

DESIGN OF HIGH-PERFORMANCE LEGGED ROBOTS

A CASE STUDY ON A HOPPING AND BALANCING ROBOT

DESIGN OF HIGH-PERFORMANCE LEGGED ROBOTS

A CASE STUDY ON A HOPPING AND BALANCING ROBOT

Dissertation

submitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy

Department of Informatics, Bioengineering, Robotics, System Engineering DIBRIS
University of Genova

and

Department of Advanced Robotics Istituto Italiano di Tecnologia

by

Josephus J. M. DRIESSEN

Advisor

Dr. R. Featherstone

Istituto Italiano di Tecnologia



Dibris



Keywords: Mechanical Design, Legged Robots, Holistic Design, Natural Dynamics, Hopping Robots, Balancing Robots

Copyright © 2019 by Josephus J. M. Driessen

*It is common sense to take a method and try it.
If it fails, admit it frankly and try another.
But above all, try something.*

Franklin D. Roosevelt, Looking Forward

CONTENTS

Abstract	xi
Acknowledgements	xiii
List of abbreviations	xv
1 Introduction	1
1.1 Objectives	2
1.2 Approach	2
1.3 Contributions	3
1.4 Outline	5
2 Designing for high physical performance	7
2.1 Common measures for physical performance	7
2.2 Jump height	9
2.2.1 Defining jump height	9
2.2.2 Jump height and size	9
2.2.3 Jump height and time	11
2.3 Robot comparisons	13
2.4 Design philosophy	15
2.4.1 Complexity	16
2.4.2 Key points	17
3 Skippy: state-of-the-art	31
3.1 Preliminary design	31
3.1.1 Design decisions	31
3.1.2 Design principles	33
3.2 Preliminary hopping studies	34
3.2.1 Strategy	35
3.3 Early balancing studies	36
3.3.1 Comparison to other work	36
3.3.2 Model	37
3.3.3 Control law	38
3.3.4 Transfer function	40
3.3.5 Leaning in anticipation	40
3.3.6 Balance offset observer	41
4 Maximizing balancing performance	43
4.1 Introduction	43
4.2 Background	45
4.3 Model	46

4.4	Strategy	48
4.5	Derivations	49
4.6	Validation	53
4.6.1	Simulation setup	53
4.6.2	Simulation results	54
4.6.3	Velocity gain	55
4.7	Effective offset angle	56
4.8	Conclusion	57
5	Tippy: a balancing-only prototype	59
5.1	Introduction	59
5.2	Requirements	60
5.3	Critical design decisions	61
5.3.1	Simplified actuation	62
5.3.2	Base joint	63
5.3.3	Customizability	63
5.4	Detailed design description	64
5.4.1	Kinematic structure	64
5.4.2	Work space	65
5.4.3	Actuators	67
5.4.4	Brace	72
5.4.5	Electronics	72
5.5	Swivel experimental setup	76
5.5.1	Control law simplifications	76
5.5.2	Implementations	78
5.6	Experimental results	80
5.6.1	Maximum recovery angle experiments	80
5.6.2	Balancing experiments	81
5.7	Reflection	84
5.8	Conclusion	86
6	Skippy mechanical design	87
6.1	Design overview	87
6.2	Design revisions	88
6.2.1	Battery placement	90
6.2.2	Main actuator	90
6.2.3	Swivel actuator	94
6.3	Physical robustness	95
6.3.1	Reducing peak decelerations: bumpers	95
6.3.2	Reducing internal peak loads	97
6.3.3	Increasing permitted loadings	100
6.4	Structural design	101
6.4.1	Design style	102
6.4.2	Tube-based design	103
6.4.3	Design of custom aluminium structural components	104

6.5	Hip design	109
6.6	Ankle and foot design	113
6.6.1	Ankle spring design	114
6.6.2	Ankle torque profile calculations.	115
6.6.3	Parameter design considerations	116
6.7	Main actuator design	118
6.7.1	Rocker	119
6.7.2	Guiding system	119
6.7.3	Main spring	120
6.7.4	Actuator length	123
6.8	Swivel actuator design	124
6.9	Head	126
6.10	IMU placement	129
6.10.1	Centre of percussion definition	130
6.10.2	Centre of percussion derivation	131
6.11	Conclusion	134
7	Modelling of curved compression leaf springs	137
7.1	Introduction	137
7.2	Derivation	139
7.2.1	Kinematics.	139
7.2.2	Statics	140
7.3	Validation.	141
7.4	Failure	143
7.5	Future considerations.	146
7.5.1	Shape efficiency	146
7.6	Conclusion	147
8	Design optimisations	149
8.1	Introduction	149
8.2	Skippy model	150
8.2.1	Dynamic model description	151
8.2.2	Hybrid dynamic and loop closure calculations.	154
8.2.3	Main actuator model.	155
8.2.4	Ankle spring	157
8.2.5	End stop models	157
8.3	Optimisation environment	157
8.3.1	Constraint function	158
8.3.2	Voltage command signal.	160
8.3.3	Optimisation parameters and bounds	161
8.4	Results and revisions	168
8.4.1	Initial results.	168
8.4.2	Updated design description	170
8.4.3	Updated design results.	171
8.5	Discussion and conclusion	174

9 Discussion, conclusion and future development	177
9.1 Discussion	177
9.2 Conclusion	180
9.3 Future developments	180
A Jumping and scaling	183
B Triangle functions	187
C Skippy's loop closure	189
References	193
Curriculum Vitae	207

ABSTRACT

The availability and capabilities of present-day technology suggest that legged robots should be able to physically outperform their biological counterparts. This thesis revolves around the philosophy that the observed opposite is caused by over-complexity in legged robot design, which is believed to substantially suppress design for high-performance. In this dissertation a design philosophy is elaborated with a focus on simple but high performance design. This philosophy is governed by various key points, including holistic design, technology-inspired design, machine and behaviour co-design and design at the performance envelope. This design philosophy also focuses on improving progress in robot design, which is inevitably complicated by the aspire for high performance. It includes an approach of iterative design by trial-and-error, which is believed to accelerate robot design through experience.

This thesis mainly focuses on the case study of Skippy, a fully autonomous monopedal balancing and hopping robot. Skippy is maximally simple in having only two actuators, which is the minimum number of actuators required to control a robot in 3D. Despite its simplicity, it is challenged with a versatile set of high-performance activities, ranging from balancing to reaching record jump heights, to surviving crashes from several meters and getting up unaided after a crash, while being built from off-the-shelf technology.

This thesis has contributed to the detailed mechanical design of Skippy and its optimisations that abide the design philosophy, and has resulted in a robust and realistic design that is able to reach a record jump height of 3.8 m.

Skippy is also an example of iterative design through trial-and-error, which has lead to the successful design and creation of the balancing-only precursor Tippy. High-performance balancing has been successfully demonstrated on Tippy, using a recently developed balancing algorithm that combines the objective of tracking a desired position command with balancing, as required for preparing hopping motions.

This thesis has furthermore contributed to several ideas and theories on Skippy's road of completion, which are also useful for designing other high-performance robots. These contributions include (1) the introduction of an actuator design criterion to maximize the physical balance recovery of a simple balancing machine, (2) a generalization of the centre of percussion for placement of components that are sensitive to shock and (3) algebraic modelling of a non-linear high-gravimetric energy density compression spring with a regressive stress-strain profile.

The activities performed and the results achieved have been proven to be valuable, however they have also delayed the actual creation of Skippy itself. A possible explanation for this happening is that Skippy's requirements and objectives were too ambitious, for which many complications were encountered in the decision-making progress of the

iterative design strategy, involving trade-offs between exercising trial-and-error, elaborate simulation studies and the development of above-mentioned new theories. Nevertheless, from (1) the resulting realistic design of Skippy, (2) the successful creation and demonstrations of Tippy and (3) the contributed theories for high-performance robot design, it can be concluded that the adopted design philosophy has been generally successful.

Through the case study design project of the hopping and balancing robot Skippy, it is shown that proper design for high physical performance (1) can indeed lead to a robot design that is capable of physically outperforming humans and animals and (2) is already very challenging for a robot that is intended to be very simple.

ACKNOWLEDGEMENTS

This thesis would not have been possible without the support of many of my colleagues during the years I worked on Tippy and the design of Skippy.

First, I would like to thank Dr. Roy Featherstone and Prof. Darwin Caldwell for giving me the opportunity to do this research at the Italian Institute of Technology (IIT) on Tippy and the design of Skippy, on which I already worked for my master thesis. In addition, I would like to thank him for supporting my idea to work on a more theoretical part of the thesis during my long-lasting recovery from surgery in 2017, which finally resulted in the acceptance of my first paper. I also would like to thank Antony and Roodra for their contributions to the project, who have taken care of most of Tippy's software implementations. It has furthermore been a great pleasure to collaborate with Elco and to provide him daily supervision on the ring screw project, which has advanced a major step to becoming a real product. Towards realisation of the robots, I would like to thank the technicians of ADVR, with a special mention to Alex, Lorenzo and Stefano for their feedback on various aspects of mechanical design and putting me in contact with IIT's mechanical workshop. In writing the thesis I gratefully used the constructive feedback from Roy Featherstone on the first half of the thesis and from my colleagues and friends Romeo, Giovanni, Michele and Elco on several other parts of my thesis.

However, this is not where it stops. It is because of many people that have made my stay in Genova a pleasant one; I could even say that Genova now feels like a second home. Thank you Romeo, for welcoming me to your family and teaching me the importance of having a pizza oven in your garden. Alex, since day one of the PhD track, my partner in crime for battling Italian bureaucracy. Claudio, for the fair squash games and good advice. Octavio and Maria for being the cutest neighbours ever. Michele, for playing the guitar and losing it. Victor, for being a true teacher of (salty) Brazilian BBQs. Marielle, Ahmed, Mesut and other Babozz, hiking would not have been the same without you. Nawid and Stefka, you were never supposed to leave Genova. Fellow Dutchmen Elco, Stijn, Martijn, Rian, Wesley and Nikol, who I'm still in touch with despite all having set sail back to the Netherlands. My fellow PhD pals Amira, Silvia, Marie and Giulia, for sharing holidays, walks and sushi. Carlos and Andreea, for all the adventures on the dance floor. Geoff and Abril, for later joining that same dance floor. Lisa, my forever source of wisdom for baking cakes. Ping, for helping me eat the cakes. Juan, my new Masetto buddy. Richard, Nora and Paul, when is the next cheese and wine session? Lidia, the only Italian that likes liquorish. Shamel, for always showing up, irrespective of who is there or for how long we are staying. Fresh arrivals Carlos and Hector, for showing interest in my work! My (former) housemates Samu, Marco, Sara, Lucia, Ieva and Luca, for the good talks, laughs and dinners. Dimi, Pietro, Marta and Wan-Soo, for the adventures abroad, at the falls and the nursing on the train. Emily, Juliette, Gennaro, Evie, Louis and Fanny for being random and throwing good parties. Olmo, Alperen, Edwin, Joao and many more, for attending with me those parties and other ventures.

Last but not least, I would like to gratefully thank my parents, grandparents and other family members who always supported me in doing my work in both happy but also in difficult personal circumstances. I know they will always be my rocks.

LIST OF ABBREVIATIONS

CAD	computer aided design
CoM	centre of mass
CoPc	centre of percussion
(T/M)CoT	(total/mechanical) cost of transport
COTS	commercial off-the-shelf
DC	direct current
DoF	degree of freedom
EoM	equation of motion
FEA	finite element analysis
(C/G)FRP	(carbon/glass) fibre reinforced plastic
GRF	ground reaction force
IIT	Istituto Italiano di Tecnologia (Italian Institute of Technology)
IMU	inertial measurement unit
LPF	low-pass filter
MCU	microcontroller unit
PCB	printed circuit board
PRBM	pseudo-rigid body model
PWM	pulse width modulation
RWP	reaction wheel pendulum
SEA	series elastic actuation
SLIP	spring-loaded inverted pendulum
SPI	serial peripheral interface
SQP	sequential quadratic programming
XCoM	extrapolated centre of mass
ZRAM	zero rate of change of angular momentum

1

INTRODUCTION

Today's experimental legged robots show lower levels of physical performance than what one would expect, given the capabilities of the technologies available to make them. This notion can be supported by comparing the various building blocks that are available in both technology and biology. For example, we have structural materials that are stronger and lighter than bone, and actuators with better efficiency, bandwidth and power-to-weight ratios than biological muscle. Nevertheless, many animals jump higher, run faster and are more robust than any legged robot ever built.

Many theories can be devised as to why this is the state of affairs, but one possible explanation is that the legged robots of today are too complex to properly design for them to exhibit truly high-performance behaviour for the relatively small teams of researchers who make them.

This hypothesis has been put forward by Dr. Featherstone, whose incentive it then was to establish a project to design, build and control a simple 3D hopping and balancing robot, called Skippy. This robot is supposed to exhibit a level of physical performance substantially superior to today's legged robots, partly due to its inherent simplicity. The project's main purpose is testing the hypothesis and is referred to as the Skippy Project. *"The objective is to explore the limits of what is physically possible with today's technology, by means of a robot that is simple, robust and fairly cheap."* — Featherstone (2014)

INTRODUCING SKIPPY

Skippy is maximally simple in having only two actuators, the minimum number of actuators required to control a 3D robot. Yet it is challenged with a wide variety of required high performance activities. Its physically most demanding task is reaching a target height of 4 m through methods of hopping. It is a design requirement that Skippy be able to reach this height without relying on more than half of the energy being recycled from the previous hop. This implies that it should be able to reach its target height after landing from a previous hop that is not higher than half the target height. Reaching a jump height of 4 m pushes the limits of what should theoretically be possible

with off-the-shelf technology. While performing its hops, it should also be able to make athletic manoeuvres, such as somersaults and twists.

Moreover, Skippy is a monopedal robot with a single-point contact foot, and as such, in order for it to start hopping, it has to be able to actively balance and control its posture. The balancing and hopping requirements challenge controller design, but also mechanical design, because the physical requirements of a good hopping machine and a good balancing machine are fundamentally different from each other: jumping is governed by generating vertical momentum, whereas balancing is governed by the ability to generate horizontal momentum (Featherstone, 2016b). Lastly, Skippy is to operate in full autonomy (i.e. without any tether, umbilical, harness, safety net or human hand), which implies that it should be sufficiently mechanically robust to survive its crashes (from its highest jump height) and that it should be able to autonomously stand up after such a crash, whether the crash was intended or by mistake.

1.1. OBJECTIVES

The ultimate objective of the Skippy Project is to contribute to the improvement in the physical performance of legged robots in general by showing how such improvements can be obtained through a combination of better design and better control.

In particular, this thesis focuses on the case study of Skippy's design, rather than on its control. This design case study of Skippy should be exemplary to the design of other robots that exhibit high level of physical performance. Since finalizing Skippy (i.e. including making it fully autonomous) does not fit within the scope of a single thesis project, an objective for this thesis has been imposed that focusses on high physical performance and not on closed-loop control. This objective states that Skippy should demonstrate reaching its highest jumping height by being dropped from controlled initial conditions, as this can be accomplished with motion commands that are mainly feed forward.

1.2. APPROACH

The Skippy Project is ambitious. It knows many unknowns and possible pitfalls on its road of completion. We seek to push the boundaries in both high performance balancing and hopping, two warring objectives that are sufficiently challenging even when considered in isolation.

Skippy has only a point foot, which means that it is always unstable and that it must always balance itself whilst performing other tasks. This includes the preparation of jumping motions. As such, the ability to balance is crucial for the ability to jump or hop. Various research on improving balance control and performance has been done, but this research has not yet been implemented in reality. For this purpose, and for the first half of the thesis program, a project has been established to design and build a balancing-only robot that precedes Skippy, called Tippy. In essence, the balancing objective is addressed firstly in isolation. This project has a secondary purpose of serving as a design challenge and exercise, essentially employing the method of design by trial-and-error, which makes part of our design philosophy: having a fully operational machine ready at an early stage aids to encapsulate issues ahead of time which can be useful for improving

or even critical for finalising Skippy's design.

The Skippy design project initiates after the Tippy project is completed to sufficient extent, which includes having obtained some experimental results. This allows us to reflect on the mechanical, electrical and software design and on balance performance in practice.

