC converter depends on the current flowing through it. Figure 6.10b plots the measured efficiency as a function of the illumination on the photovoltaic panel. It also shows the computed efficiency based on the model that we used to optimise the electronics and reduce the losses inside the system. The measures match well the model at low illumination ($G < 100 \text{ W/m}^2$). At higher illumination, some unmodeled Joules losses further decrease the practical efficiency of our converter.


The efficiency of the overall harvesting module is plotted in Fig. 6.10c. The achieved efficiency is especially good at low illumination. The harvester's efficiency is above 10% in the range between 30 and 110 W/m^2 . At higher illuminations, the high current delivered by the cell increases the Joule losses inside the inductor and the high-side transistor of the DC/DC converter, reducing the system's efficiency.

6.3 Methodology and Algorithms


We have designed our experiment in order to simulate the conditions experienced by indoor service robots. The robotic vacuum cleaners of Chap. 2 are good models because they have a high energy requirement to cover the whole surface. Such an application is very demanding and offers a good point of comparison for the viability of our approach. The experiment is divided into three different phases, each one with its own task and set of algorithms. The configuration of the solar arena for this experiment is shown in Fig. 6.11.



(a) Conversion's efficiency of the lighting power into electricity by the photovoltaic cell, under halogen lighting and at illumination levels between 100 and 400 W/m² ($N = 60$).



(b) Efficiency of the DC/DC converter, as a function of the illumination on the photovoltaic panel. Both the mathematical model (blue crosses) and the measures (green dots) are shown ($N = 5$).



(c) Overall efficiency of the harvesting module, as a function of the illumination on the photovoltaic panel ($N = 5$).

Figure 6.10 – Measured performances of the photovoltaic harvesting module. Data credits: Master thesis of M. Liniger [1].



Figure 6.11 – Picture of the features inside the solar arena during the final experiment. The aperture and the solar simulation system are in the lower-right corner.

Phase 1: Map building We have seen in Chap. 2 that a robot which relies on a localisation device can perform its task better, faster, and with less energy compared to a robot relying on a random navigation. When entering an unknown environment, the first task is thus to build a map of the environment.

On the solar marXbot robot, this is achieved by combining the Neato laser scanner of Chap. 4 with a Simultaneous Localisation and Mapping (SLAM) algorithm. To this end, the ROS framework⁵ is used with the “hector mapping” stack. This SLAM algorithm is fast, gives good results, and does not need a lot of computing power and resources [152].

We are using a simple displacement strategy to build the map. Starting from one corner, the robot follows the enclosing wall until it comes back to its initial position. The very same strategy was used by Robot 7 in Chap. 2, as it can be seen in Fig. 2.5g (p. 17). Inside a static environment, this mapping phase can be done only once, and the result can be reused many times by the following phases. However, this hypothesis might not be true in reality. This is the reason why we will always count this phase inside the energy balance.

Phase 2: Systematic coverage The map built by the previous phase is leveraged to perform the robot’s mission. In our case, we will perform a systematic coverage of the surface, as performed by a robotic vacuum cleaner or a floor scrubbing robot. This gives an upper bound on the required energy, as robots performing other tasks have less stringent requirements, for example, during a patrolling mission.

This is achieved by combining a localisation device with a higher level planner. The localisation is performed by using ROS with its Adaptive Monte Carlo Localisation (ACML) algorithm [153]. On top of this, we have made a simplistic planning strategy in order to move the robot following a boustrophedon path.

⁵Currently maintained by the Open Source Robotics Foundation. More information at <http://www.ros.org/>.




Figure 6.12 – Low-level state machine for the third phase of the experiment. The states in red require at least one active actuator and more power as a consequence. State transitions with a dashed arrow are not implemented in our experiment but should be planned in reality.

During this phase, the actual power delivered by the photovoltaic panel is recorded. It is combined with the current location to build a map of the available solar power inside the room. The displacement of the robot is expected to be fast enough, compared to the relative motion of the sun, to minimise the distortions on the actual beam's location. An important detail for this mapping is the inclination of the photovoltaic panel: It must be flat in order to be symmetric with respect to the robot's rotation axis. Failing to do so would result in an anisotropic mapping of the available photovoltaic power, depending on the robot's heading.

Phase 3: Sun tracking and recharge This phase is where the robot will actually recharge itself. The goal is thus to minimise the energy expenses and to maximise the solar yield at the same time. These two objectives are contradictory, but both are necessary. In a mobile robot, most of the energy is spent by two items: the mobility and the high-level control. The mobility is mandatory to stay aligned with the light along the day, and a good high-level control can minimise the failures to localise and track the best charging places. An effective trade-off must be found.

We decided to use a mix of reflective and reactive control. The reflective control is used at the beginning of the phase to efficiently reach the beam's location independently from the current robot's location. This step uses the light map built during phase 2, and effectively minimises the kinematic energy needed to begin the recharge process. The high-level control, namely the CPU and the laser scanner of the robot in our case, can then be switched off. Indeed, the rest of the process is performed using a purely reactive approach, only driven by the microcontroller of the solar harvesting module. This effectively minimises the energy required by the control logic.

The state machine used by the microcontroller during the recharge procedure is shown in Fig. 6.12. When entering the state machine for the first time, an initialisation step is carried out to turn off all unused resources. The following steps are used to optimise the placement of the photovoltaic panel with respect to the geometry of the beam. A

gradient ascent is first performed by the mobile base to find the brightest spot. For this purpose, the four photodiodes on the corner of the panel are used to compute the displacement vector. The panel is then inclined to a predetermined angle (36° in our case, but this is unimportant) to break the symmetry of the panel with respect to the robot's rotation axis. By doing so, the robot can spin on itself to find the rotation angle with the highest outcome. The inclination of the panel is finally fine-tuned by performing a gradient ascent on this axis.

The last optimisation step is important. The robot is moved forward, until the harvested power drops below a given threshold (80% of the maximum power in our case). This enables the robot to reach the forward edge of the moving sun spot. It is expected that the robot can stay longer inside the beam with such a strategy, consequently minimising the number of costly realignments that are necessary to move the robot back inside the beam.

Once the photovoltaic is correctly placed, the robot can enter a deep-sleep state. Motors are put into hibernation, and the main microcontroller of the solar harvesting module enters power-saving mode. The microcontroller periodically leaves this hibernation state, once every few minutes, to decide if the robot needs to move again. This decision is based on the actual power on the photovoltaic panel in our simple scenario. But this strategy may not be robust enough in reality, for example, to cope with periodic clouds. In this case, other criteria could be used in conjunction, such as the expected displacement speed of the sun spot or the daily weather forecast. An external sensor could also be placed onto the external wall. If the robot is still correctly placed, the microcontroller enters back into its hibernation; otherwise, the state of the robot is restored, the inclination of the panel is zeroed, and the state machine begins a new realignment.

This recharge phase ends when a stop condition is met. This can be a trigger on the duration, the recharge level of the battery, or some kind of user input (when a button is pressed, or when the robot is kicked). This user interface is also important in reality to wake up the robot when it is in hibernation state, as the user may want the robot to move out of his way. This behaviour is not implemented in the present work, however.

6.4 Experimental Results

This section will analyse the results of the experiment conducted inside the solar arena, using the marXbot robot equipped with the laser scanner, the solar harvesting module, and the four-channel power logger. The four channels were placed as follow: 1) on the battery; 2) at the exit of the mobile base; 3) on the way to the photovoltaic harvester; and 4) on the way to the CPU. The power consumed or provided by most of the subsystems can then be computed by performing some simple arithmetic. Each phase of our scenario was performed five times. We will present the results phase by phase.




Figure 6.13 – Qualitative results for the first phase (map building).


6.4.1 Phase 1: Map Building

During the first phase, the robot follows the wall and builds a map of its environment at the same time. The phase ends after a complete revolution around the room, and the map is stored for later use. The quality of the result depends on both the quality of the robot's odometry and the precision of the laser scanner. The precision of the laser scanner was already studied in Chap. 4. Figure 6.13a shows the resulting point cloud and odometry, when the SLAM algorithm is disabled. On the other side, Fig. 6.13b shows the map produced by the SLAM mapping, as well as the estimated robot's trajectory.


The quality of the (calibrated) odometry on the marXbot is good enough, even if it has wheels. The dominant error is on the self-rotation because the contact point between the tracks and the ground is loosely defined and can even vary over time. The SLAM algorithm successfully copes with such errors and produces a usable estimate of the robot's trajectory.

Figure 6.14a shows an example of the recorded instantaneous power for the different subsystems during the course of the mapping phase. The data were post-processed with an averaging window of 0.2 s to remove the high-frequency harmonics and ease the visualisation. The average filter was designed to conserve the energy and the average power. The sum of the power consumed by the four subsystems is the power drawn from the battery. The numerical integration of the curves gives the energy used by each subsystem during the experiment. The


Chapter 6. Experimental Validation of an Indoor Harvesting Robot




(a) Power consumed by the subsystems during the fourth run. Data are averaged on a window of 0.2 s to remove harmonics, but the average energy is conserved. The power at the battery equals the sum of the four other subsystems.



(b) Box plot of the consumed energy per subsystem, based on the numeric integration of the power ($N = 5$).



(c) Box plot of the completion time for the first phase ($N = 5$).



(d) Box plot of the power per subsystem, averaged on the whole duration of the run ($N = 5$).

Figure 6.14 – Quantitative results for the first phase (map building).

resulting statistics are shown in Fig. 6.14b, while the box plot for the completion time is in Fig. 6.14c. The median of the completion time is at 173 s, just below 3 minutes.

The average power of Fig. 6.14d is found by dividing the energy by the completion time. The median of the total power drawn from the battery is 5.6 W. This value is broken down as follows between the subsystems: the mobile base consumes the most with 2.4 W (43%), followed by the processor board with 1.9 W (34%), and finally the rest of the robot with 1.3 W (23%). Inside this last group, we find the laser scanner. From Sec. 4.4, we know that the Neato scanner consumes on average 1.1 W (20%). The remaining 0.2 W are mainly consumed by the main microcontroller of the solar harvesting module. The harvested solar power is null, as the robot never enters the sun spot and the ambient lighting is negligible.


The mapping capability, corresponding to the laser sensor and part of the CPU, is the major power consumer with an average of 3 W (54%). From the pre-run data, we can find that the power consumption of the motionless base is 0.78 W. Hence, the motors come second with $2.4 - 0.78 = 1.6$ W (29%). The remaining consumers, led by the numerous microcontrollers disseminated inside the marXbot, represent just below 1 W of power (17%). This is the price to pay for the great modularity of this robot; the functionalities are physically split into different modules, where a single microcontroller could suffice inside a monolithic design. The CAN bus used to link the controllers between them is also a major consumer.

We can note on Fig. 6.14a that the instantaneous power of the mobile base depends on the velocity. When moving forward at around 10 cm/s, the power of the mobile base is around 2.5 W. When turning on spot at lower speed, the power decreases to 1.9 W. The relation between the displacement speed of the marXbot and the power of the mobile base was already studied by Bonani et al. in [45].