Unlike tackling the balancing objective by designing and building Tippy, the hopping objective is not tackled in isolation by building a separate hopping-only machine prototype. Since a monopedal hopping machine requires the ability to balance, but not the other way round, such a machine would require more impeding mobility constraints than Tippy (e.g. a rail to limit its motion to 1D). Building such a machine was expected not to have added benefits. Instead it was decided to rely on realistic simulation studies, which mainly concern power modulation with a focus on the feasibility of performing Skippy's highest hop. Then, with the help of automated optimisations, a Skippy design should be finalised. The goal of optimisations is not to find an optimal design, but to explore Skippy's behaviour and come to a design that is good enough to fulfil the objective. We believe that this is done best by focusing on the single most physically demanding item in Skippy's repertoire of behaviours, which is why the design and related optimisation are centred on achieving the highest hop.

1.3. CONTRIBUTIONS

The general contribution of this thesis is a design study of the provisional case study of Skippy, which is exemplary for a design philosophy that is aimed to improve the design of robots that are to exhibit a high degree of physical performance, and to improve their design progress. As postulated by Dr. Featherstone, the design philosophy has its main focus on reducing complexity in robotics, which is believed to be a main factor in limiting robot physical performance below their potential, as indicated by the capabilities of present-day technology. In chapter 2 this design philosophy is further elaborated by the definition of various key points, which are intended to be incentives for making and accelerating high performance robot designs. The key points are the following, each of which are explained in detail in section 2.4.2:

- Robot designs should be holistic.
- Robot designs should be targeted on behaviour at the performance envelope.
- Robots should be a product of machine and behaviour co-design.
- Technology inspired design has priority over bio-inspired design.¹
- Progress is made through trial-and-error by making prototypes.
- Robots should have a reduced requirement for risk management.
- Robots should be designed to be physically robust.
- Robots should not be designed to be bigger than necessary.
- Robots should be designed to be fully autonomous, so that they do not require an umbilical, tether or harness.

¹This is not to say that bio-inspired design is poor design.

- Robots should be designed to be good enough rather than optimal.

The thesis has contributed to demonstrating the potentials and challenging aspects of this design philosophy. In particular, it has contributed to the developments of the two robots Tippy and Skippy.

- Tippy has been designed and built to be a good balancer, and has demonstrated high performance balancing using a recently developed balance theory. The platform can be used to implement many 2D and 3D balancing algorithms, as it has a customizable structure and simple electronics. It can also be used as an educational platform, much like the Acrobot (Hauser & Murray, 1990) or the reaction wheel pendulum (RWP) (Block, Astrom, & Spong, 2007).
- Skippy's mechanical design is governed by solutions and strategies for lightweight, physically robust and affordable robot design. In combination with its optimisation studies, the result is a realistic robot design that can reach a jump height of 3.8 m in simulation when somersaulting or 2.3 m when jumping straight up. Its usefulness as a design exercise for creating high-performance robots is further explained in chapter 9.

In addition, various contributions have emerged from the Skippy Project that can find direct applications in other high performance machine designs. In the past, these have included for example balancing theories (Azad, 2014; Azad & Featherstone, 2016) and the ring screw invention (Featherstone, 2016a). This thesis has the following such contributions:

- An analysis of balance performance and the introduction of a new actuator criterion on the balance performance of simple balancing machines (chapter 4). A criterion of this kind can be used to design robots that need to be good at balancing, and to analyse the physical balance performance of already existing balance machines.
- A generalization of the concept of the Centre of Percussion (CoPc), which can be calculated for any rigid body in a robot for any particular impact scenario (e.g. a foot strike) (section 6.10). The CoPc serves the purpose of placing components that are shock-sensitive, such as the inertial measurement unit (IMU), and could be useful for many types of robots that are subject to repeatedly receiving high impact loads (such as jumping robots).
- Formulae for modelling the force-deflection profile and failure prediction of curved compression leaf springs (chapter 7). This type of spring is generic and usable for any application that requires regressive spring profiles, which is generally desirable for systems that require (1) high control bandwidth for small loads and (2) high energy storing capabilities for high loads. In contrast to previous work, which focusses mainly on generalizing numerical models, this work presents an efficient algebraic approximation that is well suited for use in higher levels of robot design optimisations.

1.4. OUTLINE

This rest of the thesis can roughly be divided in three parts.

Chapters 2 and 3 focus on work that is fundamental to this thesis, as they reflect on previous work and ideas, including the design philosophy, that are relevant for the development of Skippy. Chapter 2 focuses on more general background and theories regarding high performance in legged robot design, and explains the design philosophy, whereas chapter 3 focuses on research related to Skippy itself.

Chapters 4 and 5 focus on balancing and Tippy. Chapter 4 introduces a new actuator design criterion and analyses physical balance performance of a simple balancing machine. Theories on balance control and performance are applied to the real robot Tippy, of which the design and experimental results are discussed in chapter 5.

Chapters 6 to 8 all focus on the design of Skippy. Chapter 6 lays a basis for Skippy's design, by reflecting on the simplified design from previous studies, the design of Tippy and new design studies. The design of Skippy is a multi-faceted project. One of the emerged sub-projects stands out and is treated in a separate chapter, namely chapter 7 on modelling of Skippy's non-linear ankle spring. Skippy's final design is the product of various iterations between computer aided design (CAD) and simplified optimisation studies. The final optimisations and the resulting design are presented and discussed in chapter 8.

Discussions, conclusions and considerations on future developments can be found in chapter 9.

2

DESIGNING FOR HIGH PHYSICAL PERFORMANCE

This chapter focuses on high physical performance in robots. Based on findings in the literature, definitions of high performance are retrieved, comparisons are done between other robots and Skippy, and a design philosophy is devised. The philosophy is the result of reflecting on employed design approaches observed in legged robot design, for robots that ought to exhibit a high level of physical performance, such as running or jumping. The design philosophy is a detailed elaboration of an earlier idea of Dr. Featherstone. The Skippy project serves as a case study for the design philosophy.

2.1. COMMON MEASURES FOR PHYSICAL PERFORMANCE

We pursue the design of a legged robot that exhibits high physical performance. In order to compare the performance of robots and to devise a design strategy, it is a requisite to firstly define physical performance. Various criteria and measures are used in the robotics community to define physical performance. [De Boer \(2012\)](#) proposed that a robot's performance can be evaluated by the criteria energy efficiency, robustness and versatility. Most of the works in legged robotics focus either on energy efficiency or versatility. This is an interesting observation because one could argue that it is more important for a robot to perform its task reliably (robustly), and that performing the task efficiently or being able to perform additional tasks (versatility) is of secondary importance.

In legged locomotion, the typical measure for energy efficiency is one or another form of the cost of transport (CoT), which relates energy consumption to velocity. The CoT has its roots in biology, where it was observed that animals' gait selection be correlated with minimizing the metabolic cost of transport ([Alexander, 1990](#)). Popular particular uses of the CoT in robotics are the mechanical cost of transport (MCoT) ([Collins, 2005](#)) and the total cost of transport (TCoT) ([Bhounsule, Cortell, & Ruina, 2012](#)). In contrast to MCoT, TCoT also considers exchange of energy prior to the generation

of mechanical work, such as Joule heating in the windings of electric motors. Typically, both the MCoT and TCoT also normalize for robot weight. Another variation of the CoT is the cost of hopping (X. Yu & Iida, 2013), specifically dedicated for vertical hopping machines, which replaces velocity in the denominator for the product of jump height and gravitational acceleration.

If physical performance was only quantified by the CoT, then the Cornell Ranger (Bhounsule et al., 2012) would be the best performing robot built to date.¹ However, this robot is neither versatile nor robust.

Versatility refers to the ability of changing gait and actions. It describes a set of skills and is as such not quantifiable by using an interval or a rational scale. Instead, some degree of versatility can be imposed as a requirement on a machine. For legged robots, such a requirement could be the ability to be fully navigational in 3D, including the ability to start and stop moving without the necessity of human intervention. A design criterion that combines this requirement with that of for instance reducing the CoT would be more comprehensive for defining high physical performance than considering the CoT alone. A comparison between high performance robots would then no longer consider robots that have no skill other than being efficient in one particular gait, such as most of the passive-dynamic-based walkers (Bhounsule et al., 2012), including also the aforementioned Cornell Ranger. However, a design criterion of this kind still leaves robustness out of the equation.

Robustness here refers to the ability of a system to recover from an unexpected disturbance, also known as disturbance rejection. As such, it refers to behavioural rather than physical robustness, the latter of which describes a robot's ability to survive impacts (such as those resulting from an unintended crash after failing to recover from an unexpected disturbance). Albeit quantifiable, it is difficult to measure behavioural robustness, and the appropriate measure depends on the application. For a balancing-only machine (i.e. a machine that cannot make steps), such as an acrobot (Hauser & Murray, 1990) or Segway (Segway, 2019), behavioural robustness can be quantified by a measure as simple as a maximum offset angle from which the robot can still recover to a balanced configuration. This subject is further treated in chapter 4. However, measuring behavioural robustness becomes more cumbersome for legged robots that can step. Research has been done on *capturability*, which describes the ability to recover from a disturbance by making one or multiple steps (Koolen, de Boer, Rebula, Goswami, & Pratt, 2012). For limit cycle behaviour (e.g. walking, running) in particular, measures such as the gait sensitivity norm have been devised (Hobbelen & Wisse, 2007) to quantify behavioural robustness. However, their calculation is not straightforward, and as such they are not commonly adapted as a design criterion. An additional complication is that there appears to be a conflict of interest between energy efficiency and robustness and versatility (De Boer, 2012). By having separate and unrelated measures for both energy efficiency and robustness, it is unclear where to set the priority in design studies.

Other common measures that are used for measuring physical performance are jump height or forward velocity (Haldane, Plecnik, Yim, & Fearing, 2016a; Kenneally, 2016). Primarily, these are measures for (explosive) energy output rather than for energy

¹To the authors knowledge, no robot has surpassed the CoT of the Cornell Ranger.

efficiency, as they do not weigh in energy cost, and—as such—are in essence simpler measures than CoT. Recovery from unexpected disturbances is also governed by energetic bursts, suggesting that a measure such as jump height is a better candidate for tackling robustness than a CoT measure. In turn, a CoT measure tackles energy efficiency better than jump height. Nevertheless, record jump heights cannot be achieved by relying only on actuator work, but also require efficient recycling of energy in the system. This implies that the measure of jump height also addresses energy efficiency to some extent.

Since Skippy is to be designed to be a good hopping machine, jump height is the most prominent measure, and forms the basis for core design decisions of the robot. Comparisons between other robots and Skippy are also based on jump height, in section 2.3. To make proper comparisons, some considerations have to be taken into account, such as the time in between jumps and robot size.

2.2. JUMP HEIGHT

Jump height knows various definitions. This section establishes its definition employed throughout the thesis and furthermore relates it to robot size and hopping time.

2.2.1. DEFINING JUMP HEIGHT

We define *jump height* h as the vertical distance that the robot's centre of mass (CoM) travels between the instance of take-off and the *apex*, i.e. the highest point the CoM reaches during the flight phase. This measure can be calculated directly from take off velocity using kinetic and potential energy. However, in some works jump height is defined as the largest vertical distance between the lowest point of the robot during the flight phase and the ground. The latter is easier to observe with the eye, and it has the obvious practical application of measuring the ability to jump over high obstacles. However, unlike the former definition, it always requires information of the actual flight phase: it is essentially a combined measure of (the former defined) jump height and the ability of retracting the legs. If such information is not available, it can be approximated by summing the *stroke length* s to the jump height. The stroke length is the distance the robot's CoM travels from its lowest point, called *nadir*, to take-off. The sum of h and s will be referred to as *total jump height* h_{tot} . See table 2.1 for a complete overview of all terms and their definitions.

2.2.2. JUMP HEIGHT AND SIZE

The performance of jump height can be related to body size. Most works consider jump height as a measure that is independent from robot mass or size (Haldane et al., 2016a), which would imply that a 1 kg jumper (whether that be a robot or animal) that jumps 1 m high is as impressive as a 100 kg jumper that jumps to the same height. The reasoning here is twofold. Firstly, it makes sense from a practical perspective, as all robots will eventually have to be operational in the same (real world) environment. Secondly, a physical perspective shows that jump height does not scale under the core assumptions that (1) mechanical work done by actuators or muscles is proportional to their mass and (2) actuator force is the limiting factor in determining the jump height. A short

Term	Definition
touch down	The point at which the robot makes contact with the ground (resulting from a previous jump or drop).
nadir	The lowest point that the robot's CoM reaches.
take-off	The point at which the robot releases the ground (initiating the jump).
apex	The highest point that the robot's CoM reaches.
jump height h	The vertical distance that the robot's CoM travels from ... take-off to apex;
stroke length s	nadir to take-off;
total jump height h_{tot}	nadir to apex (i.e. $h_{\text{tot}} = h + s$).

Table 2.1: Terminology for vertical jumping, as used throughout this thesis.

explanation on the scaling theory follows below, in addition to explanations of other theories found in literature on scaling of jump height. Corresponding derivations for scaling can be found in appendix A.

SCALING WITH S^0 (NO SCALING)

If a robot and its actuator scale with size S and, as such, their mass with S^3 , then the amount of energy required to reach the same *total* jump height also scales with S^3 , assuming that we can neglect any form of energy loss due to air resistance or slipping. The actuator or muscle force is assumed to scale with S^2 , which is based on the assumption that actuator and muscle force is proportional to their cross-sectional area (perpendicular to the axis of actuation) (Hill, 1950; Wensing et al., 2017). The stroke length of the robot also scales with S , which implies that the total work output scales with $S \times S^2 = S^3$. It follows that the *total* hopping height scales with $S^3/S^3 = S^0$ (i.e. it does not scale). Note that—following the definitions in table 2.1—this implies that the sum of the jump height and stroke length do not scale, and thus that jump height reduces as the jumper increases in size. The effect is not noticeable for small jumpers. Notably, this effect is often unmentioned in literature. However, it is worth noticing that this possibly supports the theory why the largest elephants (6 Mg) and legged robots, such as Tradinno (Zollner Elektronik AG, 2013) (11 Mg), cannot jump at all. To the best of the author's knowledge, the heaviest robot that has demonstrated jumping behaviour is LS3 (Legged Squad Support System, 590 kg) (Boston Dynamics, 2012), which reaches a jump height of several centimetres while trotting, depicted in fig. 2.1b.

SCALING WITH $S^{2/3}$

Some works on jumping of small animals and robots reason that jump height ought to scale with $S^{2/3}$ with relation to body size (equivalent to $S^{2/9}$ with relation to body mass) instead of S^0 for *small* robots and animals (Armour, Paskins, Bowyer, Vincent, & McGill, 2007; M. Gabriel, 1984). For example, it would imply that a jumper that is 8 times as big (ca. 512 as heavy) jumps 4 times as high. This $S^{2/3}$ scaling follows from assuming actuator power to be the limiting factor rather than actuator work or force, which would be the case for jumpers of sufficiently small size. In particular, this power scaling assumes that




Figure 2.1: Four autonomous quadrupeds of different size: Tradinno (Zollner Elektronik AG, 2013), LS3 (Legged Squad Support System) (Boston Dynamics, 2012), Cheetah 2 (Park, Park, & Kim, 2015), and Minitaur (Ghost Robotics, 2016). Of these robots, the two smallest have shown the highest and approximately equal physical performance in terms of jump height ($h \approx 0.5\text{m}$) and vertical jumping agility ($v \approx 1.2\text{ms}^{-1}$). LS3 has demonstrated only small jumps of several centimetres, while trotting (depicted above, in flight phase). Tradinno has not demonstrated any type of jumping behaviour. This pattern is in line with S^0 scaling (section 2.2.2) compensated for by change of stroke length.

the product of force and velocity is limited, rather than being velocity-limited in addition to being force-limited. A full explanation and derivation can be found in appendix A.

However, it is doubtful whether measuring jumping performance is done best by compensating jump height h for robot size according to $S^{2/3}$, as jumpers are only power-limited if they are of sufficiently small size. Notably, the $S^{2/3}$ scaling for small jumpers is used as a possible explanation why small jumping animals (weighing less than ca. $1 \times 10^{-1}\text{g}$)—such as locusts—employ a different type of jumping mechanism to perform high jumps. These mechanisms are based on principles of pre-harvesting energy to generate short but high peak power bursts, and is further treated in section 2.2.3.

SCALING WITH S OR S^3

A few other works normalize jump height for robot size (S) or even to robot mass (S^3) to use it as measure (Kovač, Fuchs, Guignard, Zufferey, & Floreano, 2008; Kovač, Schlegel, Zufferey, & Floreano, 2009; Nguyen & Park, 2012). No scientific justification or derivation was found for these measures, and instead seem to unjustified favour small robots. As such, these measures are not considered in this thesis.

2.2.3. JUMP HEIGHT AND TIME

Roughly speaking, jumping robots employ one of two following operating principles: *single-leap* and *continuous-hopping*. Single-leap robots generally jump higher, but their performance is commonly conceived to be lower. By also taking the time required to jump into account, a fairer comparison can be made regarding the performance.

SINGLE-LEAP ROBOTS

Single-leap robots release pre-stored or pre-harvested energy explosively to achieve their highest hop. An example of a robot with pre-stored energy is Boston Dynamic's Sandflea (Boston Dynamics, 2012), depicted in fig. 2.2a, which carries a CO2 or propane cylinder and piston as a secondary on-board energy source specifically dedicated for jumping. Other single-leap robots pre-harvest potential energy by using a catapult mechanism,




Figure 2.2: Jumping robots SandFlea (Boston Dynamics, 2012), EPFL Jumper (Kovač et al., 2009) and Salto-1P (Haldane, Yim, & Fearing, 2017). All of these robots can perform multiple jumps without physical human intervention, but only Salto-1P is a continuous-hopping robot.

which charges an elastic element at a relatively slow rate, and then releases it in a short burst. This principle is bio-inspired; it can be found in nature in for instance grasshoppers and locusts (Armour et al., 2007). The principle is an extreme example of *power modulation*, where energy is temporarily stored and subsequently released with a higher power output than actuators can provide on their own. Single-leap robots that make use of such a mechanism have been reported to jump up to heights of approximately 3.4 m (Haldane et al., 2016a).

CONTINUOUS-HOPPING ROBOTS

Continuous-hopping robots have more powerful actuators or use an elastic element to recycle a portion of the kinetic energy of the previous hop into the next one, or both, in order to jump high. Skippy is also a continuous-hopping robot, and has both a powerful actuator and elastic elements for power modulation. Prior to initiation of this thesis project, continuous-hopping robots had been able to reach jump heights of approximately 0.6 m (Haldane et al., 2016a), which is significantly lower than that of single-leap hopping robots.

VERTICAL JUMPING AGILITY

The jump height of single-leap robots has outreached that of continuous-hopping robots. However, single-leap robots do not require making controlled jumps, as they do not need to make a proper landing. The requirement of pre-harvesting energy to jump furthermore does not make them capable of responding to an unexpected disturbance. In addition, they are typically slower and do not recycle energy from previous jumps. As such, their physical performance is generally conceived to be lower than that of continuous-hopping robots.

Recently, a measure called the *vertical jumping agility* has been introduced, which considers both the jump height and the time required to make a jump (Haldane, Plecnik, Yim, & Fearing, 2016b). It is calculated as $v = h/(t_s + t_a)$, where t_s is the stance time and t_a the apogee time or half the flight time (time in flight from take-off until reaching apex). Note that the definition does not include gravity; instead it assumes an Earth-like

gravitational environment. A higher value corresponds to a better jumping performance. In effect, this measure penalizes single-leap robots for the longer stance time that is required to prepare a jump. More general, the measure is useful to compare robots (and animals) that cannot be easily classified as either single-leap or continuous-hopping robots.

The vertical jumping agility measures an average of displacement over time, with the unit of velocity (ms^{-1}), which are also the physical units of power-to-weight ratio and velocity. As such, it suggests that it is comparable to the performance measure of steady forward velocity. However, the core difference is that jumping is governed by high fluctuations of velocity (which changes sign), whereas horizontal movement can be achieved by maintaining a constant velocity. This implies that a robot with a certain value for vertical jumping agility requires more power regulation and higher peak power output than a robot that has an equal value for steady forward velocity.

Note that gravity imposes a limit on the flight time, which increases as the jump height increases. In fact, h and t_a are directly related through $h = gt_a^2/2$, with g the gravitational acceleration. As such, we can write

$$v = \frac{h}{t_s + \sqrt{2h/g}}. \quad (2.1)$$

In an ideal scenario where the stance time t_s is zero, we find $v \propto \sqrt{h}$, which states that the maximum vertical jumping agility scales with the square root of the jump height.

Note furthermore that the vertical jumping agility should not be confused with the *vertical specific agility* (Kenneally, 2016), which is the measure of jump height corrected for by gravity, calculated as hg . Since the comparisons presented below only consider terrestrial robots, this measure will not be used.

2.3. ROBOT COMPARISONS

Notable work on continuous-hopping robots is Raibert's early work in the 1980s (Raibert, 1986), which demonstrated several legged robots that could hop, run and make somersaults in 3D, while using simple designs and control systems. The robots are all tethered, and the highest hops of these robots reach several decimetres. Despite not very high, many of the robots achieved high horizontal velocities and were fully navigational in 3D. Their capabilities were deemed a large improvement in robotic physical performance for their time, and only few machines have been able to equal or surpass their performance since.

By the beginning of this thesis project, the robot that was reported to have the highest vertical jumping agility is Salto, with $v = 1.7 \text{ms}^{-1}$ and $h = 1 \text{m}$ (in 0.58 s) (Haldane et al., 2016b). The initial versions of Salto lacked measures to stabilize it in 3D, and as such, was not able to demonstrate long repetitive jumping behaviour. More recent developments on Salto have demonstrated this instead. These include the development of Salto's successor Salto-1P (Haldane et al., 2017), depicted in fig. 2.2c, which has demonstrated long repetitive and precise jumping in 3D (Yim & Fearing, 2018). This has been made possible by the addition of thrusters. Salto-1P also has an improved jump performance of $v = 1.83 \text{ms}^{-1}$ and $h = 1.25 \text{m}$. However, the results from Yim and Fearing (2018) indicate

that Salto-1P is still not fully autonomous, as it relies on a motion capture system for doing some of its experiments.

Albeit not designed as a hopping robot, the second most agile reported robot is the quadruped Minitaur (Kenneally, 2016), shown in fig. 2.1d, with $v = 1.2 \text{ m s}^{-1}$ and $h = 0.5 \text{ m}$. Minitaur has demonstrated full autonomy in 3D. Other examples of robot with high physical performance are MIT's series of Cheetah robots (Park, Park, & Kim, 2015; Park, Wensing, & Kim, 2015; Seok et al., 2015), of which Cheetah 2 is shown in fig. 2.1c. It has demonstrated a vertical jumping agility of $v = 1.1 \text{ m s}^{-1}$ with $h = 0.5 \text{ m}$ and running gaits with high steady forward velocities of 4.5 m s^{-1} .²

A few examples of single-leap robots are the EPFL Jumper (Kovač et al., 2009), Glumper (Armour et al., 2007) and TAUB (Zaitsev, Gvirsman, Ben Hanan, Ayali, & Kosa, 2015), which are able to jump up to heights of 1.4 m, 1.6 m and 3.4 m, respectively. TAUB in particular is not fully navigational in 3D, as it requires human intervention to reset it after jumping. The EPFL Jumper and Glumper also have the ability to recover and reorient themselves after landing. The EPFL jumper in particular uses a light-weight roll cage that enables it to passively realign itself after a fall, and double-functions as protection, as shown in fig. 2.2b. A few robots combine jumping with another type of locomotion, such as the jumping-crawling robot JumpRoACH (Jung, Casarez, Jung, Fearing, & Cho, 2016) and the jumping-gliding robot MultiMo-Bat (Woodward & Sitti, 2014), both able to jump up to heights of approximately 3 m. Of these robots, the most agile robot is the EPFL jumper, with a vertical jumping agility of $v = 0.34 \text{ m s}^{-1}$.

For comparison, a visual representation of the vertical jumping agility of various of the above mentioned jumpers is shown in fig. 2.3, which shows the height that each jumper would reach after consecutive jumping for a duration of 4 s. The graph is amended from fig. 1 in (Haldane et al., 2016b) with data of Salto-1P and the above mentioned hypothetical data of Skippy. It also includes data of a bushbaby (*galago senegalensis*), a high performance animal jumper.

COMPARISONS TO SKIPPY

Ballpark figures for Skippy assume a jump height of 4 m following an estimated stance time of ca. 0.2 s, as further explained in section 3.1. These numbers would imply a vertical jumping agility of $v = 3.6 \text{ m s}^{-1}$. However, it is expected that Skippy cannot make a lengthy sequence of 4 m hops. Nevertheless, even if it jumped half as high with double the stance time (this assumes an identical stroke length), so that $h = 2 \text{ m}$ and $t_s = 0.4 \text{ m}$, the vertical jumping agility would be $v = 1.9 \text{ m s}^{-1}$, which is still slightly higher than that of Salto.

If we analyse jumping performance by also correcting jump height for robot size based on limited actuator power (see section 2.2.2), then Salto becomes more competitive. By comparing Salto's mass (100 g) with the anticipated mass of Skippy (2 kg) (Driessen, 2015), it is suggested that the jump height of Salto would have to be multiplied with $20^{2/9} \approx 2$ to compare it with that of Skippy. However, it remains doubtful that power should be considered the main limiting factor for robots up to sizes as (relatively) large as Skippy (ca. 2 kg).

²The original Cheetah (Seok et al., 2015) has demonstrated even higher steady forward velocities of up to 6 m s^{-1} , but it is tethered.




Figure 2.3: Visual representation of the vertical jumping agility v of various jumpers. The graph has been copied from Haldane et al. (2016b), and has been amended with data of Salto-1P (Haldane et al., 2017) and hypothetical data of Skippy if it were jumping 4 m and 2 m high. The number of arrows correspond to the number of repeated jumps over an interval of 4 s, and the length of each arrow is proportional to the jump height h of one jump.

Regarded as a velocity, $v = 3.6 \text{ ms}^{-1}$ is lower than the steady forward velocity 4.5 ms^{-1} of Cheetah 2 (Seok et al., 2015). However, as explained in section 2.2.3, the two are difficult to compare due to high fluctuations of velocity required for jumping. In comparison, Skippy must reach a vertical lift-off (peak) velocity of 6.3 ms^{-1} or 8.9 ms^{-1} in order to make a 2 m or 4 m high hop respectively.

2.4. DESIGN PHILOSOPHY

Following the analysis on defining high physical performance and the comparisons of several high performance legged robots, this section elaborates more on the *design philosophy* to realise physically high performance robots by reflecting on state-of-the-art robots. The presented study arose from the observation that robots significantly underperform compared to their biological counter-parts (i.e. humans and other animals), despite decades of research. A possible reason could be lack of powerful building blocks in technology, such as high power density actuators. This has been the motivation for some researchers to focus the design strategy on creating more powerful actuators, such as on liquid cooled actuation (Kim, Campbell, Ahn, Sentis, & Paine, 2017; Paine & Sentis, 2015). Also SCHAFT, the robot that won the 2013 DARPA trials, makes use of liquid cooled actuation (Guizzo & Ackerman, 2014). However, other researchers claim that lack

of power density is not the main issue (Gaponov, Palli, May, & Ryu, 2016; Seok et al., 2015). After all, a simple comparison between datasheets shows that electric motors are available with a higher power-to-weight than biological muscles, with a factor of at least 10 (Gaponov et al., 2016), and that artificial materials exist that are stronger and harder than bone. The main difference between electric actuators and muscles is that the former has a much lower torque density, which calls for the requirement of transmission mechanisms. In particular, Gaponov et al. (2016) claims that transmissions are the main issue. However, whereas this may partly address the issue, it cannot be the full explanation: a (zero backlash) transmission with a mass approximately equal to that of its motor and an efficiency of 70% (Harmonic Drive AG, 2014) would reduce the actuator's power density with approximately a factor of 3, so that the power density is still in favour of the artificial actuator compared to that of the muscle. This suggests that there must be also other reasons that explain lack of robot physical performance.

Instead of blaming slow progress and the lack of robot performance on the unavailability of high performance parts, an explanation can also be sought in the currently employed approaches for robot design, which are likely suboptimal. This is investigated by analysing designs and performance of existing robots, with the purpose of establishing a design philosophy that is intended to aid in designing robots that exhibit higher level of physical performance and that is to be employed on the case study of Skippy.³

The author would like to exemplify that the presented design philosophy is hypothetical, and not based on scientific evidence. The philosophy comprises a list of key points that are believed to positively influence high performance robot design. In an attempt to substantiate the claims of the key points, several robot designs from literature are mentioned. It is not the intention of the author to compromise the valuable contributions of these designs. On the contrary, over the years impressive and pioneering work has been carried out on very different robot designs and approaches, and as a result, insight can be gained into how further improvements can be made possible.

2.4.1. COMPLEXITY

The main message of the design philosophy is that progress in robotic research should focus on making simple high-performance designs, and then proceed to more complex designs, rather than attempting to make already complex machines exhibit high performance. A robot's desired versatility can be used as an argument for complex designs. However, difficulty of robot design is easily underestimated, and the end result is a robot that is slow and not robust nor energy efficient: despite being designed to perform many different tasks, it is not able to execute any of them with satisfaction.

A robot's complexity is not easily defined, but a reasonable indication can be provided by counting the number of actuated degrees of freedom (DoFs) and sensors. Notably, most of the robots that have exhibited high levels of performance, have a low number of actuators, and are typically the results of extensive iterative design of versions with fewer or simpler actuators. These high performance robots include the previously mentioned high-performance Salto-1P, Cheetahs, Minitaur and SpotMini, which do not have

³*Acknowledgements* — The basis of said design philosophy has firstly been postulated by Dr. Featherstone. In this chapter of the thesis this idea is further devised through said comparisons between existing robots and the introduction of key points.

more than 12 actuated DoFs each. Remarkably, they are also relatively small and fully autonomous. In particular, Cheetah’s team claims that part of their success is due to employing a method of *proprioceptive* sensing in combination with high fidelity actuation (Wensing et al., 2017). This method requires fewer sensors as it relies on current sensing and modelling actuator characteristics. In interesting contrast, a reflection on the robots competing in the DARPA challenge concludes that robots—which were already quite complex—need more sensors (Atkeson et al., 2016). The robots that compete in this challenge are typically complex in having many actuated degrees of freedom, and significantly under-perform compared to the aforementioned robots. Yet, they are the products of big and professional teams and high budgets, so the blame cannot be put on the lack of resources or experience. A few examples of other expensive robots that have not demonstrated high physical performance, and are yet the result of long and expensive projects, are Valkyrie (129 kg, 32 DoFs excluding hands, 44 DoFs including hands) (Radford et al., 2015), Hydra (130 kg, 31 DoFs excluding hands, 41 DoFs including hands) (Kaminaga et al., 2016; Ko, Kaminaga, & Nakamura, 2018) and WAREC-1 (155 kg, 28 DoFs) (Hashimoto et al., 2017), shown in fig. 2.4.



(a) Valkyrie (44 DoFs)

(b) Hydra (130 kg, 41 DoFs)

(c) WAREC-1 (155 kg, 28 DoFs)

Figure 2.4: Examples of expensive and complex field robots that have not yet demonstrated autonomy or high physical performance: Valkyrie (Radford et al., 2015), Hydra (Ko et al., 2018) and WAREC (Hashimoto et al., 2017).

“Reducing complexity” is easier said than done. To further underpin the design philosophy, several key points have been devised, typically related to complexity in one way or another. The key points are results from direct comparisons in literature, and are aimed to aid in creating high performance designs whilst avoiding or dealing with complexity.

2.4.2. KEY POINTS

An overly complex design of a machine is typically (1) non-holistic, (2) poorly targeted, and (3) decoupled from its designed behaviour, which leads to unsupported decisions and loss of transparency. A popular answer to support decision making is favouring (4) bio-inspired design over technology-inspired design, but this has in turn shown proneness to being a trigger for even more complicated designs.

The few more complex robots that have also shown high degrees of physical perfor-

mance all have a rich history of (5) iterative design governed by trial-and-error. Whereas iterative design studies can be tedious, expensive and unresolving, they have also proven to be productive for robots that are (6) designed with good risk management, (7) fully autonomous, (8) physically robust, (9) small, and (10) not optimal.

The ten key points are summarized below and each are treated separately in the remainder of this section:

- Robot designs should be holistic.
- Robot designs should be targeted on behaviour at the performance envelope.
- Robots should be a product of machine and behaviour co-design.
- Technology inspired design has priority over bio-inspired design.
- Progress is made through trial-and-error by making prototypes.
- Robots should have a reduced requirement for risk management.
- Robots should be designed to be physically robust.
- Robots should not be designed to be bigger than necessary.
- Robots should be designed to be fully autonomous, so that they do not require an umbilical, tether or harness.
- Robots should be designed to be good enough rather than optimal.