6.4.2 Phase 2: Systematic Coverage

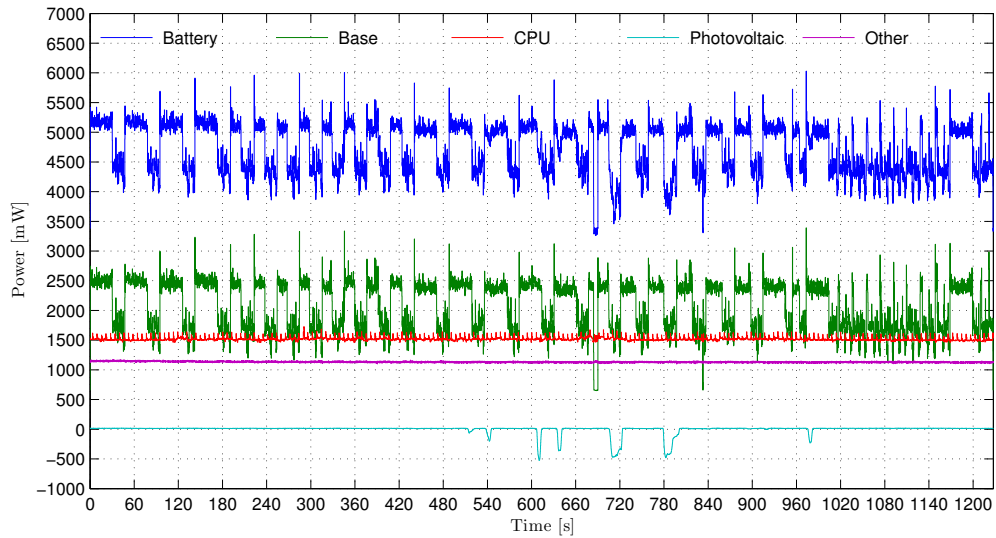
During the second phase, the robot systematically navigates inside the environment by relying on the map produced during the first phase and combining with a Monte Carlo localisation. Figure 6.15 shows the accomplishment of two distinct runs. For the first one (Fig. 6.15a), the sun is just before 10:00am. While for the second one (Fig. 6.15b), the sun is around solar noon. For each run, we show a picture from the overhead camera (top), the trajectory estimated based on the Monte Carlo localisation (middle), and the resulting light map generated by the robot. The trajectory during both runs is very similar, which shows the repeatability of the Monte Carlo localisation.

The map of the solar power inside the room is interesting in many regards. First of all, the location of the sun spot is clearly visible and detaches from the background. The shape of the spot on the map matches well the shape of the beam as seen from the overhead camera. The possible distortions due to the motion of the sun are kept low by the comparable high speed of the robot. Outside the solar spot, the background power level is really low, as expected.




(a) Top view, estimated trajectory, and measured photovoltaic power when the sun is just before 10:00am. (b) Top view, trajectory and measured photovoltaic power when the sun is around solar noon.


Figure 6.15 – Qualitative results for the second phase (systematic coverage).




(a) Power consumed by the subsystems during the third run. Data are averaged on a window of 0.2 s to remove harmonics, but the average energy is conserved. The power at the battery equals the sum of the four other subsystems.



(b) Box plot of the consumed energy per subsystem, based on the numeric integration of the power ($N = 5$).



(c) Box plot of the completion time for the second phase ($N = 5$).



(d) Box plot of the power per subsystem, averaged on the whole duration of the run ($N = 5$).

Figure 6.16 – Quantitative results for the second phase (systematic coverage).

We can note a few differences. If we compare the location of the spot on the map with the reality, there is a spatial offset. This is especially well seen in Fig. 6.15a, if we take the dust bin as a reference point. This is easy to explain though. The power is measured on the photovoltaic panel, whose plane is shifted in height from the ground by about $h_{PV} = 34$ cm. But the location on the map is reported with respect to the mobile base. Due to the incident angle of the beam, when the photovoltaic panel is inside the light, the base is already outside. The horizontal difference is thus $\|\Delta\vec{x}\| = h_{PV} / \tan(\alpha)$. The edges of the spot are also not as sharp as in reality (refer to Fig. 6.5b on p. 101 for a comparison with the measurement realised using a photometer). This is due to the physical dimension of the photovoltaic panel, which acts as a physically moving average filter.

Figure 6.16a shows an example of the recorded instantaneous power for the different subsystems during the course of the systematic coverage phase. Similarly to the mapping phase, the data were post-processed with an averaging window of 0.2 s to remove the high-frequency harmonics and ease the visualisation. The photovoltaic curve clearly shows peaks of produced energy (negative values) when the robot steps inside the sun spot. This is not enough to recharge the battery, however, as too much power is used by the rest of the robot.


The numerical integration of the instantaneous power curves gives the energy used by each subsystem during the experiment. The resulting statistics are shown in Fig. 6.16b, while the box plot for the completion time is in Fig. 6.16c. The median of the completion time is at 1229 s, just above 20 minutes. This is comparable to the coverage time of the SLAM-enabled robotic vacuum cleaners of Chap. 2 (see Fig. 2.6b on p. 19) for an arena of about the same size.

The energy required for this phase is higher than the mapping phase, but this is explained by the longer completion time. To prove this, the average power is computed in Fig. 6.16d. The figures are very similar to the mapping phase, and the same conclusions can be drawn regarding the power usage inside the robot. Even if the robot crosses the solar beam several times, this is still negligible compared to the energy required to move the robot around the room.


6.4.3 Phase 3: Sun Tracking and Recharge

During the last phase, the robot first goes to the brightest point recorded on the map built during the previous phase. Then it begins the optimisation of its placement, before switching off unnecessary resources and entering a deep sleep. The robot wakes up about every five minutes to check the amount of harvested power and possibly realign again with the beam.


An important step of this phase is the (re)alignment of the robot with the beam. We first measure the sensitivity of the photovoltaic power as a function of the misalignment of the robot with the beam. The robot needs to successfully align several degrees of freedom. Figure 6.17a shows the dependency of the power on the vertical rotation angle γ_p of the robot. The measure is performed when the robot is north-aligned with the aperture. The optimal angle is



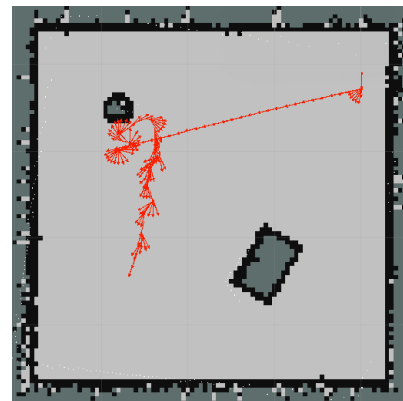
(a) Power at the output of the solar panel (in raw ADC units), as a function of the vertical orientation angle γ_p . The red dash-dotted line shows the theoretical optimum angle (180°).



(b) Power at the output of the solar panel (in raw ADC units), as a function of the inclination angle β . The red dash-dotted line shows the theoretical optimum angle (60°).



(c) Power at the output of the solar panel (in raw ADC units), as a function of the solar hour. The robot is motionless on the axis of the aperture. The black dashed line is the threshold at 80% of the maximum power.



(d) Example of a trajectory during the recharge phase, as estimated by the laser scans and the Monte Carlo localisation. The robot starts from the upper-right corner.

Figure 6.17 – Measurement of the photovoltaic power, with respect to the (mis)alignment of the panel with the beam. An example of the resulting trajectory is shown in (d).

Chapter 6. Experimental Validation of an Indoor Harvesting Robot


consequently when the robot faces the south. The power is relatively constant on the range $\gamma_p = 180 \pm 20^\circ$, and the location of the real maxima is hard to find due to some ripples on the power caused by the MPPT. During the experiments, our algorithm performed well, but a constant error between 5° and 15° was systematically observed. This has no real consequence, due to the low sensitivity of the power around the maxima.

The dependency of the power on the panel's inclination angle β is shown in Fig. 6.17b. Like with γ_p , the derivative of the power around the optimal angle (60° in this case) is close to zero. And for the same reasons, our algorithm repeatedly stops around 45° with only minor consequences on the produced power.


Finally, the dependency between the power and the robot's location inside the beam can be seen in Fig. 6.17c. The robot is kept motionless on the centre line, while the beam is moved at the sun's speed. The power is high, as long as the entire panel is inside the beam, for roughly 40 minutes between 11:40am and 12:20pm. When the panel reaches the edge of the spot, the power starts to rapidly decrease. We found that a threshold at 80% of the maximum power (black dashed line on the figure) is a reasonable choice to initiate a realignment procedure. Waiting longer does not bring any more useful power. With such a threshold, the elapsed time between two realignments should be around one hour. Figure 6.17d shows the example of a trajectory during the course of a complete 5-hour run. After the initial placement, five successive realignments were necessary in order to track the moving spot.

The power consumed by each subsystem during an entire run is shown in Fig. 6.18a, while the corresponding composite image taken with the overhead camera is depicted in Fig. 6.18b. After the initial placement (numbered 0), five successive realignments were necessary (numbered 1 to 5). The power of the battery and the solar harvester are negative between two alignments, demonstrating the recharge process. The smaller spikes at regular interval correspond to the wake-up states of the harvester's main microcontroller. It checks the current recharge power to see if it is necessary to move, which takes about 200 ms. During each realignment, the base consumes a significant amount of power but for a very short time at the scale of the experiment. Considering the losses due to the realignments, where is the balance of the recharge process? We will answer this question soon, but let us first zoom in on the important steps of the recharge process.

The power consumed by each subsystem during the initial placement is shown in Fig. 6.19a. The sequence ① corresponds to the displacement of the robot between its initial location and the estimated beam's location. The robot performs the gradient ascent ②, followed by the initial inclination of the solar panel to a predetermined angle ③. During this step, the power coming out of the photovoltaic harvester is reduced by about 700 mW. This is the power consumed by the stepper motor to actuate the inclination axis. The robot spins on itself during step ④ to align the panel with the beam, and the final inclination is adjusted during step ⑤. Finally, the robot moves toward the edge of the beam ⑥ and shuts down unnecessary resources (CPU, etc.) ⑦.




(a) Overview of the power usage during the five hours of the second run.




(b) Fused images, as seen from the overhead camera, after the first placement and each successive robot's realignment during the second run. The location of the robot is shown by the red dots.


Figure 6.18 – Proceedings of the second run. The robot realigned itself five times (numbered 1 to 5), after the initial placement (0).



(a) Power usage during the initial placement of the second run.



(b) Power usage during the first recharge slot of the second run.



(c) Power usage during the first realignment procedure of the second run.

Figure 6.19 – Zoom on the power usage during key moments of the second run. Data are averaged on a window of 0.2 s.


The recharge phase (Fig. 6.19b) is quite calm. The main microcontroller wakes up at regular intervals, about every five minutes in our case, and checks if the harvested power is below 80% of the maximum power. The microcontroller stays awake for about 200 ms, before going back to sleep. This awakening consumes about 40 mJ, while between two awakenings the module harvests roughly 180 J. These awakenings have only a very small cost, at least when the module is fully exposed to the sun. During the period outlined by (I), the harvested power decreases. This is observed only during the first recharge period (see Fig. 6.18a). This is caused by the heating up of the panel, decreasing the efficiency of the photovoltaic conversion due to changes in the semiconductor's properties. At the end of the recharge phase (II), the harvested power decreases as the cell reaches the edge of the beam, and eventually the microcontroller decides to begin a realignment procedure.