HOLISTIC DESIGN

In a holistic design all individual parts are intimately interconnected and explicable only by reference to the whole. It can be observed in designs of complex systems that decisions on a lower level of the design are done separately, rushed over, or simply neglected, whereas these choices can have major influence on the final performance.

On the lowest level, these decision could concern material choice and spring profiles. For instance, the most popular choice for mechanical springs are steel helical springs, which have a linear profile and a low gravimetric energy density. The choice behind choosing a linear profile is typically not explained, but it is in many cases suboptimal, as further explained in chapter 7.

On a higher level, these decisions concern copying of (i.e. using identical) actuators, or entire legs, despite having different load profiles. The crawling quadrupedal WAREC-1 uses only three different actuator units in a 28-DoF robot, consisting of 4 identical 7-DoF series arms (Hashimoto et al., 2017). The humanoid M2V2 (J. Pratt & Krupp, 2008) and the crawling quadruped RoboSimian (Hebert et al., 2015), shown in fig. 2.5b, use only one type of actuator for their respective 12 and 28 actuated DoFs, whereas the various actuated joints have different loading profiles. A study on limb optimisation of quadrupeds moreover suggests that the front and hind legs are to be shaped differently (Ha, Coros, Alspach, Kim, & Yamane, 2016) Nevertheless, it is common practice to copy the front and hind legs in quadrupeds, even if the robot is not symmetric about its frontal plane, such as in HyQ(2Max) (Semini, 2010; Semini et al., 2017) and the centaur-like robot CENTAURO (Project Centauro, 2014), the latter of which is depicted in fig. 2.5b.

More generally, we find examples of sub-assembly designs that were designed independently from or prior to the design of their final application (i.e. a legged robot),



(a) ATRIAS: an example of holistic design. (b) CENTAURO and RoboSimian: examples of non-holistic designs.

Figure 2.5: Examples of holistic and non-holistic design. ATRIAS' (Hubicki et al., 2018) entire design is based on a SLIP-like model, and is therefore holistic. Non-holistic design is observed in CENTAURO (Project Centauro, 2014) and RoboSimian (Hebert et al., 2015): CENTAURO uses identical pre-designed front as hind legs, despite the asymmetric upper-body design, and RoboSimian uses the same actuator for all its 28 joints, despite the different loading profiles.

such as found in the design of TORO (Englsberger et al., 2015; Ott et al., 2010), which is built from units of Kuka arms. Popular also is the design of series elastic actuators (SEAs) (Paine & Sentis, 2012; Robinson, Pratt, Paluska, & Pratt, 1999) and variable stiffness actuators (VSAs) (Groothuis, Rusticelli, Zucchelli, Stramigioli, & Carloni, 2014; Tonietti, Schiavi, & Bicchi, 2005) prior to knowing their use. Especially VSAs are known to be complicated and have very detailed characteristics, since they require at least two actuators per DoF (Vanderborght et al., 2013; Vanderborght, Van Ham, Lefeber, Sugar, & Hollander, 2009), and as such commonly end up not being used at all.

Two examples of holistic design of high-performance robots are ATRIAS (Hubicki et al., 2018, 2016), shown in fig. 2.5a, and Cheetah (Seok et al., 2015). The entirety of ATRIAS' design is tailored to closely represent a model similar to a spring-loaded inverted pendulum (SLIP) (in particular, ATRIAS stands for 'Assume The Robot Is A Sphere') (Blickhan, 1989; Schwind & Koditschek, 1997; Seyfarth, Geyer, Günther, & Blickhan, 2002). Cheetah's entire design, including its legs (Wensing et al., 2017), are based on mimicking a cheetah. Efficient bipedal robots based on passive-dynamic walkers (Collins, 2005; De Boer, 2012) are also examples of holistic design. All of these robots exploit the machine's natural dynamics to increase performance. However, their designs typically have an overfocus on energy efficiency, at the cost of sacrificing on overall performance by lacking versatility and robustness.

DESIGN TARGET AT THE PERFORMANCE ENVELOPE

A robot design should be tailored to the most extreme behaviour it is expected to deliver, so that its natural dynamics at its *performance envelope* coincide with that behaviour. The performance envelope of a design refers to its capabilities as physically limited by physical quantities, such as maximum speed, load, power, current and temperature. A robot is said to operate in saturation while one or multiple of its physical quantities are at their respective physical limits. The closer physical quantities are at their limits and

the longer the robot operates in saturation during its most extreme behaviour, the closer it is to achieving physical maximum performance.

It is not necessary to aim for limiting all physical quantities at the same time; those that limit peak behaviour typically depend on the type of application. For example, the performance of a high jumper (short peak behaviour) is likely limited by peak power, velocity, current and torque, whereas the main performance delimiter of a marathon runner (long peak behaviour) is likely temperature (or otherwise nominal current). Nevertheless, if other physical quantities stay far below their respective maxima, it is an indication that the design is unbalanced (i.e. some of its parts are over- or under-engineered, introducing a weakest link), further limiting the performance.

Nevertheless, robot design is typically neither directly focussed on maximizing physical quantities (i.e. on exploiting the performance envelope), nor on extreme behaviour. Instead, target performance in design is focussed on other imposed requirements such as stiffness, energy efficiency are minimizing joint torque. The latter in fact suggests the opposite of exploiting the performance envelope and is related to desires and conventions of controllers and low-level software, which actively avoid saturation and further limit physical performance.

Target performance Suboptimal design decisions can be the result of a missing or illogical choice of target performance that (1) is not directly linked to maximizing physical quantities (i.e. exploiting the performance envelope) and (2) does not address the most extreme behaviour the machine is expected to deliver, or any behaviour at all. Exemplary for the latter is the use of popular performance measures without providing a justification. For example, up until the revolutionary refutation of the notion “the stiffer the better” by [G. A. Pratt, Williamson, Dillworth, Pratt, and Wright \(1995\)](#), most legged robots were designed to have an as-stiff-as-possible actuator train, despite having to interact with humans and the environment. Examples of such robots include Honda’s line of robots ([Hirose & Ogawa, 2007](#)).

A popular choice in present-day robot design is the focus on energy efficiency. Whereas energy efficiency does not directly address physical quantities, it is a desired property for robots that are to exhibit long-lasting behaviours (e.g. marathon runners), and thus a reasonable starting point as a target performance.⁴ Nevertheless, the requirement of energy efficiency is in some cases over-ambitiously applied at the expense of performance, especially for cases where the desired most extreme behaviour is described by bursts (e.g. sprints or jumps).

There are many examples of robot design where the core focus is energy efficiency ([Fujimoto, 2004](#); [Guenther & Iida, 2017](#); [Haeufle, Taylor, Schmitt, & Geyer, 2012](#); [Hurst & Rizzi, 2008](#); [Papadopoulos & Cherouvim, 2004](#); [X. Yu & Iida, 2013](#)). Examples also include the aforementioned passive dynamic-based walkers. The cheetah-inspired robot Cheetah ([Seok et al., 2015](#)) is also designed for improved energy efficiency (by reducing the TCoT), which might be surprising, as the cheetah is an animal known for its high peak

⁴However, the proposed design philosophy would not take energy efficiency as a target performance. For example, the most likely physical limitation of a marathon running robot is maximum temperature. If it were designed properly so that it can finish the marathon and set records without overheating, then consequently it must have been energy efficient.

velocity (requiring peak power output) rather than for its continuous/steady velocity (requiring high efficiency).

Another used design target is joint torque reduction, which has found its way to kinematics optimisations (Ha et al., 2016). A design criterion of this kind does exactly the opposite of exploiting the performance envelope and is thus not practical as a design criterion for maximizing physical performance. Criteria such as joint torque reduction stem from control theory, which further limit physical performance, explained in more detail below.

Avoiding saturation If the robot is really at its performance envelope, then at least one actuator is saturated; but conventional control systems require that the actuators are never saturated, as many controllers fail when reaching joint or torque limits (as there is no room for over-compensation). This has led to various strategies to monitor and constrain joint torques to values (substantially) below their saturation limits (Orsolino et al., 2018). Whereas these strategies might be promising to solve control issues, they do not contribute to achieving maximum performance. In addition, many of the saturation limits as defined in low-level software are typically lower than the actual saturation limits, due to the implementation of multiple layers of (safety) margins, further distancing the robot from operating at its performance envelope. For example, these include no-go zones around end stops and conservative power and velocity limits. This issue overlaps with the issues of risk management, discussed further below.

MACHINE AND BEHAVIOUR CO-DESIGN

Machine and behaviour co-design implies that only the final action is specified, not the action sequence to achieve it. Both the system and action sequence are then to be found coherently, such as through optimisations.

Machine and behaviour co-design is associated with design freedom as it aims to reduce presumptions on the design prior to its completion. The increased design scope inevitably makes design substantially more challenging, especially for more complex robots. As such, it is more common to either (1) design a system for a desired behaviour (machine-to-behaviour design), (2) find the system's behaviour for a given system (behaviour-to-machine design), or (3) to have the design of both the machine and behaviour be decoupled (decoupled design). The associated reduced design scope facilitates the design task, but inevitably sacrifices on physical performance.

Behaviour-to-machine and decoupled design Behaviour-to-machine and decoupled design also imply that the machine's design is non-holistic. The application of forward dynamics optimisations and optimal control on pre-designed systems are examples of *behaviour-to-machine* design rather than of machine and behaviour co-design, as they focus on finding the action sequence for a given machine (Stelzer & von Stryk, 2005).

Decoupled design implies that neither the machine's natural behaviour nor the machine itself have been designed to fit the other, which is associated with imposing artificial constraints and dynamics on pre-designed machines. A motivation to do so is simplification of the control issue, at the expense of physical performance. An extreme example in legged locomotion is the use of controllers that aim for static ability, where

the robot is artificially constrained to move quasi-statically while keeping its CoM above the support polygon. Other motivations are desire for theoretically-sound control (e.g. mathematically provable stability), which has led to the development of for instance virtual constraints controllers (Chevallereau et al., 2003; Poulakakis & Grizzle, 2007, 2009).

Machine-to-behaviour design ATRIAS and Cheetah, previously mentioned examples of holistic design, are both cases where the machine is designed to obtain a predefined behaviour, and are thus examples of machine-to-behaviour design. In ATRIAS and its precursor MABEL (Sreenath, Park, & Grizzle, 2013), the machine is designed to resemble the behaviour of a SLIP-like model. Whereas this is a reasonable basis because the SLIP model does well what it's supposed to do (i.e. natural walking or running), a real machine's dynamics cannot coincide exactly with that of a SLIP model. In order for it to be a better example of machine and behaviour co-design, the model and its behaviour would have to be updated coherently with the machine's design, towards the common objective of walking or running.

In Cheetah, the machine's design is based on that of a cheetah and its behaviour, whereas the machine and the animal cannot possess the exact same characteristics due to differences in building blocks of technology and nature, which implies that also the natural behaviour of the machine will be different from that of the animal. This is further explained in the next key point on bio-inspired design.

Co-design Legged robot designs in which both the machine (e.g. limb lengths, mass distributions) and behaviour are coherently optimised have become increasingly popular only recently (Digumarti, Gehring, Coros, Hwangbo, & Siegwart, 2014; Ha et al., 2016; Spielberg, Araki, Sung, Tedrake, & Rus, 2017). A possible reason for this is the increased availability of optimisation software and computational power that are required for the complexity of the optimisation problems. *Evolutionary robotics* is an umbrella term that encompasses the field of robotic research that focuses on iterative optimisation approaches where an optimal behaviour, machine design or both (co-design) is found. A literature study has recently been published on this topic by Prabhu, Seals, Kyberd, and Wetherall (2018).

Ha et al. (2016) demonstrated performing optimisations of a 12-DoF quadruped in under a few hours on a household desktop computer. The model used here is still simplified, for being planar and using simple actuation characteristics; it nonetheless demonstrates the feasibility of the technique. Applying such (global) optimisations on more complex robot designs (e.g. 40-DoF humanoids) might still be intractable (and are as such a reason to be avoided) due to significantly increased (1) computational complexity and (2) difficulty of defining optimisation constraints.

Early work on Skippy's design (Driessen, 2015) is furthermore an example that machine and behaviour co-design not need be an outcome of optimisations necessarily, and that it is possible to find a feasible solution for a high performance design by having a general understanding of dynamics and power modulation. Canaday, Zapolsky, and Drumwright (2017) furthermore present a method of automated iterative machine and behaviour co-design, which does not rely on demanding optimisation problems. The

outcome is a machine that is good in what it is supposed to do, but it is not necessarily globally optimal. This is perfectly acceptable: it should not be an endeavour for the system to be globally optimal, especially not if it would hinder progress of design by trial-and-error. Additional explanations are provided under the last key point.

TECHNOLOGY-INSPIRED PRIORITISES BIO-INSPIRED

A common solution to the many choices that have to be made in a complex design, is to apply biological methods and systems found in nature to the design of robots, also known as bio-inspired design or bionics. Since nature has undergone millions of years of iterative optimisation (i.e. evolution), this is not a bad starting point per se. After all, various high performance robots such as Cheetah, are the products of bio-inspired designs. Also Spring Turkey and Spring Flamingo (J. E. Pratt, Chew, Torres, Dilworth, & Pratt, 2001) are bio-inspired and have exhibited a high level of performance for their time, despite being planar robots. Nature is very useful as a source of inspiration, and we can learn a lot from her.

However, technology has different building blocks from nature (Egan, Sinko, Leduc, & Keten, 2015), implying that not every principle from nature is ideally suited to engineering (Ha et al., 2016). For instance, nature has excelled in making muscles scalable and in having the nervous system synergize hundreds of (small) muscles and senses. On the contrary, electric motors are not very scalable (bigger motors have a higher power density and efficiency), require transmission mechanisms, are less physically robust and each requires its own control loop. Yet, they generally have a higher power-to-weight ratio than biological muscle, and have other advantages in offer that nature has not, such as unconstrained rotation and higher bandwidth. In addition, the products of nature have evolved for more extreme environments than the relatively highly controlled environments of most robots, characterized by survival and indefinite richness (Pfeifer, Lungarella, & Iida, 2007). Animals have the additional complex requirement that they must be self-reproducible and evolvable and growable from a single cell, which machines do not. These factors can contribute to the temptation of designing robots to be more complex than necessary.

Number of actuators As related to nature's ability to easily scale and synergize many muscles, we find that many joints in animals are over-actuated, comprising multiple muscles that form antagonistic pairs and that enable bi-articulation. Note that the presence of antagonistic pairs is a necessity since muscles are—unlike most artificial actuators—asymmetric, being able to exert force only in one direction. Antagonistic pairs also allow for stiffness modulation through co-contraction, which has been an inspiration for robotics researchers to design previously mentioned VSA's (Vanderborght et al., 2013). These comprise at least two actuators per actuated degree of freedom, and as such are governed by complex designs. The resulting designs are prone to performance reduction due to (1) poor scalability of artificial actuators (2) complicated non-holistic design (more parameters to tune) (3) complicated control (more actuators) or (4) introduced fragility due to extra parts.

A popular family of robots with many actuated DoFs are humanoids, which—per definition—are bio-inspired. An example of a humanoid that is built to mimic a hu-

man in both morphology and actuation is Robot Athlete (Niiyama & Kuniyoshi, 2010; Niiyama, Nishikawa, & Kuniyoshi, 2010), which makes use of bi-articulated and antagonistically paired artificial pneumatic (McKibenn-type) muscles (6 actuators per 2 joints). While being powered through an umbilical and protected by an external harness, to the author's knowledge, the robot has not yet demonstrated a stable running gait in 3D. Humanoids with similar actuation include the bipedal robot introduced in (Kaneko, Ogata, Sakaino, & Tsuji, 2015) and Kojiro (Mizuuchi et al., 2008), both of which have not demonstrated autonomous locomotion.

Actuator characteristics Torque-velocity profiles and heating characteristics of artificial actuators are different from that of biological muscle. As such, the expected ideal behaviour of a machine is different from that of an animal, even if the animal had similar kinematics and dynamics. This was also observed by Seok et al. (2015), who report that differences caused by behaviour in Cheetah (the robot) and the cheetah (the animal) could be significantly influenced by the differences in characteristics of an electric motor and biological muscle. In addition, most joints in animals that appear to be revolute, such as knees and elbows, have significantly different kinematics; and many even exhibit load-dependent variation in kinematic properties, as caused by the relatively soft tissues (e.g. cartilage) found in animals compared to the hard (e.g. chrome steel) bearings found in machines.