The realignment procedure is detailed in Fig. 6.19c. This is similar in many regards to the procedure followed during the initial placement. The major difference is the first step (A), where the panel is inclined back to the horizontal to perform the anisotropic gradient ascent (B). The panel is then inclined again (C), followed by the robot's spin (D) and the fine-tuning of the inclination (E). The final step moves the robot toward the edge of the beam (F). Compared to the initial placement, the overall power is also lower because some resources, such as the processor, are not used. This realignment procedure is purely reactive, which consumes less energy compared to a reflective strategy.


The statistical assessment of the efficiency of the recharge phase is presented in Fig. 6.20. The first important metric is the overall balance, as seen in Fig. 6.20a. During the five-hour trials, the median of the harvested energy is 11000 J, while the battery recharged by an amount of 8000 J (2.2 Wh). Put the other way round, on the total of the harvested energy, 73% was effectively used to recharge the system. The rest was lost during the initial placement, the realignments steps, and the recharge cycles. The median of the average recharge power (Fig. 6.20b) reached 440 mW inside the battery, while at the same time the harvester delivered 610 mW. The losses inside the other part of the robot represented an average of 170 mW.

It is interesting to go into the details of the process. A single run can be broken up into a succession a three steps: 1) the initial placement, which happens only once at the beginning, 2) a recharge period, and 3) a realignment procedure. The robot alternates between the recharge and the realignment steps, until the end of each trial. The detailed statistics for each of these steps can be computed. During the initial placement (Fig. 6.20c), the mobile base, the processor, and the laser sensor are all part of the energy expenses. Some energy is harvested at the end of the run ($\tilde{E} = 43$ J), but far from enough to compensate for the energy going out of the battery ($\tilde{E} = 418$ J).


During a single realignment step (Fig. 6.20d), the situation is similar, but with a few important differences. Now, the base (thus the mobility of the robot) represents almost the totality of the energy budget. This is due to the reactive strategy, as no high-level control is involved. The energy required for a single realignment ($\tilde{E} = 63$ J) represents only 15% of the energy required




(a) Energy consumed by each subsystem during the totality of phase 3 ($N = 5$).




(b) Average power consumed by each subsystem during the totality of phase 3 ($N = 5$).



(c) Energy consumed by each subsystem during the initial placement step ($N = 5$).



(d) Energy consumed by each subsystem during a single realignment step ($N = 21$).



(e) Energy consumed by each subsystem during a single recharge slot ($N = 21$).

Figure 6.20 – Box plots of the energy-related metrics for each subsystem, when considering the totality ((a) and (b)) or parts ((c) to (e)) of the third phase (solar recharge).

for the initial placement. Even if we consider the multiple realignment steps required on a whole run ($\bar{E} = 286 \text{ J}$), the energy used is still 30% lower. The reactive approach proves to be energy efficient, while enabling a simple and successful tracking of the sun.

Finally, the recharge step (Fig. 6.20e) is the longest of all, but it is also where energy is the most efficiently harvested. The median time of a recharge, before proceeding to a realignment, is 3862 s (64 minutes and 22 seconds) on 21 single recharge steps performed during the 5 trials. This corresponds to the 3600 s that were predicted earlier when choosing the activation threshold. The median of the energy injected into the battery is 1865 J per recharge step, but there is a substantial variance. This variance is chiefly explained by the variance on the recharge duration between two realignments (see Fig. 6.18a on p. 119 as an example of the variability) and by some small random misalignments between the photovoltaic panel and the incoming beam.

6.5 Concluding Analysis


6.5.1 Summary of the Results

Figure 6.21 summarizes the energy balance among the three phases of our experiment when considering several important subsystems. If we take a look at the battery (Fig. 6.21a), phase 1 (map building) requires 1000 J, and phase 2 (systematic coverage) requires 5400 J. On the other hand, during the last phase, we are able to recharge the battery by almost 8000 J. So, at the end of the day, we have a net gain of about 1600 J (0.45 Wh). This is a small amount of energy, but it represents a good and valuable achievement. If we had such sunny conditions for every day of the year, the robot would be able to map the room early in the morning, perform a complete coverage of the room, and then spend the rest of the day tracking the sun, while still accumulating more and more energy.


If we focus on the sole mobility (Fig. 6.21b), the most intensive task is the systematic coverage of the floor. The recharge phase comes second, due to the frequent realignments required to track the position of the beam. Finally, when it comes to the harvested energy (Fig. 6.21c), only the third phase gives any consequent amount of energy. During the two other phases, the occasional entry of the robot inside the solar beam is not enough to bring us any useful energy.

Earlier, we raised the question regarding the solar tracking on the horizontal axis of the photovoltaic panel. Is the gain worth the energy lost by the motor? In our scenario, the inclination angle β of the panel is not continuously adjusted, which is contrary to what is done in some outdoor photovoltaic plants. The adjustment is performed once during each realignment procedure, in which case, we are forced to put the panel back to the horizontal first in order to start the gradient ascent⁶. We can estimate the energy lost by this back-and-


⁶ For the final design, the photodiodes should not be fixed to the mobile panel to avoid this back-and-forth.



(a) Energy consumed by the battery.



(b) Energy consumed by the mobile base.



(c) Energy consumed by the solar harvester.

Figure 6.21 – Summary of the energy consumed by several subsystems, during the three phases of the experiment.

forth movement. We already concluded in Sec. 6.4.3 that our stepper motor was using 700 mW of power. The total time for the inclination procedure can be estimated from Fig. 6.19c (p. 120) by summing up steps (A), (C), and (E). This equals to 13 s. Thus a single adjustment costs 9.1 J. On the whole experience, we performed up to five realignments, which gives a total cost of 46 J. Compared to the 11000 J that we harvested, this is a loss of at most 0.4%. On the other side, we predicted with our theoretical framework (Sec. 5.5) a possible gain of up to 3.2% with the addition of a tracking axis, compared to a fixed system. Thus our tracking system could be worth the gain.

From another point of view, the tracking axis adds a significant complexity to the system. This could be a disadvantage when considering the use of this technology inside a real product. The system will cost more, without mentioning the fragility and the possible breakage. A gain of a few percentage points might be not enough to be worth this complexity. If we push this reasoning further, a panel inclined at a fixed angle also has some drawbacks. One advantage of the robotic vacuum cleaners of Chap. 2 is their ability to go under the pieces of furniture. A bulky device fixed on their top would severely limit this ability. So what would happen with a fixed horizontal panel? From Fig. 6.19c (p. 120), we see a gain of about 200 mW when tilting

the panel (between the end of (B) and the end of (E)). A fixed horizontal panel would mean a loss of about 30%, or 3300 J less inside the battery at the end of the day. This is a big loss in our energy balance but a gain on the usefulness of the robot in everyday life. Both arguments should be weighted when designing a real robot, but a horizontal panel will probably win in most cases, at the expense of a greater charging time.

6.5.2 From Simulation to Experiment

We can compare our experimental results with the results of the simulation performed in Sec. 5.4.2, and compute the global efficiency when moving from the simulation to the reality. We set our solar arena to reproduce the conditions experienced during the 14th of February ($d = 45$). If we refer to the simulated results for case 7 (Fig. 5.9c on p. 88), the daily energy income reaches $6367 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$ when $d = 45$, for a total of 6 hours and 21 minutes of sunshine. The energy received by our photovoltaic harvester is deduced by scaling down this value to the 5 hours of our experiment, taking into account the area of the panel (190 cm^2) and correcting for the lower power density of our light bulb (515 W/m^2 instead of 1000 W/m^2). At the end, the predicted amount of energy on the panel is 49 Wh. During the third phase of our experiment, we harvested an amount of 11000 J (Fig. 6.20a), or 3.1 Wh. The overall efficiency of the harvester, based on this predicted illumination, equals 6.3%. This is very close to the experimental efficiency that we characterised in Fig. 6.10c on p. 107 (7.8% at 354 W/m^2).

This demonstrates two major points. First, there is a good match between the theoretical framework developed in Chap. 5 and the experimental results conducted inside our solar arena. This validates the experimental conditions with respect to the solar mechanics defined in Sec. 5.3. This also reveals the good performance of our embedded control to track the solar spot. The overall efficiency is limited by the efficiency of the harvester and not by a suboptimal tracking. Thus, future efforts to increase the efficiency should focus in priority on the harvester.

6.5.3 From Experiment to Reality

Up to now, we have considered only ideal sunshine conditions. This hypothesis was used all along the theoretic framework of Chap. 5, as well as during the experiments conducted in the present chapter. How big is the step between these idealized conditions and reality? Figure 6.22 shows the statistical yearly relative sunshine duration in Switzerland for two distinctive years. During year 2010 (Fig. 6.22a), the amount of sunshine was particularly low relative to the past decade. On the contrary, year 2011 (Fig. 6.22b) was especially good. This gives us the two possible extremes when considering the statistical data of the past decade. For the worst, some places saw only 30% of sunshine in 2010. For the best, this values can rise up to 65% during a good year. If we transpose this to the experimental results of phase 3, the realistic harvested energy could be comprised between 3300 and 7150 J (instead of 11000 J). This linearly reduces the energy injected into the battery, as the losses inside the rest of the system remains. At the




Figure 6.22 – Comparison of the statistical yearly relative sunshine duration in Switzerland (in % of the theoretic maximum), when considering the worst and the best year of the past decade. Year 2010 saw a low amount of sunshine, while 2011 saw a high amount. Sunshine is taken into account when the direct solar irradiance exceeds 200 W/m^2 . Picture credits: ©MeteoSwiss.




Figure 6.23 – Comparison of the statistical monthly relative sunshine duration in Switzerland during February for two distinctive years. Picture credits: ©MeteoSwiss.

end, the battery could be recharged by 300 to 4150 J, respectively (instead of 8000 J). While a few thousands Joules are still interesting, we are on a sharp edge.

If we consider the statistical monthly relative sunshine duration for February (Fig. 6.23), the situation can be even worse. During February 2010 (Fig. 6.23a), some places saw less than 20% of relative sunshine for the whole month. Under such conditions, the equilibrium is broken, and we can no longer recharge the battery. On the contrary, some places saw more than 70% of relative sunshine during February 2012 (Fig. 6.23b). Our harvesting process is thus highly unpredictable, with respect to the variability of the weather from one year to the other. Our recharge process cannot be the only source of energy. It is still necessary to use a recharge station connected to a wall socket. Otherwise, our robot could potentially starve for complete months when experiencing harsh conditions. A large battery could smooth out




Figure 6.24 – Simulated daily energy gained by the solar robot for two distinctive cases of yearly amount of sunshine. The dashed line is our energy needs (20 kJ) for one cleaning task. If the energy gained by the battery is below this target, the robot will have to perform several consecutive days of recharge before being able to clean again.

these statistical variations, but with current battery technologies it is impossible to store the energy needed for multiple consecutive tasks in such a limited volume. As a consequence, energy shortages are still possible.

6.5.4 A Concrete Case

What happens if we transpose these results to a real service robot? If we adapt our harvesting technology to robotic vacuum cleaners, for example, how many days of recharge would be necessary after completing a cleaning cycle? Let us first estimate the needs of such a robot. If we go back to the robotic vacuum cleaners of Chap. 2, and particularly Fig. 2.7 (p. 20), we measured the specific energy used by these robots. Robots with a navigation technology had a specific energy between 500 and 2000 J/m². Scaled to the area of our arena (16 m²), this means a need in energy between 8 and 32 kJ. Let us settle on the average value: 20 kJ per cleaning task. As a point of comparison, the specific energy of the marXbot robot during phase 2 was 340 J/m². This lower value is coherent for two reasons. First, the marXbot has no cleaning functionalities, which means no additional motors. Then, we used high-end components with a higher efficiency, compared to the low-cost components used on commercial products, most notably for the motors.

Now, we have to compute the achievable recharge of energy. We derive this information from the results of case 7 (Fig. 5.9c on p. 88) and by applying a number of constraints. We take the same photovoltaic panel (190 cm², $\eta = 6.3\%$) with the same algorithms under standard conditions ($G = 1000 \text{ W/m}^2$) and with one hour of sunshine less per day (this hour is used for

cleaning). The unknown here is the relative amount of sunshine per day. Figure 6.24 plots the harvested energy and the energy used to recharge the battery for two different values of sunshine. Figure 6.24a uses a pessimistic value of 30% (but keep in mind that worse is occasionally possible), and Fig. 6.24b uses an optimistic value of 65%. On both figures, the bold dashed line shows our need of 20 kJ.

During a good year (Fig. 6.24b), we can have up to 22 kJ of recharge per day. We would be able to use our robot every day, for about 150 days in the year. In winter, the energy would be lower but still enough to use our robot once every two or three days. However, during bad years (Fig. 6.24a), the situation is less perfect. With a relative sunshine of only 30%, the recharge energy reaches its limit around 8.7 kJ per day in summer and goes down to 1.8 kJ in winter. We can thus use the robot once every 2.3 days in summer and only once every 11 days in winter.

Finally, the plots of Fig. 6.24 can be used to compute the economical worthiness of our approach. How much money do we save with this harvester, compared to a recharge station connected to the wall socket? Let us take a price for the electricity of 0.13 \$/kWh⁷. A yearly relative sunshine of 30% would result in a total recharge of 570 Wh on the whole year, or \$0.07 worth of electricity at the current price. A good year, with 65% of yearly relative sunshine, would give 1.6 kWh of energy, or save us \$0.21. At an industrial scale, this harvester would probably cost at least \$5. For the best case, it would take us 23 years to make our investment profitable or more than 70 years if we are out of luck. From an economic point of view, this project is not worth the investment due to several factors: we have a lower amount of sunshine compared to outdoors, the maximum size of the panel is small, and the actual electricity price is low.

6.6 Discussion

In this chapter, we presented the requirements and the experimental validation of the hardware used to challenge the results of Chap. 5. The final experiment was divided into three phases to duplicate the usage pattern of a robotic vacuum cleaner. This application was chosen due to its high requirement in energy and its broad adoption by people. The results for each phase showed the energy used by each major subsystem and how simple algorithms were able to achieve an autonomous harvesting of the solar power indoors.

Most importantly, our final analysis showed how the results of our theoretical framework closely matches the experimental results and how these results can be transferred to a real robotic vacuum cleaner. We cannot completely drop the recharge station, due to the daily variability of the weather conditions. But years with a high sunshine ratio can bring us an interesting amount of energy for such applications.

⁷ Source: U.S. Energy Information Administration (EIA) for residential consumers (August 2014): <http://www.eia.gov/electricity/>.

6.7. People Who Contributed to This Work

The configuration of the building will play a key role in the feasibility of such an approach. An aperture oriented towards the equator is an absolute requirement, and it will not always be satisfied. It should also be free of any obstacle that may cast a shadow. We have a higher probability of satisfying such conditions inside modern buildings, especially if they follow a passive design. Such designs try to maximize the passive solar gain, especially in winter, which implies large south-facing windows. The location of the building is equally important. A residential house is a better candidate compared to a densely packed building bloc. In the latter case, a photovoltaic installation on the roof is the best option.

Another essential aspect of our approach is the adoption by users. The adoption of a regular robotic vacuum cleaner inside the home ecosystem is not always ensured [12, 154]. With our approach, a robot will have to spend most of its time by the window, where most of the natural light is. The first problem is the potential presence of pieces of furniture, such as a dining-room table. People are like solar robots, they try to make a good use of the outside lighting. The robot can walk around obstacles, but this will increase the energy needs and probably decrease the amount of daily sunshine.

The second problem is the higher footprint of the robot inside the home ecosystem. The robot may hinder the habits of people inside the building and, consequently, decrease the adoption ratio of this technology by people. More study is needed to understand the real impact and the possible issues. Unfortunately, it seems that little can be done to minimise this inconvenience from a technical point of view.

6.7 People Who Contributed to This Work

The experimental solar simulator (Sec. 6.1) was designed and assembled by Alain Vuille, with the help of Norbert Crot and under my supervision. The electronics, the mechanics, and the MPPT controller of the solar harvester (Sec. 6.2) were done by Martin Liniger during his Master thesis, under the supervision of Philippe Rétornaz and myself. The firmware for the 16-bit microcontroller was done by Philippe Rétornaz.

7 Conclusion

Energy resources at our disposal are abundantly used by an increasing number of electrical appliances inside our households. Sales of domestic robots are likewise increasing from one year to the next. Domestic robots offer a number of appreciated services to working single people and busy families but at the cost of increased stress on our limited resources. Energy is generally not considered as a design parameter by manufacturers, unless they are constrained by legal means. The robotics community is similarly little concerned by this topic and consequently few researches are done.

At the beginning of this work, we set two major goals. Our first goal was to find metrics and technologies to design more energy-efficient domestic robots. By studying actual robotic vacuum cleaners in Chap. 2, we were able to prove that the addition of a localisation technology is beneficial for the overall energy balance. We introduced the concept of *specific energy* to measure the energy efficiency of robots performing a complete coverage of the floor. With the addition of a localisation device, the *specific energy* can be lowered by 2 to 12 times compared to robots relying on a pseudo-random navigation. At the end, a robotic vacuum cleaner can even consume less energy compared to a hand-operated one. On the other hand, due to the lower cleaning efficiency of robots on some surfaces, most people will probably use both devices in conjunction but reduce the amount of manual vacuuming in counterpart.

The choice of the mapping sensor is also crucial, as noted in Chap. 4. The performance of scanner sensors can be linked to the energy with simple metrics, such as the *energy per bit* of information. These metrics can be useful when selecting a sensor for a specific application, but the designer should also take into account other metrics, including the start-up and warm-up energies. Even if the exact choice depends on the final scenario, this work provides the necessary toolbox.

Another major goal was to increase the ratio of renewable energies used by domestic robots. We evaluated a number of energy sources in Chap. 5. We also assessed the feasibility boundary with respect to a realistic scenario based on the vacuum cleaners of Chap. 2. Indoor environments offer few powerful supplies as most physical phenomena are attenuated by the

building envelope. One energy stands out, however: the sunshine that naturally comes inside through the apertures of the buildings with an intensity that is only slightly diminished by the windows. Within a south-oriented room, this can offer enough energy for most applications as demonstrated by the upper bound ($7651 \frac{\text{Wh}}{\text{m}^2 \cdot \text{day}}$) settled by numerical simulations. Based on seven case studies, we established the critical influence of the mobility of the base. A motionless object has no way to profit from the sun, whereas a mobile robot can take full advantage of this source due to the changing spatial locality as the day progresses and around the year. It was established that active tracking on the panel's inclination has less impact on the overall balance, but it can increase the energy income by a few percentage points.

The experimental validation of Chap. 6 was performed under controlled conditions inside the solar arena designed for this occasion. A solar harvesting module was developed for the marXbot robot and thoroughly characterised. The marXbot was programmed to accurately reproduce the behaviour of a robotic vacuum cleaner by using a localisation technology. The last phase was devoted to the solar recharge with a mapping-assisted initial placement. The successive realignment procedures used a low-level reactive behaviour to consume the least energy possible. The success of this approach was proved through a number of results:

1. At the end of the three-phase experiment, we have a net gain inside the battery of 1600 J, on average, under ideal sunny conditions.
2. Using the theoretic energy income, the closed-loop efficiency of the harvester is estimated to be 6.3%. This closely matches the experimental characterisation of the harvesting module (7.8%). It shows the good performance of our sun-tracking algorithm.
3. With a realistic scenario based on the robotic vacuum cleaners of Chap. 2 (20 kJ per cleaning task), we estimated the number of consecutive days needed to recharge the battery between two tasks. Depending on the statistical sunshine conditions and the day of year, a recharge could be fulfilled in 1 to 14 days.

This technology is deemed feasible, even if the recharge station cannot be completely abandoned due to sporadic weather conditions below the statistical conditions. The integration of this technology with a recharge station is yet to be studied. Acceptance by the user is also a key factor for the success of this robotic advance, as the robot will have to share the sunny places with the user's plants and other pieces of furniture. Finally, the configuration of the building is another crucial aspect that will inherently limit the utility of our technology. Ideal households have a single-family detached home without any outside obstacle on the south side. In the case of an apartment building inside a densely packed district, the best option remains photovoltaic panels on the rooftop.

7.1 Contributions

This work provides the following contributions to the state of the art:

- We introduced new energy-wise metrics to assess the impact of technologies on the energetic balance of robots. The *specific energy* metrics is useful when comparing robots that perform the complete coverage of an area, while the *energy per bit* of information metrics can be used for a variety of sensors, including mapping sensors.
- We proved the benefit of localisation technologies with respect to the energy required to accomplish a spatially bound task. The supplementary energy required for such devices is largely compensated for by the faster completion time.
- We provided a systematic assessment of the suitability of renewable energy sources applied to indoor robots. By comparing the feasibility boundary conditions, the sunshine coming through a building's aperture is shown to be a promising, yet challenging, energy source.
- We performed extensive simulations for several hypothetical solar harvesting systems and proved the major impact of the base's mobility on the daily energy income. Important design parameters were likewise studied.
- We demonstrated a successful implementation on a real robot evolving inside a controlled environment. A mixed strategy is proposed with an initial placement performed by a cognitive control, while the realignment steps as the day progresses are executed by a low-level reactive behaviour.
- Based on a comparison between the theoretical framework and the experimental validation, we provided a method to predict the feasibility of this innovative technology within the scope of a real application.

7.2 Outlook

This work logically focused on robotic vacuum cleaners, as they are by far the main application as measured by sales. It is straightforward to extend the findings of Chap. 2 to similar robots performing a complete coverage. However, more work is needed in the case of other applications, such as patrolling tasks. It would be interesting to broaden the scope of this work by analysing different appliances. As pointed out in Chap. 4, the proposed metrics do not account for the full complexity of sensors. The extension of these metrics could be likewise studied.

In Chap. 5, we focused on the most promising source of energy, but combinations can also be envisioned. For example, a decreased photovoltaic conversion's efficiency was noted in Sec. 6.4.3 due to the heating of the solar cell. Authors have already proposed to combine

Chapter 7. Conclusion

photovoltaic panels with thermoelectric generators to increase the system's efficiency [155, 156]. The implications of such hybrid technology should be studied within our theoretical framework.