Unconstrained rotation An example where technology has excelled in compared to nature is unconstrained (infinite) rotation, which has led to the invention of wheels and vehicles that reach maximum velocities exceeding that of any animal. Wheels have found their way to legged robotics too. (X-)RHex (Altendorfer et al., 2001; Galloway et al., 2010) and Quince (Nagatani et al., 2011) are examples of technology-inspired yet simple designs that combine legged with wheeled or tracked locomotion, which have both demonstrated the ability to cross rough terrain and climb stairs. A recent example of a high performance legged robot with wheels is Boston Dynamic's technology-inspired Handle (Boston Dynamics, 2017). Other examples of legged robots with wheels are CENTAURO (Project Centauro, 2014) and a wheeled version of ANYmal (Bjelonic et al., 2019).

Infinite rotation has also enabled the invention of thrusters, reaction wheels and control moment gyroscopes. Thrusters are used in Salto-1P for stabilisation and in-flight control (Haldane et al., 2017). Reaction wheels and control moment gyroscopes find many applications for stabilisation. For instance, reaction wheels can be found in the balancing robots Cubli (Gajamohan, Merz, Thommen, & D'Andrea, 2012; Muehlebach & D'Andrea, 2017) and Murata Boy (muRata, 2015), and control robot gyroscopes in walking assist devices (Matsuzaki & Fujimoto, 2013; Vallery, Bögel, O'Brien, & Riener, 2012).

DESIGN BY TRIAL-AND-ERROR

A majority of projects that rapidly lead to successful high performance robots have a rich history of rapid iterative design through trial-and-error and prototyping. An argument against design by trial-and-error is that it costs additional resources and time to build

and test prototypes. However, it appears that the experience gained from this strategy pays off for the inherent complexity of legged robotics, as the teams behind a majority of robots that are ranked highest in physical performance have employed it. In addition, it should be recalled that also simulation studies have a cost (labour, time and computing equipment).

These include the projects on UC Berkeley's Salto, OSU's ATRIAS and Cassie, MIT's Cheetah and all of Boston Dynamic's robots. Salto's team has a history on experimenting with structural building blocks for small robots (smart composite micro-structures), then proceeded to building a leg mechanism to jump vertically with zero angular momentum at take-off (Haldane et al., 2016a; Plecnik, Haldane, Yim, & Fearing, 2016), prior to building various versions of Salto, each with added functionality (Haldane et al., 2017). Salto is proof that iterative design not necessarily be more expensive. The high performance bipedal Cassie (Agility Robotics, 2018) is a commercialized spinoff of ATRIAS (Hubicki et al., 2018, 2016), which has undergone various iterative design cycles (Grimes & Hurst, 2012; Hubicki et al., 2018, 2016). ATRIAS was furthermore preceded by several other walkers and hoppers, including MABEL (Grizzle, Hurst, Morris, Park, & Sreenath, 2009) and RABBIT (Chevallereau et al., 2003). Boston Dynamics' robots all have a history of design by trial-and-error, and refer to the design strategy as "make it, break it, fix it", claiming it to be a key factor in their success.


Also Honda's ASIMO has a rich history of iterative design, which has spanned more than two decades and has resulted in at least 12 bipedal walkers. To date, ASIMO is considered to be among the most versatile robots. However, it lacks speed, robustness and energy efficiency, for which its physical performance is not considered to be ranked high. Possible explanations for its lack of performance, despite years of research, are its completely stiff drive train, as explained previously, but also its inherent complexity, totalling 30 actuated DoFs.

RISK MANAGEMENT

Projects on robots that do not require risk management and have a reduced penalty of damage have generally shown faster progress. Two examples of extremities are the design of robots for space exploration (Bogue, 2012) and Lego Mindstorms (Radford et al., 2015). Space projects require much more resources to assure the functionality of technology, rather than to develop it, because the penalty of damage (i.e. a failed space mission) is in the millions or billions of Euro.

Salto is exemplary of a project that shows fast and cheap progress with little risk involved. The small robots are not prone to damaging themselves or the environment when falling, and it is built from cheap components, so that if it breaks, then the broken parts are easily replaced. This enables unconcerned and swift execution of experiments. The desire for risk-free robotic platforms for speeding up robotics research has gained increased popularity, and lead to the development of for instance the affordable and rugged quadruped Stoch (Dholakiya et al., 2019).

Projects on robots that require measures to reduce risk have shown slower progress, such as those on the quadrupedal HyQ (Semini, 2010; Semini et al., 2017) and humanoid Petman (Nelson et al., 2012). These robots are built from more expensive components, use more dangerous actuators and are larger. The robots are hydraulically actuated,



(a) Protection on Elios (cage) and Leo (sponge) (b) Exposed actuators on FROG and Wabian-2R

Figure 2.6: Physical protection on robots (Elios (Flyability, 2018) and Leo (Schuitema, 2012)) that have been designed to crash, and lack thereof on robots (FROG (Wang & Yang, 2011) and Wabian-2R (Ogura et al., 2006)) that are not designed to crash, consequently requiring more risk management.

which require safety measures. The robots are also larger, and therefore more powerful, which enable them to inflict damage to their surroundings (including humans) and themselves. Experiments are therefore to be executed in a more controlled environment, which requires additional resources. The experiments would also be executed more cautiously, such as by keeping a low locomotion speed and by avoiding risky manoeuvres, essentially preventing the robot from operating at its performance envelope, as discussed before.

The requirement for risk management can be reduced by making robots more physically robust, and by keeping them small, which are the next two discussed key points.

PHYSICAL ROBUSTNESS

By physical robustness is meant the ability to withstand impacts, whether desired (e.g. by making a step) or undesired (e.g. from a fall), as opposed to behavioural robustness, which refers to the ability to respond to undesired disturbance (as described in section 2.1). From a certain perspective, physical and behavioural robustness are similar, as they both have the goal of letting a robot perform its task after being disturbed: a robot that falls over but is able to get back up unaided and undamaged could continue performing its original task. As such, a robot that is physically robust can be considered to have a higher physical performance than a robot that is not.

Physical protection is mainly popular on robots that are designed to crash or fall frequently, such as protective cages on flying robots (Briod, Klapotcz, Zufferey, & Floreano, 2012; Briod, Kornatowski, Zufferey, & Floreano, 2014; Flyability, 2018; Klapotcz, Briod, Daler, Zufferey, & Floreano, 2013) or the sponge bumpers on the humanoid Leo that learns by falling (Schuitema, 2012), as shown in fig. 2.6a. However, physical robustness does not seem to be commonly addressed as a design requirement on robots that are not *intended* to crash or fall, whereas they are still able to. In addition to lacking physical protection, it is frequently observed that vulnerable components are placed on the exterior of a robot's skeleton (exposed to the environment). For example, the autonomous quadrupedal FROG (Shao, Huang, Wang, Cai, & Wang, 2014; Wang & Yang, 2011) and Wabian-2R (Ogura et al., 2006) have their motors and transmissions, respectively, protrude from their hulls, as shown in fig. 2.6b.

Making a robot physically robust avoids the necessity of a harness, and contributes

more progressively to gaining experience in robot design (by paying attention to how it could break) and experimentation (by having reduced risk management). The next key points elaborates more on the use of a harness. Robot projects that allow their robots to fall have demonstrated fast and steady progress in both robot design for performance and physical robustness. Examples are the Salto project and Boston Dynamics' projects.

DESIGN NOT BIGGER THAN NECESSARY

Most of the successful robots are smaller than average and smaller than their precursors. Examples of robots that are smaller than their precursors and yet exhibit higher or similar performance include SpotMini (precursor: Spot), ASIMO (precursors: P2, P3, et al.) and Cheetah 3 (precursors: Cheetah, Cheetah 2). Benefits of small design are higher physical robustness and better risk management, which subsequently lead to faster progress. Other high performance lightweight quadrupeds that have been developed in a short amount of time include mini Cheetah (Bosworth, Kim, & Hogan, 2015), Minitaur (Ghost Robotics, 2016; Kenneally, 2016), Laikago (Unitree) (Unitree, 2017) and ANYmal (Hutter et al., 2016). Salto is another example of a robot that exhibits high performance and that is lightweight. Small robots furthermore seem to be constructable from cheaper materials, of which Salto and Stoch (Dholakiya et al., 2019) are examples, and subsequently facilitate design by trial-and-error.

An argument that is used for making robots bigger is that they have to be operational in the same urban environment as humans. Most of the robots from the DARPA challenge, which are designed for a human environment, are heavy too. They are typically designed to have a size similar to that of a human, but end up being heavier because metal has a higher density than biological tissue. For example, DARPA contestants JAXON (Kojima et al., 2015) and WALK-MAN (Negrello et al., 2016) weigh 127 kg and 135 kg respectively. The search and rescue robot WAREC-1 weighs 155 kg and is 169 cm tall when standing upright (Hashimoto et al., 2017). However, Boston Dynamic's SpotMini has proven that it too can navigate in human environments (Boston Dynamics, 2018), despite being much smaller. Minitaur, weighing 5 kg has shown that it can climb stairs and open doors.

Another argument that is used for making robots bigger than necessary is the greater availability of commercial off-the-shelf (COTS) industrial components, which have the typical disadvantage that they must abide various types of (safety) regulations, for which they are typically over-designed and heavier than necessary. However, limited availability of smaller components should not necessarily be an issue, as in most cases the design leaves many parameters that can be chosen so as to tailor the design to the available COTS components. In addition, the increased popularity of robotics and manufacturing technologies has opened new doors for robot design. For instance, some companies have specialized in motor units for robots, and the rise of commercial 3D printing and milling services and machines allows for the creation of custom yet affordable shapes.

FULL AUTONOMY

Autonomy here refers to the lack of external physical support, such as an umbilical to power the robot or a harness to prevent it from damaging itself or the environment. The harness is introduced for risk management and damage control, but its required use is an indication that the robot is inadequately designed by lacking physical robustness

or forming a danger to the environment in the first place, for which it is seen as a monument to a wrong type of technical solution.

A distinction can furthermore be made between robots designed to be fully autonomous, but use a harness for (some) experiments, and those that were not designed to be fully autonomous in the first place. The latter is an example of acquiring oneself a luxury that is unavailable in the final application. Various projects have shown to struggle to upgrade their subsequent designs to being fully autonomous, once having grown attached and accustomed to having the support of an umbilical or harness: making the step to their removal has shown to be bigger of a challenge than anticipated. A few examples of robots and their successors that have not demonstrated operability without their harness or umbilical are (Kaminaga et al., 2016), Athlete Robot (Niiyama et al., 2010), TaeMu (Hyon, Suewaka, Torii, & Oku, 2017), Kojiro (Mizuuchi et al., 2008) and HyQ (successor: HyQ2Max) (Semini, 2010; Semini et al., 2017).

Additional reasons to avoid an umbilical or harness are described below.

- An umbilical or harness hinders the robot from performing its motions, and thus physically prevents the robot from being used to its fullest potential. The resulting behaviour is thus also not going to be representative for a fully autonomous successor.
- A robot that lacks protective parts (to assure physical robustness) or an on-board power supply is governed by a dynamic model that is not representative for a fully autonomous successor.
- If the robot is not permitted to fall, and thus also not to break, then no experience will be gained in finding its weak spots, to be taken over to successive designs.

GOOD ENOUGH PRIORITISES OPTIMALITY

Optimization is popular in most industrial fields, such as in car industry, in which for instance crashworthiness research focuses on weight-optimising frames to dissipate as much as possible energy for a maximum peak deceleration during crashes (Patel, Kang, Renaud, & Tovar, 2009; Pedersen, 2003). This type of research is justified by the fact that millions of cars are produced every month, and that it might save human lives. Instead, most of the legged robots are experimental and only produced once or a few times, especially if they are part of a design by trial-and-error program. To improve the effectiveness of trial-and-error research, and to reduce the necessity for risk management by reducing cost, it is advised to strive for a design that is good enough rather than for one that is optimal, especially in early design cycles. High performance robots such as Salto, ANYmal and Cheetah are not designed to be optimal, and the designers of Cheetah have openly claimed this approach to be part of their quick success. Also Boston Dynamics only recently started to experiment with complex optimised parts on newer versions of Atlas by integrating multiple functions and by performing automated topology optimisations.

The purpose of optimisation studies in an early stage of the design, such as those described under the key point of machine and behaviour co-design, is to expand the horizon of feasible solutions and to obtain an understanding of desired properties that fit the behaviour, rather than to find an optimal design. Another pitfall of optimisation,

especially at an early stage and for complex machines, is that it is prone to overlooking scenarios and introducing ill-defined constraints, subsequently leading to weak or infeasible solutions.

3

SKIPPY: STATE-OF-THE-ART

This chapter describes activities and decisions that are relevant for the development of Skippy but that are not the author's own work or have been executed prior to initiation of the thesis project.

OUTLINE

This section summarizes a recently developed balance control theory that is described in more detail in Featherstone (2015a, 2017), and that has its roots in (Azad, 2014; Azad & Featherstone, 2016). Preliminary design decisions and principles of Skippy are summarized in section 3.1. The proposed design is used as a basis to perform accurate simulation studies towards hopping 4 m high. The study leads to a control strategy and more detailed design choices of the robot, and is summarized in section 3.2. Other work focuses on balancing instead of hopping. The theory behind a newly proposed balance controller is summarized in section 3.3.

3.1. PRELIMINARY DESIGN

This section presents relevant design decisions and related design principles that were made prior to the study presented in this thesis.

3.1.1. DESIGN DECISIONS

Design decisions can roughly be enveloped by three core subjects: the mechanism, mass distribution and the use of series springs.

SKIPPY'S MECHANISM

Figure 3.1a shows the tentative design of Skippy's mechanism, as described in Driessen (2015). It is, essentially, a planar double pendulum with a crossbar at its head, with the

Acknowledgements — Section 3.3 includes excerpts from Driessen, Gkikakis, Featherstone, and Singh (2019), some of which are from the writings of Dr. Featherstone.




Figure 3.1: Skippy's tentative design in 2015 and the simplified representation of the hip's 4-bar linkage used for simulation.

second (hip) joint replaced for a 4-bar linkage mechanism. The torso and leg together make a planar hopping and balancing machine, and the crossbar stabilizes and steers it in 3D. The feasibility of hopping and full 3D balancing was established in (Azad, 2014; Azad & Featherstone, 2013, 2014a).

The 4-bar linkage enables the hip joint to have approximately double the angular range of the input lever, e.g. a range of ca. 180° for a lever range of 90° . The linkage mechanism furthermore allows for a beneficial non-constant transmission ratio.

MASS DISTRIBUTION

The robot's mass distribution is a compromise between the conflicting needs of hopping and balancing. For good hopping performance, the CoM must be as close as possible to the head; but for good balancing performance (as measured by velocity gain (Featherstone, 2016b)) the torso needs a large radius of gyration. The compromise is to place most of the mass near the head, but to have a small concentration of mass at the hip (i.e., the 4-bar linkage), and to make the middle portion of the torso as light as possible. This is achieved by placing both the main motor and crossbar actuator close to the head. The main motor then drives the hip lever through a ball screw and the main spring.

SERIES SPRINGS

The presence of the main spring makes the actuator an SEA. The main spring is necessary for three reasons:

1. to recycle energy from one hop to the next,
2. to allow the main motor to do positive work during the entire stance phase, and
3. to allow the hip joint to reach the necessary speed for lift-off.

The robot has a secondary series elastic element: the ankle spring, which is driving the unactuated ankle joint. This spring also stores some energy, but the main purpose of the




Figure 3.2: A few frames of a simplified depiction of Skippy during its stance phase while preparing its highest hop. The function of the ankle joint during steering is also displayed by inclusion of approximate indications of the CoM's location and the GRF's direction and magnitude: build-up of positive angular momentum during the landing and thrust phase can be (partly) compensated for during the steering phase. The activity of the springs and nut are shown too, as their displacements are drawn approximately to scale.

spring-loaded ankle is to allow the robot to control its angular momentum at lift-off.

3.1.2. DESIGN PRINCIPLES

Skippy's design principles find their roots in an analysis of the required behaviour during the stance phase towards making the highest hop, based on which some ballpark figures are derived.