A solar-enhanced service robot will have to compete for the sunny places inside the house. As outlined in Chap. 6, the adoption by the final user will be critical to the success of our technology. This should be studied, for example, by extending the adoption patterns found by [154]. Such a work can bring solutions to reduce the major causes of friction between the robot's behaviour and the user's habits. Before introducing such types of robots inside households, however, the prototype of Chap. 6 should first be converted into a real product. This implies considerations regarding the inclination of the panel (fixed or actuated), the useful user interactions, or the fine-tuning of the tracking algorithms. This is, however, mostly engineering work, and our study provides the necessary tools to evaluate the consequences of these choices on the final energy outcome.

A Four-Channel Power Datalogger

The hardware of the power datalogger was developed during this work¹. The principle characteristics are shown in Table A.1. This chapter presents the design requirements in Sec. A.1. The calibration procedure, as well as the achieved accuracy, are discussed in Sec. A.2.

A.1 Hardware

The module has the following design requirements:


- Four identical channels, each channel measuring both the voltage and the current.
- A high input impedance when measuring the voltage and a low input impedance when measuring the current.
- A measurement frequency of at least 1 kHz.
- Autonomous in energy. The module must not be connected to the supply of the measured system.
- All calculations should be performed on-board.

The simplified block diagram of the designed system is shown in Fig. A.1a. For each channel, the current is measured with a shunt resistor and amplified with a instrumentation amplifier. The voltage is measured through a buffering electronics to guarantee a high input impedance. Both the voltage and current are sampled by the Analog Digital Converter (ADC) of the microcontroller. Analog low-pass filters are used to avoid aliasing effects due to the sampling process.


The microcontroller communicates with the outside world through a Controller Area Network (CAN) bus. A Universal Asynchronous Receiver/Transmitter (UART) interface is also available.

¹ The measurement electronics is based on a previous design by Daniel Burnier. This electronics had only one channel, a lower voltage input impedance and a lower bandwidth limit.


Appendix A. Four-Channel Power Datalogger



(a) Simplified block diagram of the power datalogger.



(b) Assembled power datalogger.



(c) Hardware interface to connect the datalogger between modules of the marXbot.

Figure A.1 – Overview of the electronics.

It is possible to use Bluetooth and Universal Serial Bus (USB) interfaces with the help of dedicated communication modules developed by the Laboratoire de Systèmes Robotiques (LSRO) laboratory.

The final hardware is pictured in Fig. A.1b. The wires on the right side of the module are the inputs of the four channels. The power supply is performed separately, for example, with a Li-ion battery. The piece of hardware shown in Fig. A.1c is a small module dedicated to the easy connection of the datalogger between modules of the marXbot robot.

Table A.1 – Summary of the principle characteristics of the power datalogger.

	Typical	Unit
Sampling frequency	1.5	kHz
Voltage measure		
– Absolute maximum voltage	26	V
– Maximum voltage	6.5	V
– Input impedance	5	GΩ
– Resolution	1.6	mV
Current measure		
– Maximum current	11	A
– Maximum continuous current	5	A
– Input impedance	6	mΩ
– Resolution	2.7	mA

A.2 Calibration

For each channel, the voltage and current measured by the microcontroller (v_m, i_m) linearly depend on the input voltage and current (v_n, i_n). Based on the k -th sample, we can estimate the voltage and current (\hat{v}_n, \hat{i}_n) using

$$\hat{v}_n[k] = \text{gain}_v \cdot v_m[k] - \text{offset}_v \quad (\text{A.1})$$


$$\hat{i}_n[k] = \text{gain}_i \cdot i_m[k] - \text{offset}_i \quad (\text{A.2})$$

With some care, this calculation can be performed on the microcontroller in real-time. The theoretic gains in voltage and current equal 1.6 mV and 2.7 mA, respectively. However, it is necessary to proceed to a calibration in order to increase the precision of the conversion. Points were collected inside the working range, the true value was measured with a precision multimeter, and the resulting error was minimised by performing a least square optimisation.


The error on the estimated current $\Delta i[k] = \hat{i}_n[k] - I_n$ is shown in Fig. A.2a in the case of channel 1. The error is bounded in any case between -8 and +8 mA on all channels. If we consider the median of the samples, the error is even bounded between -4 and +4 mA.

The error on the estimated voltage $\Delta v[k] = \hat{v}_n[k] - V_n$ is shown in Fig. A.2b in the case of channel 1. Unlike the current, the measure of the voltage is not perfectly linear and a small quadratic component can be observed. But we know that this datalogger will be used only with electronics powered by a Li-ion battery. Consequently, we can limit the range of the measure between 3.5 and 4.2 V. We optimised the offset around 3.9 V. The relative error is shown in Fig. A.2c. The error is bounded between +0.2% and -0.2% inside this working range.


Appendix A. Four-Channel Power Datalogger



(a) Absolute error on the current when considering the full range (0 – 3 A).



(b) Absolute error on the voltage when considering the full range (0 – 4.2 V).



(c) Relative error on the voltage when zooming in the working range (3.5 – 4.2 V) and after the adjustment of the offset around the working point (3.9 V).

Figure A.2 – Box plots showing the accuracy on the voltage and current after calibration for channel 1 ($N = 15000$ for each bin). The other channels have a similar response.

Bibliography

- [1] M. Liniger, “Conception of a Solar, Indoor, Autonomous Mobile Robot,” Master’s thesis, Ecole Polytechnique Fédérale de Lausanne (EPFL), 2012.
- [2] E. Mandfield, “The diffusion of industrial robots in Japan and the United States,” *Research Policy*, vol. 18, no. 4, pp. 183–192, 1989.
- [3] IFR Statistical Department, *World Robotics 2014 Industrial Robots*. International Federation of Robotics, 2014.
- [4] R. K. Miller, “Robots in Industry,” in *Industrial Robot Handbook*, pp. 1–9, Springer, 1989.
- [5] IFR Statistical Department, *World Robotics 2014 Service Robots*. International Federation of Robotics, 2014.
- [6] J. Fink, V. Bauwens, F. Kaplan, and P. Dillenbourg, “Living with a Vacuum Cleaning Robot - A 6-month Ethnographic Study,” *I. J. Social Robotics*, vol. 5, no. 3, pp. 389–408, 2013.
- [7] IFR Statistical Department, *World Robotics 2010 Service Robots*. International Federation of Robotics, 2010.
- [8] IFR Statistical Department, *World Robotics 2011 Service Robots*. International Federation of Robotics, 2011.
- [9] IFR Statistical Department, *World Robotics 2012 Service Robots*. International Federation of Robotics, 2012.
- [10] IFR Statistical Department, *World Robotics 2013 Service Robots*. International Federation of Robotics, 2013.
- [11] “iRobot: Our History.” <http://www.irobot.com/About-iRobot/Company-Information/History>, 2014.
- [12] F. Vaussard, J. Fink, V. Bauwens, P. Rétoznaz, D. Hamel, P. Dillenbourg, and F. Mondada, “Lessons learned from robotic vacuum cleaners entering the home ecosystem,” *Robotics and Autonomous Systems*, vol. 62, no. 3, pp. 376–391, 2014.

Bibliography

- [13] F. Vaussard, P. Réturnaz, D. Hamel, and F. Mondada, “Cutting Down the Energy Consumed by Domestic Robots: Insights from Robotic Vacuum Cleaners.,” in *TAROS* (G. Herrmann, M. Studley, M. J. Pearson, A. T. Conn, C. Melhuish, M. Witkowski, J.-H. Kim, and P. Vadakkepat, eds.), vol. 7429 of *Lecture Notes in Computer Science*, pp. 128–139, Springer, 2012.
- [14] F. Yasutomi, M. Yamada, and K. Tsukamoto, “Cleaning robot control,” in *1988 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1839–1841, IEEE, 1988.
- [15] E. Prassler, A. Ritter, C. Schaeffer, and P. Fiorini, “A Short History of Cleaning Robots.,” *Auton. Robots*, vol. 9, no. 3, pp. 211–226, 2000.
- [16] H. Choset and P. Pignon, *Coverage Path Planning: The Boustrophedon Cellular Decomposition*. Springer, 1998.
- [17] N. L. Doh, C. Kim, and W. K. Chung, “A Practical Path Planner for the Robotic Vacuum Cleaner in Rectilinear Environments.,” *IEEE Trans. Consumer Electronics*, vol. 53, no. 2, pp. 519–527, 2007.
- [18] S. C. Wong and B. A. MacDonald, “A topological coverage algorithm for mobile robots,” in *2003 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, vol. 2, pp. 1685–1690, IEEE, 2003.
- [19] M. Cirillo, L. Karlsson, and A. Saffiotti, “Human-aware Task Planning: An Application to Mobile Robots,” *ACM Trans. Intell. Syst. Technol.*, vol. 1, pp. 15:1–15:26, December 2010.
- [20] M.-C. Chiu, L.-J. Yeh, and Y. Lin, “The design and application of a robotic vacuum cleaner,” *Journal of Information and Optimization Sciences*, vol. 30, no. 1, pp. 39–62, 2009.
- [21] T. Seifried, C. Rendl, F. Perteneder, J. Leitner, M. Haller, D. Sakamoto, J. Kato, M. Inami, and S. D. Scott, “CRISTAL, Control of Remotely Interfaced Systems Using Touch-based Actions in Living Spaces,” in *ACM SIGGRAPH 2009 Emerging Technologies*, (New York, NY, USA), ACM, 2009.
- [22] J. Forlizzi, “How robotic products become social products: an ethnographic study of cleaning in the home.,” in *HRI* (C. Breazeal, A. C. Schultz, T. Fong, and S. B. Kiesler, eds.), pp. 129–136, ACM, 2007.
- [23] J. Forlizzi and C. F. DiSalvo, “Service robots in the domestic environment: a study of the roomba vacuum in the home.,” in *HRI* (M. A. Goodrich, A. C. Schultz, and D. J. Bruemmer, eds.), pp. 258–265, ACM, 2006.
- [24] P. Fiorini and E. Prassler, “Cleaning and Household Robots: A Technology Survey.,” *Auton. Robots*, vol. 9, no. 3, pp. 227–235, 2000.

-
- [25] E. Prassler and K. Kosuge, "Domestic Robotics.," in *Springer Handbook of Robotics* (B. Siciliano and O. Khatib, eds.), pp. 1253–1281, Springer, 2008.
- [26] I. Alonso, "Service Robotics," in *Service Robotics within the Digital Home*, vol. 53 of *Intelligent Systems, Control and Automation: Science and Engineering*, p. 89–114, Springer Netherlands, 2011.
- [27] T. Breuer, G. R. G. Macedo, R. Hartanto, N. Hochgeschwender, D. Holz, F. Hegger, Z. Jin, C. Müller, J. Paulus, M. Reckhaus, J. A. Á. Ruiz, P.-G. Plöger, and G. K. Kraetzschmar, "Johnny: An Autonomous Service Robot for Domestic Environments.," *Journal of Intelligent and Robotic Systems*, vol. 66, no. 1-2, pp. 245–272, 2012.
- [28] I. Carrera, H. A. Moreno, R. J. Salterén, C. Pérez, L. Puglisi, and C. E. G. Cena, "ROAD: domestic assistant and rehabilitation robot.," *Med. Biol. Engineering and Computing*, vol. 49, no. 10, pp. 1201–1211, 2011.
- [29] F. Yuan, L. Twardon, and M. Hanheide, "Dynamic path planning adopting human navigation strategies for a domestic mobile robot.," in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3275–3281, IEEE, 2010.
- [30] H. Durrant-Whyte and T. Bailey, "Simultaneous localization and mapping: part I," *Robotics Automation Magazine, IEEE*, vol. 13, pp. 99–110, June 2006.
- [31] J. Leonard, "Challenges for Autonomous Mobile Robots," in *International Machine Vision and Image Processing Conference (IMVIP 2007)*, p. 4, September 2007.
- [32] S. Thrun, "Simultaneous Localization and Mapping," in *Robotics and Cognitive Approaches to Spatial Mapping* (M. Jefferies and W.-K. Yeap, eds.), vol. 38 of *Springer Tracts in Advanced Robotics*, pp. 13–41, Springer Berlin Heidelberg, 2008.
- [33] T. Palleja, M. Tresanchez, M. Teixido, and J. Palacín, "Modeling floor-cleaning coverage performances of some domestic mobile robots in a reduced scenario.," *Robotics and Autonomous Systems*, vol. 58, no. 1, pp. 37–45, 2010.
- [34] A. T. de Almeida and J. Fong, "Domestic Service Robots," *Robotics Automation Magazine, IEEE*, vol. 18, pp. 18–20, September 2011.
- [35] H. Seki, K. Ishihara, and S. Tadakuma, "Novel Regenerative Braking Control of Electric Power-Assisted Wheelchair for Safety Downhill Road Driving," *IEEE Transactions on Industrial Electronics*, vol. 56, pp. 1393–1400, May 2009.
- [36] K. Konolige, J. Augenbraun, N. Donaldson, C. Fiebig, and P. Shah, "A low-cost laser distance sensor," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3002–3008, May 2008.
- [37] W. Jeong and K. Lee, "CV-SLAM: A new ceiling vision-based SLAM technique," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3195–3200, IEEE, 2005.

Bibliography

- [38] A. Piebalgs, “Commission Regulation (EC) No 1275/2008,” *Official Journal of the European Union*, 2008.
- [39] P. KaewTraKulPong and R. Bowden, “An improved adaptive background mixture model for real-time tracking with shadow detection,” in *Proc. 2nd European Workshop on Advanced Video Based Surveillance Systems*, vol. 25, pp. 1–5, 2001.
- [40] M. Dorigo, D. Floreano, L. Gambardella, F. Mondada, S. Nolfi, T. Baaboura, M. Birattari, M. Bonani, M. Brambilla, A. Brutschy, D. Burnier, A. Campo, A. Christensen, A. Decugniere, G. Di Caro, F. Ducatelle, E. Ferrante, A. Forster, J. Martinez Gonzales, J. Guzzi, V. Longchamp, S. Magnenat, N. Mathews, M. Montes de Oca, R. O’Grady, C. Pinciroli, G. Pini, P. Retornaz, J. Roberts, V. Sperati, T. Stirling, A. Stranieri, T. Stutzle, V. Trianni, E. Tuci, A. Turgut, and F. Vaussard, “Swarmanoid: A Novel Concept for the Study of Heterogeneous Robotic Swarms,” *Robotics & Automation Magazine, IEEE*, vol. 20, no. 4, pp. 60–71, 2013.
- [41] S. Magnenat, R. Philippsen, and F. Mondada, “Autonomous construction using scarce resources in unknown environments - Ingredients for an intelligent robotic interaction with the physical world.,” *Auton. Robots*, vol. 33, no. 4, pp. 467–485, 2012.
- [42] F. RoCHAT, P. Schoeneich, M. Bonani, S. Magnenat, F. Mondada, H. Bleuler, and H. Christoph, “Design of magnetic switchable device (MSD) and applications in climbing robot,” in *Emerging trends in mobile robotics*, pp. 375–382, World Scientific, 2010.
- [43] A. Gribovskiy, *Animal-Robot Interaction for Ethological Studies*. PhD thesis, Lausanne, 2011.
- [44] M. Bonani, *Robotique collective et auto-assemblage*. PhD thesis, Lausanne, 2010.
- [45] M. Bonani, V. Longchamp, S. Magnenat, P. Rétornaz, D. Burnier, G. Roulet, F. Vaussard, H. Bleuler, and F. Mondada, “The marXbot, a miniature mobile robot opening new perspectives for the collective-robotic research.,” in *2010 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, p. 4187–4193, IEEE, 2010.
- [46] ISO Standard, “ISO 11898, 1993,” *Road vehicles—interchange of digital information—Controller Area Network (CAN) for high-speed communication*, 1993.
- [47] Semiconductors, NXP, “UM10204,” *I2C-bus specification and user manual*, 2007.
- [48] S. Magnenat, P. Rétornaz, M. Bonani, V. Longchamp, and F. Mondada, “ASEBA: a modular architecture for event-based control of complex robots,” *IEEE/ASME Transactions on Mechatronics*, no. 99, p. 1–9, 2009.
- [49] F. Vaussard, P. Rétornaz, S. Roelofsen, M. Bonani, F. Rey, and F. Mondada, “Towards Long-Term Collective Experiments.,” in *IAS (2)* (S. Lee, H. Cho, K.-J. Yoon, and J. Lee, eds.), vol. 194 of *Advances in Intelligent Systems and Computing*, pp. 683–692, Springer, 2012.

- [50] S. Magnenat, V. Longchamp, M. Bonani, P. Rétornaz, P. Germano, H. Bleuler, and F. Mondada, "Affordable SLAM through the co-design of hardware and methodology," in *Proceedings of the 2010 IEEE International Conference on Robotics and Automation (ICRA)*, p. 5395–5401, IEEE, 2010.
- [51] M. Ferguson and N. Webb, "An intelligent low-cost scanning range finder," in *2011 IEEE Conference on Technologies for Practical Robot Applications (TePRA)*, pp. 168–172, IEEE, 2011.
- [52] P. Shah, K. Konolige, J. Augenbraun, N. Donaldson, C. Fiebig, Y. Liu, H. Khan, J. Pinzarrone, L. Salinas, H. Tang, *et al.*, "Distance sensor system and method," Feb. 4 2010. US Patent App. 11/780,017.
- [53] B. Sofman, V. Ermakov, M. Emmerich, S. Alexander, and N. Monson, "Method and apparatus for simultaneous localization and mapping of mobile robot environment," Feb. 27 2014. US Patent App. 14/067,705.
- [54] H. Kawata, A. Ohya, S. Yuta, W. Santosh, and T. Mori, "Development of ultra-small lightweight optical range sensor system," in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 1078–1083, August 2005.
- [55] F. Pomerleau, A. Breitenmoser, M. Liu, F. Colas, and R. Siegwart, "Noise characterization of depth sensors for surface inspections," in *2012 2nd International Conference on Applied Robotics for the Power Industry (CARPI)*, pp. 16–21, IEEE, 2012.
- [56] Y. Okubo, C. Ye, and J. Borenstein, "Characterization of the Hokuyo URG-04LX laser rangefinder for mobile robot obstacle negotiation," in *SPIE Defense, Security, and Sensing*, pp. 733212–733212, International Society for Optics and Photonics, 2009.
- [57] J. Pascoal, L. Marques, and A. de Almeida, "Assessment of Laser Range Finders in risky environments," in *2008 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3533–3538, IEEE, 2008.
- [58] L. Kneip, F. Tâche, G. Caprari, and R. Siegwart, "Characterization of the compact Hokuyo URG-04LX 2D laser range scanner," in *Proceedings of the 2009 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1447–1454, IEEE, 2009.
- [59] C. A. Rockey, E. M. Perko, and W. S. Newman, "An evaluation of low-cost sensors for smart wheelchairs.," in *2013 IEEE International Conference on Automation Science and Engineering (CASE)*, pp. 249–254, IEEE, 2013.
- [60] C. Ye and J. Borenstein, "Characterization of a 2-D Laser Scanner for Mobile Robot Obstacle Negotiation.," in *Proceedings of the 2002 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 2512–2518, IEEE, 2002.
- [61] C. Stachniss, G. Grisetti, W. Burgard, and N. Roy, "Analyzing gaussian proposal distributions for mapping with rao-blackwellized particle filters," in *2007 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3485–3490, IEEE, 2007.

Bibliography

- [62] J. Guivant and F. Masson, "Using Absolute Non-Gaussian Non-White Observations in Gaussian SLAM," in *Proceedings of the 2005 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 336–341, IEEE, 2005.
- [63] Y. Yi and L. Ding, "Robot simultaneous localization and mapping based on non-linear interacting multiple model concerning statistical property mutation," in *Measuring Technology and Mechatronics Automation in Electrical Engineering* (Z. Hou, ed.), vol. 135 of *Lecture Notes in Electrical Engineering*, pp. 441–450, Springer New York, 2012.
- [64] N. Aspert, D. Santa Cruz, and T. Ebrahimi, "MESH: measuring errors between surfaces using the Hausdorff distance.," in *IEEE International Conference on Multimedia and Expo (ICME)*, pp. 705–708, 2002.
- [65] C. F. Olson and L. Matthies, "Maximum Likelihood Rover Localization by Matching Range Maps.," in *Proceedings of the 1998 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 272–277, IEEE Computer Society, 1998.
- [66] R. De Maesschalck, D. Jouan-Rimbaud, and D. L. Massart, "The mahalanobis distance," *Chemometrics and intelligent laboratory systems*, vol. 50, no. 1, pp. 1–18, 2000.
- [67] J. Neira and J. D. Tardós, "Data association in stochastic mapping using the joint compatibility test.," *IEEE T. Robotics and Automation*, vol. 17, no. 6, pp. 890–897, 2001.
- [68] F. Vaussard, M. Bonani, P. Rétoznaz, A. Martinoli, and F. Mondada, "Towards Autonomous Energy-Wise RObjects.," in *TAROS* (R. Groß, L. Alboul, C. Melhuish, M. Witkowski, T. J. Prescott, and J. Penders, eds.), vol. 6856 of *Lecture Notes in Computer Science*, pp. 311–322, Springer, 2011.
- [69] S. Beeby and N. White, *Energy Harvesting for Autonomous Systems*. Artech House Publishers, 2010.
- [70] S. Priya and D. J. Inman, *Energy Harvesting Technologies*. Springer Publishing Company, Incorporated, 1st ed., 2008.
- [71] J. Randall, *Designing indoor solar products*. Wiley Online Library, 2006.
- [72] S. Roundy, P. Wright, and J. Rabaey, *Energy scavenging for wireless sensor networks: with special focus on vibrations*. Springer Netherlands, 2004.
- [73] J. Thomas, M. Qidwai, and J. Kellogg, "Energy scavenging for small-scale unmanned systems," *Journal of Power sources*, vol. 159, no. 2, pp. 1494–1509, 2006.
- [74] J. Crisp, M. Adler, J. Matijevic, S. Squyres, R. Arvidson, and D. Kass, "Mars exploration rover mission," *Journal of Geophysical Research*, vol. 108, no. E12, p. 8061, 2003.
- [75] C. Melhuish, I. Ieropoulos, J. Greenman, and I. Horsfield, "Energetically autonomous robots: Food for thought," *Autonomous Robots*, vol. 21, no. 3, pp. 187–198, 2006.