STANCE PHASE

The *stance phase* preceding Skippy's highest hop is depicted in fig. 3.2. This figure shows several frames of interest of Skippy's stance phase; the time step in between each of the frames is not constant. The stance phase is divided into a *landing phase* and *launch phase*. The beginning of the launch phase is defined by the nadir, at which Skippy's CoM is at its lowest point. The launch phase is divided into a *thrust phase*, during which the main motor is in saturation and the hip torque is high, followed by a *steering phase*, during which the main motor can be either in saturation or not, and the hip torque is lower.

The presence of the series elastic main spring allows the motor to do positive work during the landing phase and at the nadir, while the system is being slowed down. This implies that the spring reaches maximum elongation after the nadir, which is visually indicated in frame 3 of fig. 3.2.

The beginning of the thrust phase is depicted in frame 2, at which the ankle joint is (nearly) at its flexion end stop. It is held there by the high hip torque. This continues

until the end of the thrust phase, depicted in frame 4, at which point the reduced hip torque during the steering phase allows the ankle to extend. The timing and rate of extension is influenced by the activity of the motor. Throughout the thrust phase, the ground reaction force (GRF) passes slightly to the right of the CoM, implying a build-up of positive angular momentum. This build-up of positive momentum is (partly) compensated for by the change of motion as the ankle joint extends, as it causes the GRF to change direction so that it passes to the left of the CoM. This offers the control system the possibility to bring the robot's net angular momentum to zero, or to a desired non-zero value, before take-off, depicted in frame 5.

BALLPARK FIGURES

Skippy's target mass is 2 kg and it is approximately 60 cm in length when flexed (as in fig. 3.1 (Driessen, 2015)). At this mass, it requires 80 J of kinetic energy at lift-off to make a 4 m hop. As mentioned in chapter 1, it is a design requirement that Skippy be able to reach this height without relying on more than half of the energy being recycled from the previous hop. This implies that the hip actuator must be able to supply at least 40 J of new kinetic energy during the stance phase. Assuming a stroke length of 0.45 m (the vertical component of CoM travel from its lowest point to lift-off), the stance phase will last approximately 0.2 s. This implies 9g average vertical acceleration, 200 N average vertical GRF, and 200 W average power output from the hip actuator.

The limiting factor on the performance of this design is the maximum speed of the ball screw. The study in Driessen (2015) uses a 8 mm diameter, 2 mm pitch ball screw with a thrust rating of 3 kN from the Steinmeyer catalog (Steinmeyer, 2015), paired with a Maxon's direct current (DC) brushed RE35 90 W 24 V motor. The Steinmeyer ball screws can be operated at 6000 rpm,¹ which equates to a nut speed of 0.2 m s^{-1} . At this speed, the nut must pull at 1000 N in order to deliver 200 W, and – assuming 90% efficiency – the necessary driving torque would be 0.35 N m^{-1} . Two more limits are the 30 V available from the battery pack and the 50 A current limit of the motor drivers. Within these limits, the motor can deliver 0.58 Nm at 6000 rpm, and the maximum available output power is 383 W at 4889 rpm. The motor will overheat if Skippy makes a lengthy sequence of 4 m hops.

3.2. PRELIMINARY HOPPING STUDIES

Planar simulation studies have been done on the launch phase leading to a controllable 4 m high hop (Driessen, 2015). The study results in a proposed control strategy and tentative values for critical system parameters (e.g. those defining the spring profiles) that enable Skippy to reach a 4 m high hop controllably.

The essence of the design process in the study is as follows: the robot is made to operate at its performance limit by saturating the motor, the resulting motion is observed and compared with the desired motion, and the robot is redesigned so as to reduce the difference. The end result is a robot whose natural dynamics at its performance envelope coincides with the most demanding behaviour it is designed to deliver; and the role of feedback control is merely to fine-tune the behaviour rather than to try and impose an

¹The operable speed of 6000 rpm was recommended by colleagues in the robotics community. The catalog recommends a top speed of only 4500 rpm.

artificial dynamics on the robot. The design process is therefore holistic, as it considers the behaviour of the whole robot. It is also an example of machine and behaviour co-design, because only the final outcome (a 4 m hop) is specified in advance, not the action sequence to achieve it, and the objective is to discover both the robot and its behaviour.

MODEL AND OBJECTIVE

The simulations assume realistic limitations of the system, including limited nut velocity of the ball screw, motor saturation (where both voltage and current are limited) and mechanical friction. To facilitate the study, two simplifications have been made to Skippy's mechanism:

1. The crossbar is removed and its inertia added to the torso; and
2. the 4-bar linkage at the hip is replaced by a pair of revolute joints geared 1 : 1, as shown in fig. 3.1b, which approximates the kinematics of the 4-bar linkage.

Using this simplified mechanism, the chosen motor and ball screw, the proposed robot size and an estimate of the mass distribution, only a few model parameters are left to be determined. These model parameters include the lever length l_1 and angle ϕ_1 , as shown in fig. 3.1b. The objective is then to find suitable values for these model parameters, and a suitable control strategy, such that an accurate simulation shows that the robot has reached a sufficient velocity at lift-off to make a 4 m hop.

3.2.1. STRATEGY

To achieve a 4 m hop, the main challenge is to move energy very quickly from the battery to the mechanism. The strategy is therefore to focus on maximizing energy flow. This is accomplished as follows.

1. Maximum energy flows from the battery to the main motor when the full battery voltage is applied to the motor. The motor is then said to be operating *in saturation*.
2. The motor's mechanical load must be designed such that when the motor is in saturation it is operating at a speed close to its maximum-power speed.
3. The mechanism, the springs and their initial conditions must be designed so as to convert as much as possible of the incoming mechanical power and stored elastic energy into vertical motion of the CoM.

In addition, since the study also has an interest in controllability, there is one more requirement:

4. The steering phase must be able to reach a range of possible angular momenta at lift-off, which must include zero.

The control strategy for the hop is then as follows. During the thrust phase the motor is kept in saturation. During the steering phase, bang-bang control is applied on the input voltage, on the grounds that it is simple and it works. However, it is not claimed that this is optimal. Simulation studies confirmed that it is possible to use only one motor for generating vertical momentum (thrusting) and controlling rotational momentum

(steering). Steering for maximum positive angular momentum required the motor to be kept in saturation also during the steering phase, for which the robot consequently jumps higher than for jumping with zero angular momentum at lift-off.

3.3. EARLY BALANCING STUDIES

The theory concerns balancing of a machine on one or multiple unactuated revolute base joints (implying that the system is under-actuated and intrinsically unstable), while also tracking an arbitrary position command. The base joints represent a point contact in 2D or 3D, a line contact in 3D, or—in practice—a small or narrow area. In practice, the theory is applicable to humanoids with point feet, but also to those with flat feet standing on their feet's edges or on a convex surface such as a hump. The theory is also applicable to quadrupeds balancing on two of their four feet, which could be of practical use for quasi-static creeping gaits on—for example—the edge of a wall or the ridge of a steep mountain.

The theory has furthermore been amended by the implementations of the method of *leaning in anticipation* and a *balance offset observer*, the former of which is already introduced in Featherstone (2017). The method of leaning in anticipation is largely responsible for the high balancing performance. It is based on a simple low-pass filter running backwards in time, and offers a solution to the inevitable non-minimum-phase behaviour that is intrinsic to the physics of balancing: an unanticipated disturbance or desired motion in one direction will always require an initial response in the opposite direction to recover balance.

The balance offset observer is a new theoretical contribution. It implements a method to observe and compensate for the discrepancy between the estimated and true balanced configurations. This extension improves the accuracy of the control system and reduces its reliance on having an accurate model of the robot and accurate sensor signals.

In this section the controller is explained for the special case of a single-actuated 2R (2 revolute joints) balancing mechanism, as this is a mechanism similar to Skippy in the sagittal plane. Skippy in the frontal plane – also referred to as swivel plane – is a sub-case of this already special case, which is further explained in chapter 4. An extension to more general planar mechanisms and 3D can be found in Featherstone (2017).

3.3.1. COMPARISON TO OTHER WORK

Robotic balancing is not a new topic, and has already been solved before. The claimed contribution of the recently developed control theory is that it is simpler, more generalizable and that it leads to a higher level of balance performance, enabling the robot to make faster and larger movements than earlier balance controllers. The theory is claimed to be simpler because it is based on controlling a plant that (1) considers only the exact physical problem of balancing, and as such extracts from the robot's dynamics only the non-linear characteristics of the robot that are important for balancing, and (2) embeds those characteristics in only two gains. It is more generalizable because it only requires *numeric*—rather than (manipulated) *symbolic*—values of the coefficients of the equation of motion (EoM) (e.g. those from the joint-space inertia matrix and joint-space

vector of force terms), which implies that the theory is scalable to more complex mechanisms, i.e. those that have more degrees of freedom and that have closed-chain linkages (kinematic loops). In simulation, the theory has shown to outperform previously published balance control solutions, including those described in Azad (2014); Azad and Featherstone (2016); Berkemeier and Fearing (1999); Grizzle, Moog, and Chevallereau (2005); Spong (1995), which can be seen in figure 12 of Featherstone (2017). The higher performance here is also due to the introduction of the aforementioned method of leaning in anticipation, which makes the robot lean in anticipation of motion commands.

Few other works have demonstrated high-performance balancing on *real* robots. However, in many cases these works have simpler control objectives and systems. For example, in Cubli (Muehlebach & D'Andrea, 2017) and Murata Boy (muRata, 2015), the objective of maintaining balancing is combined with controlling a velocity variable, which is less demanding than making the robot track a generalized position variable command signal. The difference is that between controlling a third- and fourth-order plant. Note that these systems furthermore describe the special case of a reaction wheel balancer, where the CoM of the second body coincides with the second joint. This is also the system treated in chapter 4 and section 5.6.

3.3.2. MODEL

The controller requires the numeric values of the coefficients of the equation of motion (EoM). These are the numeric values of the elements in the joint space inertia matrix \mathbf{H} and the joint-space vector of force terms \mathbf{C} , if the EoM are written as

$$\boldsymbol{\tau} = \mathbf{H}\ddot{\mathbf{q}} + \mathbf{C} \quad (3.1)$$

where $\ddot{\mathbf{q}}$ and $\boldsymbol{\tau}$ are the generalized joint acceleration and force vector respectively. For the special case of a planar 2R mechanism where only the second joint is actuated, the EoM can be written as

$$\begin{bmatrix} 0 \\ \tau_2 \end{bmatrix} = \begin{bmatrix} H_{11} & H_{12} \\ H_{21} & H_{22} \end{bmatrix} \begin{bmatrix} \ddot{q}_1 \\ \ddot{q}_2 \end{bmatrix} + \begin{bmatrix} C_1 \\ C_2 \end{bmatrix}. \quad (3.2)$$

The planar 2R mechanism is depicted in fig. 3.3. It consists of two rigid bodies and two revolute joints, of which only the second joint is actuated. Body 1 is connected to the ground via (the unactuated) joint 1, which is equivalent to assuming the system is a robot with point feet that neither slip nor loose contact with the ground. Body 2 is connected to body 1 via (the actuated) joint 2. In the field of control theory the system is known as the Acrobot (Hauser & Murray, 1990).

$[c_x \ c_y]$ are the coordinates of the robot's centre of mass with respect to the ground. The robot is in a balanced position if $c_x = 0$ and if it is stationary, i.e. if $\dot{q}_1 = 0$ and $\dot{q}_2 = 0$. The robot is also subject to the position command signal $q_2 = q_c(t)$. This implies that the robot cannot remain in a balanced position while tracking the command signal, and that the command signal cannot be tracked to perfection if the robot is not permitted to lose its balance. The new balance method revolves around tracking $q_c(t)$ while keeping the robot about its balanced position as a side effect.




Figure 3.3: The mechanism of a 2R planar balancing machine, where only the second joint is actuated, like the Acrobot.




Figure 3.4: Control plant

3.3.3. CONTROL LAW

The plant shown in fig. 3.4 forms the basis of the new balance theory. By controlling the actuated joint position q_2 via this plant, then as a side effect the robot will also maintain its balance. L , \dot{L} and \ddot{L} are state variables and \ddot{L} is the input, and are the robot's angular momentum about the support point and its first three derivatives. It can be derived that

$$L = H_{11}\dot{q}_1 + H_{12}\dot{q}_2 \quad (3.3)$$

$$\dot{L} = -mgc_x \quad (3.4)$$

$$\ddot{L} = -mg\dot{c}_x \quad (3.5)$$

$$\ddot{\ddot{L}} = -mg\ddot{c}_x, \quad (3.6)$$

with m and g the robot mass and gravitational acceleration respectively. The two gains are

$$Y_1 = \frac{1}{mgT_c^2 G_v} \quad \text{and} \quad Y_2 = \frac{-1}{mgG_v}. \quad (3.7)$$

T_c is the robot's natural time constant of toppling (Featherstone, 2015a, 2017), treating the whole robot as a rigid body, defined by

$$T_c^2 = \frac{H_{11}}{mc_y g}. \quad (3.8)$$

G_v is the robot's linear velocity gain as defined in Featherstone (2015b, 2016b). Velocity gains are ratios that measure the degree to which an actuated joint's motion influences the motion of the robot's CoM. For our 2R robot shown in fig. 3.3, we have

$$G_v = \frac{\Delta \dot{c}_x}{\Delta \dot{q}_2}. \quad (3.9)$$

where the change in horizontal velocity of the center of mass $\Delta \dot{c}_x$ and change in velocity of the actuated joint $\Delta \dot{q}_2$ are caused by an impulse at the actuated joint. Note that also G_v can be written as function of numeric values of the coefficients of the EoM, as described in Featherstone (2017) by solving the impulsive equations of motion.

Note that both T_c and G_v are not constant, but they can be easily calculated from the robot's (impulsive) equations of motion, and they are easily measured properties of a robot mechanism because of their physical meanings. They are nearly constant for Skippy in the swivel plane, as further discussed in section 5.5.1.

Following the steps in Featherstone (2017), the simplified plant can be controlled using the control law

$$\ddot{L} = k_{dd}\dot{L} + k_d\dot{L} + k_M L + k_q(q_2 - u), \quad (3.10)$$

in which the feedback gains are given by

$$\begin{aligned} k_{dd} &= -a_3 & k_d &= -a_2 + a_0 Y_2 / Y_1 \\ k_M &= -a_1 & k_q &= -a_0 / Y_1 \end{aligned} \quad (3.11)$$

and

$$\begin{aligned} a_0 &= \lambda_1 \lambda_2 \lambda_3 \lambda_4 \\ a_1 &= -\lambda_1 \lambda_2 \lambda_3 - \lambda_1 \lambda_2 \lambda_4 - \lambda_1 \lambda_3 \lambda_4 - \lambda_2 \lambda_3 \lambda_4 \\ a_2 &= \lambda_1 \lambda_2 + \lambda_1 \lambda_3 + \lambda_1 \lambda_4 + \lambda_2 \lambda_3 + \lambda_2 \lambda_4 + \lambda_3 \lambda_4 \\ a_3 &= -\lambda_1 - \lambda_2 - \lambda_3 - \lambda_4, \end{aligned} \quad (3.12)$$

where $\lambda_1, \lambda_2, \lambda_3$ and λ_4 are the closed-loop poles of the system. Suitable choices of poles and zeros or their relations are given below.

The input signal u , is given by

$$u = q_c + \alpha_1 \dot{q}_c + \alpha_2 \ddot{q}_c, \quad (3.13)$$

where q_c, \dot{q}_c and \ddot{q}_c are the commanded values of q_2, \dot{q}_2 and \ddot{q}_2 , respectively, and α_1 and α_2 are feed-forward gains. These allow the inclusion of up to two additional zeros in the closed-loop response of the system. For the case of two zeros, these gains are set as follows:

$$\begin{aligned} \text{no zeros:} & & \alpha_1 &= \alpha_2 = 0 \\ \text{zero at } \mu_1 \text{ only:} & & \alpha_1 &= -1/\mu_1, \alpha_2 = 0 \\ \text{zeros at } \mu_1 \text{ and } \mu_2: & & \alpha_1 &= -(\mu_1 + \mu_2)/(\mu_1 \mu_2) \\ & & \alpha_2 &= 1/(\mu_1 \mu_2). \end{aligned}$$

in which the zeros are located at μ_1 and μ_2 .