- [76] M. Laibowitz and J. Paradiso, "Parasitic mobility for pervasive sensor networks," *Pervasive Computing*, pp. 255–278, 2005.
- [77] F. Rey, M. Leidi, and F. Mondada, "Interactive mobile robotic drinking glasses," in *Distributed Autonomous Robotic Systems 8 (DARS)*, pp. 543–551, Springer, 2009.
- [78] C. Knight and J. Davidson, "Thermal Energy Harvesting for Wireless Sensor Nodes with Case Studies," *Advances in Wireless Sensors and Sensor Networks*, pp. 221–242, 2010.
- [79] C. Yu and K. Chau, "Thermoelectric automotive waste heat energy recovery using maximum power point tracking," *Energy Conversion and Management*, vol. 50, no. 6, pp. 1506–1512, 2009.
- [80] G. Solbrekken, K. Yazawa, and A. Bar-Cohen, "Thermal management of portable electronic equipment using thermoelectric energy conversion," in *The Ninth Intersociety Conference on Thermal and Thermomechanical Phenomena in Electronic Systems (ITHERM'04)*, pp. 276–283, IEEE, 2004.
- [81] M. Kishi, H. Nemoto, T. Hamao, M. Yamamoto, S. Sudou, M. Mandai, and S. Yamamoto, "Micro thermoelectric modules and their application to wristwatches as an energy source," in *Eighteenth International Conference on Thermoelectrics*, pp. 301–307, IEEE, 2002.
- [82] V. Leonov, T. Torfs, P. Fiorini, and C. Van Hoof, "Thermoelectric converters of human warmth for self-powered wireless sensor nodes," *Sensors Journal, IEEE*, vol. 7, no. 5, pp. 650–657, 2007.
- [83] J. Weber, K. Potje-Kamloth, F. Haase, P. Detemple, F. Völklein, and T. Doll, "Coin-size coiled-up polymer foil thermoelectric power generator for wearable electronics," *Sensors and Actuators A: Physical*, vol. 132, no. 1, pp. 325–330, 2006.
- [84] K. Sangani, "The sun in your pocket," *Eng. Technol*, vol. 2, no. 8, pp. 36–38, 2007.
- [85] T. O'Donnell and W. Wang, "Power Management, Energy Conversion and Energy Scavenging for Smart Systems," *Ambient Intelligence with Microsystems*, pp. 241–266, 2009.
- [86] A. Virtuani, E. Lotter, and M. Powalla, "Influence of the light source on the low-irradiance performance of Cu (In, Ga) Se₂ solar cells," *Solar Energy Materials and Solar Cells*, vol. 90, no. 14, pp. 2141–2149, 2006.
- [87] M. Green, K. Emery, Y. Hishikawa, and W. Warta, "Solar cell efficiency tables (version 34)," *Progress in Photovoltaics: Research and Applications*, vol. 17, no. 5, pp. 320–326, 2009.
- [88] J. Randall and J. Jacot, "The performance and modelling of 8 photovoltaic materials under variable light intensity and spectra," in *World Renewable Energy Conference VII Proceedings, Cologne, Germany, 2002*.

Bibliography

- [89] M. Grätzel, "Solar energy conversion by dye-sensitized photovoltaic cells," *Inorg. Chem*, vol. 44, no. 20, pp. 6841–6851, 2005.
- [90] A. Barnett, D. Kirkpatrick, C. Honsberg, D. Moore, M. Wanlass, K. Emery, R. Schwartz, D. Carlson, S. Bowden^{1a}, D. Aiken, *et al.*, "Milestones toward 50% efficient solar cell modules," in *Presented at the 22nd European Photovoltaic Solar Energy Conference*, vol. 3, 2007.
- [91] A. Smith, *Radio Frequency Principles & Applications*. Universities Press, 1998.
- [92] M. Philipose, J. Smith, B. Jiang, A. Mamishev, S. Roy, and K. Sundara-Rajan, "Battery-free wireless identification and sensing," *IEEE Pervasive Computing*, pp. 37–45, 2005.
- [93] B. Obrist and S. Hegnauer, "A microwave powered data transponder," *Sensors and Actuators A: Physical*, vol. 46, no. 1-3, pp. 244–246, 1995.
- [94] J. Reutter, "Horloge à remontage automatique par les variations de température ou de pression atmosphérique," Jan. 15 1929. Swiss Patent CH130941A.
- [95] S. Phillips, "Temperature responsive self winding timepieces," Oct. 1 2002. US Patent 6,457,856.
- [96] E. Callaway, *Wireless sensor networks: architectures and protocols*. CRC press, 2004.
- [97] J. Lebet and F. Jequier, *Living on Air: History of the Atmos Clock*. Jaeger-LeCoultre, 1997.
- [98] J. Bean, "Energy saving shower head," June 9 2005. US Patent App. 11/148,524.
- [99] P. Mitcheson, E. Yeatman, G. Rao, A. Holmes, and T. Green, "Energy harvesting from human and machine motion for wireless electronic devices," *Proceedings of the IEEE*, vol. 96, no. 9, pp. 1457–1486, 2008.
- [100] C. Federspiel and J. Chen, "Air-powered sensor," in *Sensors, 2003. Proceedings of IEEE*, vol. 1, pp. 22–25, IEEE, 2004.
- [101] D. Rancourt, A. Tabesh, and L. Fréchet, "Evaluation of centimeter-scale micro windmills: aerodynamics and electromagnetic power generation," *Proceedings of Power-MEMS 2007*, pp. 28–29, 2007.
- [102] A. Flammini, D. Marioli, E. Sardini, and M. Serpelloni, "An autonomous sensor with energy harvesting capability for airflow speed measurements," in *2010 IEEE Instrumentation and Measurement Technology Conference (I2MTC)*, pp. 892–897, IEEE, 2010.
- [103] E. Sardini and M. Serpelloni, "Self-Powered Wireless Sensor for Air Temperature and Velocity Measurements With Energy Harvesting Capability," *IEEE Transactions on Instrumentation and Measurement*, vol. PP, no. 99, pp. 1–7, 2010.
- [104] S. Beeby, M. Tudor, and N. White, "Energy harvesting vibration sources for microsystems applications," *Measurement science and technology*, vol. 17, p. R175, 2006.

-
- [105] B. Op het Veld, D. Hohlfeld, and V. Pop, "Harvesting mechanical energy for ambient intelligent devices," *Information Systems Frontiers*, vol. 11, no. 1, pp. 7–18, 2009.
- [106] S. Arms, C. Townsend, D. Churchill, S. Moon, and N. Phan, "Energy Harvesting Wireless Sensors for Helicopter Damage Tracking," in *Annual Forum Proceedings-American Helicopter Society*, vol. 62, p. 1336, American Helicopter Society, Inc, 2006.
- [107] X. Cao, W.-J. Chiang, Y.-C. King, and Y.-K. Lee, "Electromagnetic energy harvesting circuit with feedforward and feedback DC–DC PWM boost converter for vibration power generator system," *Power Electronics, IEEE Transactions on*, vol. 22, no. 2, pp. 679–685, 2007.
- [108] J. Donelan, Q. Li, V. Naing, J. Hoffer, D. Weber, and A. Kuo, "Biomechanical energy harvesting: generating electricity during walking with minimal user effort," *Science*, vol. 319, no. 5864, pp. 807–810, 2008.
- [109] J. Feenstra, J. Granstrom, and H. Sodano, "Energy harvesting through a backpack employing a mechanically amplified piezoelectric stack," *Mechanical Systems and Signal Processing*, vol. 22, no. 3, pp. 721–734, 2008.
- [110] J. M. Renno, M. F. Daqaq, and D. J. Inman, *On the Optimal Energy Harvesting from a Vibration Source Using a Piezoelectric Stack*. Springer, 2009.
- [111] G. Zhu, R. Yang, S. Wang, and Z. L. Wang, "Flexible high-output nanogenerator based on lateral ZnO nanowire array," *Nano letters*, vol. 10, no. 8, pp. 3151–3155, 2010.
- [112] I. L. Baginsky, E. G. Kostsov, and A. A. Sokolov, "Electrostatic microgenerators of energy with a high specific power," *Optoelectronics, Instrumentation and Data Processing*, vol. 46, no. 6, p. 580, 2010.
- [113] P. D. Mitcheson, T. Sterken, C. He, M. Kiziroglou, E. Yeatman, and R. Puaers, "Electrostatic microgenerators," *Measurement and Control*, vol. 41, no. 4, pp. 114–119, 2008.
- [114] B. Yen and J. H. Lang, "A variable-capacitance vibration-to-electric energy harvester," *Circuits and Systems I: Regular Papers, IEEE Transactions on*, vol. 53, no. 2, pp. 288–295, 2006.
- [115] S. Roundy, P. Wright, and J. Rabaey, "A study of low level vibrations as a power source for wireless sensor nodes," *Computer Communications*, vol. 26, no. 11, pp. 1131–1144, 2003.
- [116] T. Starner, "Human-powered wearable computing," *IBM Systems Journal*, vol. 35, no. 3&4, 1996.
- [117] J. Kymissis, C. Kendall, J. Paradiso, and N. Gershenfeld, "Parasitic power harvesting in shoes," in *Second International Symposium on Wearable Computers*, pp. 132–139, 1998.
- [118] I. Mandal and P. Patra, "Renewable Energy Source," *International Journal of Computer Applications IJCA*, vol. 1, no. 17, pp. 44–53, 2010.

Bibliography

- [119] M. J. T. Lewis, *Millstone and Hammer: the origins of water power*. University of Hull, 1997.
- [120] J. Paulides, J. Jansen, L. Encica, E. Lomonova, and M. Smit, "Human-powered small-scale generation system for a sustainable dance club," in *IEEE International Electric Machines and Drives Conference (IEMDC '09)*, pp. 439–444, IEEE, 2009.
- [121] S. Hollar, *A solar-powered, milligram prototype robot from a three-chip process*. PhD thesis, University Of California, 2003.
- [122] A. Boletis, W. Driesen, J.-M. Breguet, and A. Brunete, "Solar Cell Powering with Integrated Global Positioning System for mm3 Size Robots.," in *2006 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 5528–5533, IEEE, 2006.
- [123] J. Lever and L. Ray, "Revised solar-power budget for Cool Robot polar science campaigns," *Cold Regions Science and Technology*, vol. 52, no. 2, pp. 177–190, 2008.
- [124] H.-T. Lee, W.-C. Lin, and C.-H. Huang, "Indoor Surveillance Security Robot with a Self-Propelled Patrolling Vehicle," *Journal of Robotics*, vol. 2011, 2011.
- [125] T. Marco, V. Cristina, A. Paolo, P. Luca, G. Dario, and P. Massimo, "Design considerations about a photovoltaic power system to supply a mobile robot," in *2010 IEEE International Symposium on Industrial Electronics (ISIE)*, pp. 1829–1834, IEEE.
- [126] C. Mei, "Self-Sustaining Robot," Master's thesis, National University of Singapore, 2008.
- [127] "BEAM Online Community." <http://www.beam-online.com>, 2014.
- [128] A. Asgaonkar, "Design of an Efficient Solar Engine Circuit for Autonomous Robotics," *International Journal of Computer Applications IJCA*, vol. 7, no. 7, pp. 1–4, 2010.
- [129] S. Jain and N. Tiwari, "BEAM Dragster Solar Robot with it's Strategic Winning Design Analysis for the Nexus Photoroller Competition," *International Journal of Computer Applications*, vol. 1, no. 28, 2010.
- [130] J. Lever, L. Ray, A. Streeter, and A. Price, "Solar power for an Antarctic rover," *Hydrological processes*, vol. 20, no. 4, pp. 629–644, 2006.
- [131] J. Lever, A. Streeter, and L. Ray, "Performance of a solar-powered robot for polar instrument networks," in *Proceedings of the 2006 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4252–4257, IEEE, 2006.
- [132] G. Zhao, L. Wang, S. Cheng, and J. Zhang, "Design and realization of a mobile robot for power transmission lines inspection," in *2011 International Conference on Mechatronics and Automation (ICMA)*, pp. 1509–1514, IEEE, 2011.
- [133] "Husqvarna Automower Hybrid." <http://www.husqvarna.com/de/produkte/automower/automower-solar-hybrid/>, 2014.