The output of the balance controller has to be either a torque τ_a or an acceleration command \ddot{q}_a for joint 2. Either can be computed from \ddot{L} using eq. (3.6), the EoM in

eq. (3.2), and an expression describing \ddot{c}_x as function of the generalized joint position variables and their respective derivatives or forces. A method to obtain the latter numerically for a general system is explained in Featherstone (2017), which involves expanding the EoM to include a fictitious prismatic joint between joint 1 and the base in the x direction, so that its joint force equals the horizontal ground reaction force (GRF) $m\ddot{c}_x$.

3.3.4. TRANSFER FUNCTION

It can be shown that the closed-loop transfer function of the plant under the action of the control law in eq. (3.10) is

$$q_2(s) = \frac{a_0(1 - T_c^2 s^2)(1 + \alpha_1 s + \alpha_2 s^2)}{s^4 + a_3 s^3 + a_2 s^2 + a_1 s + a_0} q_c(s). \quad (3.14)$$

As such, there are four poles in total, all of which determined by the feedback gains, two zeros determined by the mechanism, and up to two more zeros determined by the feed-forward gains. The simulation studies in Featherstone (2017) show that it is a good idea to keep the feedback gains low; so a good choice of poles and zeros is as follows: set $\lambda_1 = -1/T_c$ so as to cancel the mechanism-dependent zero at $-1/T_c$; choose slow frequencies for λ_2 and λ_3 so as to keep the feedback gains low; set $\mu_1 = \lambda_2$ and $\mu_2 = \lambda_3$, so that the two optional zeros cancel these two poles. response of the closed-loop system. One possibility is choosing $\lambda_2 = \lambda_3 = \mu_1 = \mu_2 = -1/T_c$, so that the transfer function simplifies to

$$q_2(s) = \frac{1 - T_c s}{1 - s/\lambda_4} q_c(s). \quad (3.15)$$

This is also the choice of poles and zeros for the experiments reported in chapter 5. λ_4 determines the response speed of the closed-loop system, which is typically desirably maximized, for which it is chosen as high as possible in magnitude and limited mainly by physical capabilities of the robot.

3.3.5. LEANING IN ANTICIPATION

The one remaining zero in eq. (3.15) is problematic: it is the zero that is responsible for the non-minimum-phase behaviour of the system, and it cannot be cancelled by a pole because the corresponding pole is unstable. However, there is a technique that will solve this problem: pass the command signal through a first-order low-pass filter with time constant T_c , but which is running backwards in time from a point sufficiently far in the future. This works because the transfer function $1/(1 + T_c s)$ in reverse time behaves like $1/(1 - T_c s)$ in forward time. This scheme requires knowledge of what the command signal is expected to be at a time approximately one or two seconds ahead of the current instant, but such information is often available. A filter whose output depends on future values of its input is called an acausal filter.

Let q_f be the result of passing q_c through the acausal filter. If we use q_f in place of q_c

as the input to the balance controller then

$$\begin{aligned} q_2(s) &= \frac{1 - T_c s}{1 - s/\lambda_4} q_f(s) \\ &= \frac{1 - T_c s}{1 - s/\lambda_4} \frac{1}{1 - T_c s} q_c(s) \\ &= \frac{1}{1 - s/\lambda_4} q_c(s). \end{aligned} \quad (3.16)$$

The observed effect of this filter is that the non-minimum-phase behaviour disappears, and the robot leans in anticipation of future commanded motions.

3

3.3.6. BALANCE OFFSET OBSERVER

The theory described above assumes that the robot is balanced when $q_1 = 0$. However, simulation studies have shown that balance controllers in general are quite sensitive to errors in their estimate of a balanced configuration (Azad, 2014; Azad & Featherstone, 2016). In practice these errors are introduced by IMU drift or an offset bias in the base encoder (measuring q_1). Therefore, we introduce here a simple method for observing and compensating for this error.

First, we retain the idea that the robot is balanced when $q_1 = 0$, but introduce two new quantities, \hat{q}_1 and q_o , which satisfy

$$q_1 = \hat{q}_1 - q_o. \quad (3.17)$$

\hat{q}_1 is the measured value of q_1 , and q_o is the unknown balance offset error, which is assumed to be a slowly varying quantity.

In particular, it is assumed that q_o varies sufficiently slowly that \dot{q}_o is too small to significantly affect the measured (or calculated) value of \hat{q}_1 that is used in eq. (3.3) to calculate L . We now have two ways to calculate \dot{L} , one using eq. (3.4) and one by differentiating eq. (3.3). From eq. (3.3) we have

$$\dot{L} = \frac{dL}{dt} = \frac{d}{dt} (H_{11} \dot{q}_1 + H_{12} \dot{q}_2). \quad (3.18)$$

where d/dt denotes numerical differentiation.

Equation (3.4) and c_x can generally be expressed as a function of both q_1 and q_2 (all position state variables). If we assume that these expressions can be manipulated to a form so that q_1 can be written as function of \dot{L} and q_2 , so that

$$q_1 = f(\dot{L}, q_2), \quad (3.19)$$

then eqs. (3.17–3.19) can be combined to provide the following formula for estimating q_o :

$$q_o = \text{LPF} \left(\hat{q}_1 - f \left(\frac{dL}{dt}, q_2 \right) \right), \quad (3.20)$$

where ‘LPF’ stands for a low-pass filter. Having obtained an estimate of q_o , it can be incorporated in the balance controller by replacing q_1 with $\hat{q}_1 - q_o$ in eqs. (3.4–3.6).

4

MAXIMIZING BALANCING PERFORMANCE

To analyse a machine's physical ability to balance, one could consider its behaviour at the performance envelope, as determined by saturating its actuators, which does not require the involvement of closed-loop controllers that are responsible for actually balancing the machine. This chapter performs such an analysis on a simplified model of a balancing robot, corresponding also to that of Tippy and Skippy in their frontal plane. The results are a formula that predicts the largest balance disturbance from which the robot can recover and an actuator design criterion that can be used to maximize this largest disturbance. The study assumes that the robot does not make a step, and that the robot can be approximated by the model of a reaction wheel pendulum, which has a single major joint that it uses to recover from the balance disturbance. The presented theories are used to design for Tippy and Skippy to be good physical balancers.

4.1. INTRODUCTION

The importance for Skippy to balance has been addressed in chapter 1, which has been a motivation to develop the new balancing control theory that is summarized in section 3.3 and subsequently to initiate the project of developing the balancing-only precursor Tippy, as further treated in chapter 5. However, in order to properly experimentally validate the claimed high performance of a balancing control algorithm, it is a prerequisite that also the robot can physically exhibit high levels of balancing performance. After all, the observed performance is a product of the balancing capabilities of both the control algorithm and the machine. However, despite the popularity of balancing in legged robot research, most work focuses on the control aspect, and little on the machine's aspects. The machine itself contains information about its ability to balance, without hav-

Acknowledgements — Parts of this chapter have been published in *IEEE/RSJ International Conference on Intelligent Robots and Systems 2018* (Driessen, Featherstone, & Gkikakis, 2018), and includes excerpts from that paper.

ing to be tied to a specific balance control algorithm. If this information is quantifiable in one or multiple performance measures, then it can be useful for the design and comparison of high performance balancing machines.

In particular, the physical performance of the system can be analysed by considering the robot's behaviour at the performance envelope. This would require one or more of its actuators to operate in saturation, and furthermore depends on the action taken to recover from a balance disturbance. A legged robot can either make one or multiple steps, or it can use its body's inertias to recover without making a step, providing the balance disturbance is not too big. This work focuses on the latter, which is also the appropriate scenario for Tippy, as it is hinged to the ground.

An algebraic formula is found that predicts the largest balance disturbance from which the robot can still recover, as function of the robot's and its actuator's physical parameters, by assuming a simplified robot model that consists of one actuator acting in saturation. In addition, an actuator design criterion that maximizes this largest permissible balance disturbance is derived from this formula. The criterion essentially says that the robot's time constant of toppling should be equal to the actuator's time constant under load. Balance disturbance here is initially defined as a maximum offset angle from which the robot can still recover to a balanced configuration, but it is also generalized to being a combination of a position and velocity offset.

To keep the problem tractable and allow for analytic derivations, a simplified robot model is used to represent the legged robot, but the actuator model is realistic. The model is governed by having a single major revolute actuated joint, such as a hip joint, which is the joint used for recovering from the balance disturbance, and the second body that is connected to this joint has its mass centralized at that joint. It is also assumed that most of the effective motion for recovering from the balance disturbance is done before hitting an end stop, as the actuated joint is not assumed to have joint rotation limits. In addition, it is assumed that the robot has point feet, like in section 3.3, and that the balance recovery movement is planar, so that a 2D analysis remains valid in 3D. The resulting dynamic model is known as a reaction wheel pendulum (RWP), which is also that of Tippy and Skippy in their frontal plane (using their crossbar to balance), and effectively that of several other robots, such as Cubli ([Gajamohan et al., 2012](#)).

The realistic actuator model is a DC motor with a non-perfect transmission mechanism, which is the most commonly used actuator in legged robots. Effects such as maximum voltage, back-EMF and friction are taken into account. Taking the realistic aspects of actuator dynamics and saturation into account to find a machine's physical capabilities of balance recovery, and to consider the full time trajectory in doing so, is a major novel aspect of this work.

OUTLINE

The rest of this chapter is organized as follows. First the background to this research is given in section 4.2, and then the model of the RWP including its actuator is defined in section 4.3. A formula for predicting the maximum recoverable balance error is derived in sections 4.4 and 4.5, as well as the conditions for this formula to have maximum value, which result in the design criterion. Section 4.6 then compares the formula's predictions with accurate dynamic simulations, as its derivations involve significant

assumptions. These comparison shows that the formula (and therefore also the actuator design criterion) is accurate under a broad range of conditions, but not universally so. Section 4.7 generalizes the maximum recoverable balance error from a position offset to a combined position and velocity offset.

4.2. BACKGROUND

Design criteria result from performance requirements, which can be quantified by performance measures. Various performance measures for balancing have been introduced previously. Some measures are designed specifically to be an input for control systems, such as the instantaneous capture point (Koolen et al., 2016), zero rate of change of angular momentum point (ZRAM)(Goswami & Kalleem, 2004) and the extrapolated center of mass (XCoM)(Hof, 2008). These are instantaneous measures that return a non-zero value if the robot is falling over; higher absolute values typically indicate the robot is falling over faster and that it is more difficult to retain balance without making a step. As such, controllers can be used to drive the measures to 0 in order to retain balance. The XCoM considers the robot as a composite rigid body, where bodies of the robot are not moving with respect to each other and where actuation does not come into play. As such, it does not provide information on the robot's capability to actively recover from an unbalanced configuration. The ZRAM uses the robot's Center of Mass (CoM) position and the ground reaction force (GRF) as input, which are both instantaneous variables, and the GRF is not straightforward to obtain. As such, the ZRAM has little use for describing overall balancing performance.

More recently the velocity gain has been introduced, as described in section 3.3.3. The velocity gain describes a machine's ability to alter its state of balance (Featherstone, 2015b, 2016b). In general, the velocity gain is defined as the ratio between an instantaneous change in a velocity coordinate of the CoM and an instantaneous change in the velocity of an actuated joint, both changes being caused by an impulse applied at the actuated joint. In particular, section 3.3.3 has defined a *linear* velocity gain G_v (eq. (3.9)), where the velocity coordinate of the CoM is the horizontal (translational) velocity of the CoM. The *angular* velocity gain G_ω is similar, which instead considers the rotational velocity of the CoM about the support joint. For the machine shown in fig. 4.1, the angular velocity gain is

$$G_\omega = \frac{\Delta \dot{q}_1}{\Delta \dot{q}_2}, \quad (4.1)$$

and for this machine it is related to the linear velocity gain as follows:

$$G_\omega = -\frac{G_v}{c \cos q_1} \quad (4.2)$$

In general, both G_v and G_ω are a function of the machine's position state, but G_ω is constant for the particular machine discussed in this work.

Unlike the XCoM or ZRAM, velocity gain does not measure a robot's state of imbalance, but rather its ability to correct that imbalance. If $G_v = 0$ in some particular configuration then it means that the robot is physically incapable of maintaining its balance. This makes velocity gain a better measure of a robot's *ability* to balance. Its chief disadvantage is that it is an instantaneous measure, and therefore does not take into account

the effect of gravity, or changes in configuration, over the finite time interval needed to make a balance recovery movement. Nor does it take into account actuator torque limits and their dependence on speed. These deficiencies are addressed in this chapter.

Apart from theories that focus on control or dynamics, but not on actuation, other work is based on actuators in isolation, which serve the purpose of improving mechanical design while decoupling the particular objective of (for instance) balancing. Classical approaches focus on choosing the gear ratio to match the effective actuator inertia to that of the load (Chen & Tsai, 1993; Pasch & Seering, 1984). For the family of direct drive robots, it is assumed that motor maximum continuous torque output is the main limiting factor (Kenneally, 2016), for which specific torque $K_{sp} = K_\tau \sqrt{R}^{-1} m_m^{-1} = K_m m_m^{-1}$ or thermal specific torque $K_{thsp} = K_{sp} \sqrt{R_{th}}^{-1}$ is used as measure. K_τ , m_m , R_{th} , R and K_m are the torque constant, motor mass, thermal resistance, winding resistance and motor constant respectively. Thermal resistance R_{th} measures the motor's capability to dissipate heat. It is often omitted in performance measures because it is difficult to determine and largely dependent on the motor environment (e.g. the heat sink, fan, or other cooling system). Naturally, these measures provide no use as overall performance measure as they only focus on actuation, but have nonetheless been proven useful for motor and transmission selection and design.

4.3. MODEL

A model has to be used with both dynamic and actuator properties in order to find a criterion or measure that does take into account both system dynamics and actuator properties towards the single goal of improving balancing performance. To do so, a simple system known as the reaction-wheel pendulum (RWP) with actuator limits is investigated.

An RWP is a planar multibody system with two revolute joints and two bodies, see fig. 4.1, where only the second joint is actuated; the unactuated first joint resembles a point foot. It is essentially a 2R planar double pendulum like the system defined in section 3.3.2 (fig. 3.3) with the special property that the CoM of the second body coincides with the second joint.

USE OF THE RWP

The RWP is a balancing system popular in control education (Aguilar-Avelar, Rodríguez-Calderón, Puga-Guzmán, & Moreno-Valenzuela, 2017; Andrievsky, 2011; Block et al., 2007; Meyer, Delson, & De Callafon, 2009). Most work focuses on control, although some previous work has partial focus on mechanics, where it is shown through numerical simulation that there is an optimal transmission ratio between the electric motor and reaction wheel (Meyer et al., 2009). The RWP has also been adopted as basis for centroidal momentum models, which are used to model simplified dynamics of more complicated legged robots (Lee & Goswami, 2007). A common model in robotics is the spring-loaded inverted pendulum (SLIP), of which some extended models include also upper body inertia (Cherouvim & Papadopoulos, 2005), effectively corresponding to a reaction wheel on a pendulum similar to the RWP. Other extended SLIP models, such as that of the virtual pivot point (Maus, Lipfert, Gross, Rummel, & Seyfarth, 2010), include also an offset of the upper body CoM with respect to the second joint.




Figure 4.1: The reaction wheel pendulum (RWP). Only the second joint is actuated. Note that the robot's CoM does not have to coincide with the line that passes through both joints, and that the second body does not have to take the shape of a wheel, as long as its CoM coincides with the second joint.

Reaction wheels, sometimes referred to as actuated flywheels, are moreover used in control moment gyroscopes, which provide – other than the reaction moment – also gyroscopic torques, and have found applications in human assisting exoskeletons (Chiu & Goswami, 2014; Matsuzaki & Fujimoto, 2013; Vallery et al., 2012). The RWP also has direct applications for robots that use reaction wheels like the Cubli (Gajamohan et al., 2012) and the balancing and hopping robot Skippy (Driessen, 2015) and its preceding balancing-only version named Tippy. Tippy is a spatial balancing machine, and is discussed in more detail in chapter 5. With the ability to lock some of its joints, the system resembles a planar RWP in one of its vertical planes. This robot is used to experimentally validate the results of this chapter in section 4.6.

MODEL DEFINITION

The two bodies of the RWP are named pendulum and reaction wheel, and are numbered 1 and 2 respectively. The ground is considered as body 0. Joint i connects body i to body $i - 1$. The two joint positions are described by joint variables of the vector $\mathbf{q} = [q_1 \ q_2]^T$, where q_1 is the angle between the vertical and pendulum and q_2 the angle between the pendulum and reaction wheel, as shown in fig. 4.1.