- [134] G. T. Sibley, M. H. Rahimi, and G. S. Sukhatme, "Robomote: A Tiny Mobile Robot Platform for Large-Scale Ad-Hoc Sensor Networks.," in *Proceedings of the 2002 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 1143–1148, IEEE, 2002.
- [135] M. H. Rahimi, H. Shah, G. S. Sukhatme, J. S. Heidemann, and D. Estrin, "Studying the Feasibility of Energy Harvesting in a Mobile Sensor Network.," in *Proceedings of the 2003 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 19–24, IEEE, 2003.
- [136] P. Hartono, K. Tabe, K. Suzuki, and S. Hashimoto, "Strategy acquirement by survival robots in outdoor environment.," in *Proceedings of the 2003 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 3571–3575, IEEE, 2003.
- [137] A. Afarulrazi, W. Utomo, K. Liew, and M. Zarafi, "Solar Tracker Robot using microcontroller," in *Business, Engineering and Industrial Applications (ICBEIA), 2011 International Conference on*, pp. 47–50, IEEE, 2011.
- [138] D. Shirley and J. Matijevic, "Mars pathfinder microrover," *Autonomous Robots*, vol. 2, no. 4, pp. 283–289, 1995.
- [139] P. Stella, R. Ewell, and J. Hoskin, "Design and performance of the MER (Mars Exploration Rovers) solar arrays," in *Conference Record of the Thirty-first IEEE Photovoltaic Specialists Conference*, pp. 626–630, IEEE, 2005.
- [140] M. Smart, R. Bugga, K. Chin, L. Whitcanack, and S. Surampudi, "Performance characteristics of lithium-ion technology under extreme environmental conditions," in *1st International Energy Conversion Engineering Conference*, 2003.
- [141] ASTM G173-03(2012), "Standard Tables for Reference Solar Spectral Irradiances: Direct Normal and Hemispherical on 37° Tilted Surface," standard, ASTM International, West Conshohocken, PA, 2012.
- [142] C. Gueymard, "Critical analysis and performance assessment of clear sky solar irradiance models using theoretical and measured data," *Solar Energy*, vol. 51, no. 2, pp. 121–138, 1993.
- [143] S. R. Wenham, *Applied photovoltaics*. Routledge, 2011.
- [144] U.S. Naval Observatory; H.M. Nautical Almanac Office, *The Astronomical Almanac for the Year 2010*. US Government Printing Office, 2009.
- [145] J. E. Braun and J. C. Mitchell, "Solar Geometry for Fixed and Tracking Surfaces," *Solar Energy*, vol. 31, no. 5, pp. 439–444, 1983.
- [146] F. Vaussard, P. Rétornaz, M. Liniger, and F. Mondada, "The Autonomous Photovoltaic MarXbot.," in *IAS (2)* (S. Lee, H. Cho, K.-J. Yoon, and J. Lee, eds.), vol. 194 of *Advances in Intelligent Systems and Computing*, pp. 175–183, Springer, 2012.

Bibliography

- [147] ASTM E490-00a(2014), "Standard Solar Constant and Zero Air Mass Solar Spectral Irradiance Tables," standard, ASTM International, West Conshohocken, PA, 2014.
- [148] B. G. Streetman and S. Banerjee, *Solid state electronic devices*, vol. 2. Prentice-Hall Englewood Cliffs, NJ, 1995.
- [149] T. Eswam and P. Chapman, "Comparison of Photovoltaic Array Maximum Power Point Tracking Techniques," *Energy Conversion, IEEE Transactions on*, vol. 22, pp. 439–449, June 2007.
- [150] R. Faranda and S. Leva, "Energy comparison of MPPT techniques for PV Systems," *WSEAS transactions on power systems*, vol. 3, no. 6, pp. 446–455, 2008.
- [151] G. Bunea, K. Wilson, Y. Meydbray, M. Campbell, and D. De Ceuster, "Low Light Performance of Mono-Crystalline Silicon Solar Cells," in *Photovoltaic Energy Conversion, Conference Record of the 2006 IEEE 4th World Conference on*, vol. 2, pp. 1312–1314, May 2006.
- [152] S. Kohlbrecher, J. Meyer, O. von Stryk, and U. Klingauf, "A Flexible and Scalable SLAM System with Full 3D Motion Estimation," in *Proc. IEEE International Symposium on Safety, Security and Rescue Robotics (SSRR)*, IEEE, November 2011.
- [153] S. Thrun, W. Burgard, and D. Fox, *Probabilistic robotics*. MIT Press, 2005.
- [154] J. Fink, *Dynamics of Human-Robot Interaction in Domestic Environments*. PhD thesis, Lausanne, 2014.
- [155] X.-Z. Guo, Y.-D. Zhang, D. Qin, Y.-H. Luo, D.-M. Li, Y.-T. Pang, and Q.-B. Meng, "Hybrid tandem solar cell for concurrently converting light and heat energy with utilization of full solar spectrum," *Journal of Power Sources*, vol. 195, no. 22, pp. 7684–7690, 2010.
- [156] H. Yu, Y. Li, Y. Shang, and B. Su, "Design and investigation of photovoltaic and thermo-electric hybrid power source for wireless sensor networks," in *3rd IEEE International Conference on Nano/Micro Engineered and Molecular Systems (NEMS 2008)*, pp. 196–201, IEEE, 2008.

Curriculum Vitae

Florian Vaussard

Chemin de la Biolle 4
Batterie de la Biolle 511
1807 Blonay – Switzerland

Born 12th of March 1984
Swiss citizen
Married

Education

2010 – 2014	Ph.D. in Mobile Robotics, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
2002 – 2007	M.Sc. in Microengineering, École Polytechnique Fédérale de Lausanne (EPFL), Switzerland
1999 – 2002	Maturité fédérale , Gymnase cantonal de Burier, Switzerland

Professional Experience

2009 – 2010	R&D Engineer (Swiss civilian service), Distributed Systems Laboratory (EPFL – DISAL), Switzerland
2008 (1 month)	R&D Engineer , Justlife, Switzerland
2008	R&D Engineer , BlueBotics, Switzerland
2007 – 2008	Engineer Intern , JTEKT, Japan

Languages

French	C2
English	C1
Japanese	B1
German	B1

Peer-Reviewed Publications

1. **F. Vaussard**, J. Fink, V. Bauwens, P. Réturnaz, D. Hamel, P. Dillenbourg and F. Mondada: *Lessons learned from robotic vacuum cleaners entering the home ecosystem*. In Robotics and Autonomous Systems, 62(3), pages 376 – 391 (2014)
2. J. Fink, S. Lemaignan, P. Dillenbourg, P. Réturnaz, **F. Vaussard**, A. Berthoud, F. Mondada, F. Wille and K. Franinovic: *Which Robot Behavior Can Motivate Children to Tidy up Their Toys? Design and Evaluation of “Ranger”*. In Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction (HRI’14), pages 439-446 (2014)
3. M. Dorigo, D. Floreano, L. Gambardella, F. Mondada, S. Nolfi, T. Baaboura, M. Birattari, M. Bonani, M. Brambilla, A. Brutschy, D. Burnier, A. Campo, A. L. Christensen, A. Decugnière, G. Di Caro, F. Ducatelle, E. Ferrante, A. Förster, J. Guzzi, V. Longchamp, S. Magnenat, J. M. Gonzales, N. Mathews, M. Montes de Oca, R. O’Grady, C. Pinciroli, G. Pini, P. Réturnaz, J. Roberts, V. Sperati, T. Stirling, A. Stranieri, T. Stützle, V. Trianni, E. Tuci, A. E. Turgut and **F. Vaussard**: *Swarmanoid: a novel concept for the study of heterogeneous robotic swarms*. In IEEE Robotics & Automation Magazine, 20(4), pages 60-71(2013)
4. **F. Vaussard**, P. Réturnaz, M. Liniger and F. Mondada: *The Autonomous Photovoltaic MarXbot*. In Proceeding of the International Conference on Intelligent Autonomous System (IAS-12), pages 175-183. Advances in Intelligent Systems and Computing (194), Springer Berlin Heidelberg (2013)
5. **F. Vaussard**, P. Réturnaz, S. Roelofsen, M. Bonani, F. Rey and F. Mondada: *Towards Long-Term Collective Experiments*. In Proceeding of the International Conference on Intelligent Autonomous System (IAS-12), pages 683–692. Advances in Intelligent Systems and Computing (194), Springer Berlin Heidelberg (2013)
6. P. Réturnaz, F. Riedo, S. Magnenat, **F. Vaussard**, M. Bonani and F. Mondada: *Seamless Multi-Robot Programming for the People: ASEBA and the Wireless Thymio II Robot*. In 2013 IEEE International Conference on Information and Automation (ICIA), pages 337 - 343, IEEE (2013)
7. **F. Vaussard**, P. Réturnaz, D. Hamel and F. Mondada; *Cutting Down the Energy Consumed by Domestic Robots: Insights from Robotic Vacuum Cleaners*. In Joint Proceedings of the 13th Annual TAROS Conference and the 15th Annual FIRA RoboWorld Congress. Lecture Notes in Computer Science (7429), pages 128–139, Springer (2012)
8. **F. Vaussard**, M. Bonani, P. Réturnaz, A. Martinoli and F. Mondada: *Towards Autonomous Energy-Wise RObjects*. In Proceedings of the 12th Conference Towards Autonomous Robotic Systems (TAROS). Lecture Notes in Computer Science (6856), pages 311–322, Springer (2011)
9. M. Bonani, V. Longchamp, S. Magnenat, P. Réturnaz, D. Burnier, G. Roulet, **F. Vaussard**, H. Bleuler and F. Mondada: *The MarXbot, a Miniature Mobile Robot Opening new Per-*

spectives for the Collective-robotic Research. In International Conference on Intelligent Robots and Systems (IROS), 2010 IEEE/RSJ, pages 4187-4193, IEEE Press (2010)

10. S. Magnenat, P. Schoeneich, F. Rochat, P. Rétornaz, M. Bonani, V. Longchamp, M. Voelke, T. Barras, D. Burnier, P. Noirat, T. Baaboura, **F. Vaussard** and F. Mondada: *Autonomous Construction by a Mobile Robot in Unknown Environments with Scarce Resources*. In 4th International Conference on Cognitive Systems, ETH Zurich, Switzerland, January 27 & 28 (2010)