Since the CoM of the wheel coincides with the second joint, the position of joint 2 does not influence the robot's CoM coordinates and the location of joint 2 on the pendulum is not important. As such – compared to the robot in fig. 3.3 – q_1 is redefined to be measured from the line that passes through the robot's CoM rather than through the second joint. This is also illustrated in fig. 4.1. The robot is balanced when $q_1 = 0$.

The equation of motion (EoM) for this robot is

$$\begin{aligned} H_{11}\ddot{q}_1 + H_{12}\ddot{q}_2 &= mcg \sin(q_1) \\ H_{21}\ddot{q}_1 + H_{22}\ddot{q}_2 &= \tau, \end{aligned} \quad (4.3)$$

where \ddot{q}_i are the joint acceleration variables, H_{ij} are the elements of the joint-space inertia matrix, τ is the joint torque acting at joint 2, and g is the acceleration due to

gravity (a positive number). For this particular robot, it can be shown that

$$H_{11} = I_0 = mc^2 + I_1 + I_2 \quad (4.4)$$

and

$$H_{12} = H_{21} = H_{22} = I_2 = -G_\omega H_{11} \quad (4.5)$$

with I_0 the robot's composite rigid body inertia about the support joint, I_1 and I_2 the rotational inertia of bodies 1 and 2 about their CoM respectively, $m = m_1 + m_2$ the whole robot mass and c the distance between joint 1 and the robot's CoM. Note that this system has only three parameters that influence its behaviour (taking g as a constant): I_0 , mc and I_2 , where mc corresponds to the first moment of mass of the robot about joint 1, and I_0 to the robot's rotational inertia about joint 1 assuming that joint 2 is locked.

The motor that actuates the second joint is an electric motor. The motor rotor torque τ_r , motor rotor angular velocity ω_r , and applied voltage U can be related as follows, neglecting armature inductance:

$$\tau_r = U \frac{K_T}{R} - J_0 K_T \text{sign} \omega_r - \frac{K_T^2}{R} \omega_r$$

with K_T the torque constant, R the winding resistance and J_0 the no-load current. The second term describes the friction torque and is assumed to be constant, except for a sign dependence on rotor velocity¹.

By introducing a transmission with ratio $n = \omega_r / \dot{q}_2$ and constant forward efficiency η (corresponding to coulomb kinetic friction), where $0 < \eta < 1$, and by assuming that rotor and transmission inertia are negligible², then

$$\tau = \eta n \left(U \frac{K_T}{R} - J_0 K_T \text{sign} \dot{q}_2 - \frac{K_T^2}{R} n \dot{q}_2 \right).$$

4.4. STRATEGY

This section introduces a quantification for balancing performance and explains the required system behaviour to maximize balancing performance.

In this section balancing performance is measured as the maximum offset angle $q_1 = \phi_{\max}$ at rest ($\dot{q}(0) = \mathbf{0}$) from which the robot can return to a balanced configuration $q_1 = 0$. In section 4.7 the measure is generalized to allow for a non-zero initial velocity.

It is required to understand how the applied motor torque τ influences q_1 . A positive torque τ produces a positive acceleration at joint 2 and a negative acceleration at joint 1; so the balance recovery requires τ to be positive in order to bring q_1 down from ϕ_{\max} to zero. Torque is positively correlated with the applied voltage. Recovery from ϕ_{\max} requires the applied voltage to the motor to equal the maximum voltage that can be supplied by the batteries $U = U_{\max}$, at which point the motor is said to be operating in saturation.

The joint torque is linearly related to the joint velocity if the motor operates in saturation and if the joint velocity does not change sign. This linear torque profile can be

¹Details of the friction torque for $\omega_r = 0$ are not further treated.

²See section 4.6 for the numbers.

defined by two parameters (rather than using the six parameters K_τ , R , U_{\max} , J_0 , η , n); for instance using \bar{P}_m and \bar{K}_m , or using $\bar{\tau}_s$ and \bar{K}_m :

$$\tau = 2\bar{K}_m\sqrt{\bar{P}_m - \bar{K}_m^2\dot{q}_2} = \bar{\tau}_s - \bar{K}_m^2\dot{q}_2 \quad (4.6)$$

with

$$\bar{K}_m = n\sqrt{\eta}K_m = n\sqrt{\eta}\frac{K_\tau}{\sqrt{R}} \quad (4.7)$$

$$\bar{P}_m = \eta P_m = \eta\frac{(U_{\max} - J_0 R)^2}{4R} \quad (4.8)$$

$$\bar{\tau}_s = 2\bar{K}_m\sqrt{\bar{P}_m} = n\eta\tau_s = n\eta K_\tau\left(\frac{U_{\max}}{R} - J_0\right). \quad (4.9)$$

K_m is the motor constant, P_m the motor maximum power output and τ_s the stall torque minus the friction torque. A bar ($\bar{\quad}$) on top of symbols is used to denote ‘effective’ versions of those parameters. Effective here implies that the transmission ratio and friction effects are taken into account (i.e. the parameters apply to transmission output rather than to motor output).

Equation (4.6) can be incorporated in the EoM, which can then be analysed to find both experimental and algebraic relations between various of the physical parameters and the balancing performance. Incorporating eq. (4.6) in eq. (4.3) and linearizing at $q_1 = 0$ gives

$$\begin{bmatrix} I_0 & I_2 \\ I_2 & I_2 \end{bmatrix} \ddot{\mathbf{q}} = \begin{bmatrix} mcgq_1 \\ \bar{\tau}_s - \bar{K}_m^2\dot{q}_2 \end{bmatrix} \quad (4.10)$$

which describes the system with five parameters (I_0 , I_2 , mc , $\bar{\tau}_s$ and \bar{K}_m). An algebraic solution can be found. However, this solution occupies several pages and is not usable for further analyses, i.e. finding ϕ_{\max} .

Instead, a strategy is utilized that involves solving a simplified EoM in two steps with specific assumptions.

If $\dot{q}_1 \ll \dot{q}_2$, the possibility is introduced to first estimate reaction wheel dynamics, which can then be used to estimate pendulum dynamics. Two possible approaches are as follows. It can be assumed that either

- I pendulum dynamics do not influence reaction wheel dynamics, or
- II only gravity does not influence reaction wheel dynamics.

Derivations for both approaches are made and then compared with simulation studies.

4.5. DERIVATIONS

For the first approach (I), the equation of motion of the reaction wheel are formulated assuming that joint 1 does not move, resulting in

$$I_2\ddot{q}_2 + \bar{K}_m^2\dot{q}_2 = \bar{\tau}_s \quad (4.11)$$

which is a first-order dynamical system. For the second approach (II), the EoM are solved in the absence of gravity to obtain the reaction wheel dynamics, i.e.

$$\begin{bmatrix} I_0 & I_2 \\ I_2 & I_2 \end{bmatrix} \dot{\mathbf{q}} = \begin{bmatrix} 0 \\ \bar{\tau}_s - \bar{K}_m^2 \dot{q}_2 \end{bmatrix}$$

which can be rewritten to

$$I_a \dot{q}_2 + \bar{K}_m^2 \dot{q}_2 = \bar{\tau}_s \quad (4.12)$$

where I_a is to the apparent inertia of joint 2, and is given by $I_a = (I_0 - I_2)I_2/I_0$. Note that the only difference between eq. (4.12) compared to eq. (4.11) is that I_2 is replaced with I_a .

The remainder of this section shows only derivations for the first approach. For the second approach, eq. (4.13) and eqs. (4.15–4.24) require replacement of I_2 with I_a .

The solution of eq. (4.11) is

$$\dot{q}_2 = \frac{\bar{\tau}_s}{\bar{K}_m^2} \left(1 - e^{-t/\bar{T}_m}\right) \quad \text{with} \quad \bar{T}_m = \frac{I_2}{\bar{K}_m^2}. \quad (4.13)$$

\bar{T}_m is the actuator time constant under load, and therefore also the effective motor (mechanical) time constant. From simulation studies it was found that this result indeed predicts dynamic behaviour of the reaction wheel accurately.

Since eq. (4.13) turns out to be a reasonable approximation of reaction wheel dynamic behaviour, it can be used to approximate the behaviour of the pendulum by inserting the solution for \dot{q}_2 in the first equation of the EoM. However, the first equation in eq. (4.10) contains \ddot{q}_2 and no \dot{q}_2 . The approximation of \dot{q}_2 can be differentiated, but this possibly introduces additional error. To avoid this, the EoM can be written to contain \dot{q}_2 instead of \ddot{q}_2 in the first equation:

$$\begin{bmatrix} \check{I}_0 & 0 \\ I_2 & I_2 \end{bmatrix} \dot{\mathbf{q}} = \begin{bmatrix} mcg q_1 - \bar{\tau}_s + \bar{K}_m^2 \dot{q}_2 \\ \bar{\tau}_s - \bar{K}_m^2 \dot{q}_2 \end{bmatrix} \quad (4.14)$$

with $\check{I}_0 = mc^2 + I_1$. \check{I}_0 is the robot's articulated-body inertia about joint 1, which is the apparent inertia at joint 1 assuming that $\tau = 0$. By inserting the solution for \dot{q}_2 from eq. (4.13) in the first equation of the EoM in eq. (4.14) we obtain

$$\check{I}_0 \dot{q}_1 = \bar{K}_m^2 \left[\frac{\bar{\tau}_s}{\bar{K}_m^2} \left(1 - e^{-t/\bar{T}_m}\right) \right] + mcg q_1 - \bar{\tau}_s. \quad (4.15)$$

Reordering gives

$$\check{I}_0 \dot{q}_1 - mcg q_1 = -\bar{\tau}_s e^{-t/\bar{T}_m}. \quad (4.16)$$

The solution for this differential equation with initial conditions $q_1(0) = \phi$ and $\dot{q}_1(0) = 0$ is:

$$q_1(t) = A \left(e^{-t/\bar{T}_m} - \cosh(t/\check{T}_c) + \frac{\check{T}_c}{\bar{T}_m} \sinh(t/\check{T}_c) \right) + \underbrace{\phi \cosh(t/\check{T}_c)}_{\text{particular solution}} \quad (4.17)$$

with the time constant of toppling

$$\check{T}_c = \sqrt{\frac{\check{I}_0}{mcg}} \quad (4.18)$$

and

$$A = \frac{\bar{\tau}_s \bar{T}_m^2}{mcg \bar{T}_m^2 - \check{I}_0} = \frac{\bar{\tau}_s \check{T}_c^2 \bar{T}_m^2}{\check{I}_0 (\bar{T}_m^2 - \check{T}_c^2)}.$$

Note that the above definition of the time constant of toppling \check{T}_c differs from that of $\sqrt{I_0/(mcg)}$ or $T_c = \sqrt{I_0/(mc_y g)}$ in eq. (3.8) and Featherstone (2015a, 2017). \check{T}_c represents that of an articulated-body inertia, whereas T_c represents that of a rigid-body inertia (assuming the whole robot to be a single body). In addition, eq. (3.8) defines T_c to be valid in any configuration, whereas the definition in eq. (4.18) corresponds only to that of the balanced configuration, for which $c_y = c$.

The hyperbolic sine and cosine terms in eq. (4.17) consist of both exponential decay and growth terms, of which the presence of the latter conform with system instability. Note how all growth terms have an identical time constant (\check{T}_c), and how the term of the particular solution is directly scaled to the offset angle ϕ , from which the following can be concluded. The system will fall back ($q_1 \rightarrow \infty$) if the offset angle ϕ is chosen too big such that the system does not reach its equilibrium point, indicating that the robot is unable to recover. It will fall forward ($q_1 \rightarrow -\infty$) if ϕ is chosen sufficiently small such that it overshoots the equilibrium point, indicating that the robot is able to recover. ϕ can also be selected such that all exponential growth terms cancel out, for which the system theoretically comes to a halt at the equilibrium point $q_1 = 0$. This value of ϕ corresponds to the theoretical maximum offset angle $\phi_{\max|I}$ from which the robot can still recover. Therefore, by defining $\phi_{\max|I}$ to be the value of ϕ for which all of the growth terms cancel, we find that

$$\phi_{\max|I} = A \left(1 - \frac{\check{T}_c}{\bar{T}_m} \right) = \frac{\bar{\tau}_s \bar{T}_m \check{T}_c^2}{\check{I}_0 (\bar{T}_m + \check{T}_c)} \quad (4.19)$$

$$= \frac{2n\eta K_m \sqrt{P_m} J_2 \check{T}_c^2}{\check{I}_0 (I_2 + \check{T}_c n^2 \eta K_m^2)}. \quad (4.20)$$

The solution for q_1 in eq. (4.17) then simplifies to

$$q_1(t) | \phi_{\max|I} = A \left(e^{-t/\bar{T}_m} - \frac{\check{T}_c}{\bar{T}_m} e^{-t/\check{T}_c} \right). \quad (4.21)$$

Equation (4.19) is the formula for the maximum recovery angle computed according to approach I. For any given mechanism and motor, there exists an optimal transmission ratio, $n_{\text{opt}|I}$, at which $\phi_{\max|I}$ reaches its maximum value, referred to as $\phi_{\text{opt}|I}$. The values of these two quantities can be found by solving $d\phi_{\max|I}/dn = 0$ for $n > 0$. In computing

$d\phi_{\max|I}/dn$, remember that \bar{T}_m , $\bar{\tau}_s$ and \bar{K}_m are dependent on n (see eq. (4.20)). We find:

$$\begin{aligned} \frac{d\phi_{\max|I}}{dn} &= \frac{2I_2\eta K_m \sqrt{\bar{P}_m} \check{T}_c^2 (I_2 - n^2\eta K_m^2 \check{T}_c)}{\check{I}_0 (\check{T}_c n^2 \eta K_m^2 + I_2)^2} \\ &= \frac{\bar{\tau}_s \bar{T}_m \check{T}_c^2 (\bar{T}_m - \check{T}_c)}{\check{I}_0 (\bar{T}_m + \check{T}_c)^2 n}. \end{aligned}$$

By solving $d\phi_{\max|I}/dn = 0$ for $n > 0$ we find the one solution

$$\boxed{\bar{T}_m = \check{T}_c} \quad (4.22)$$

$$n_{\text{opt|I}} = \frac{1}{K_m} \sqrt{\frac{I_2}{\check{T}_c \eta}} = \frac{1}{K_\tau} \sqrt{\frac{I_2 R}{\check{T}_c \eta}}. \quad (4.23)$$

It suggests that *the time constant of toppling (\check{T}_c) be equal to the effective motor time constant (\bar{T}_m) to optimise the balancing performance*. By substituting $n = n_{\text{opt|I}}$ in eq. (4.20) we find the optimal offset angle

$$\phi_{\text{opt|I}} = \frac{\check{T}_c \sqrt{\check{T}_c I_2 \bar{P}_m}}{\check{I}_0} \quad (4.24)$$

$$= \frac{I_2^{1/2} \bar{P}_m^{1/2}}{(mcg)^{3/4} \check{I}_0^{1/4}}. \quad (4.25)$$

For the second approach (II), using apparent inertia I_a , we find the following solutions equivalent to eqs. (4.19), (4.23) and (4.24):

$$\phi_{\max|II} = \frac{\bar{\tau}_s \bar{T}_{ma} \check{T}_c^2}{\check{I}_0 (\bar{T}_{ma} + \check{T}_c)} \quad \text{with } \bar{T}_{ma} = \frac{I_a}{\bar{K}_m^2} \quad (4.26)$$

$$n_{\text{opt|II}} = \frac{1}{K_\tau} \sqrt{\frac{I_a R}{\check{T}_c \eta}} \quad (4.27)$$

$$\phi_{\text{opt|II}} = \frac{\check{T}_c \sqrt{\check{T}_c I_a \bar{P}_m}}{\check{I}_0} = \frac{I_a^{1/2} \bar{P}_m^{1/2}}{(mcg)^{3/4} \check{I}_0^{1/4}}. \quad (4.28)$$

For any given actuator, there is a specific value of ϕ such that the output torque exactly balances the gravity torque. This value equals

$$\phi_{\text{bal}} = \frac{\bar{\tau}_s \check{T}_c^2}{\check{I}_0} = \frac{\bar{\tau}_s}{mcg} \quad (4.29)$$

and corresponds to solving eq. (4.19) for $\bar{T}_m \gg T_c$, e.g. for $n \rightarrow 0$.

Since $\phi_{\text{bal}} > \phi_{\max}$, and since calculating ϕ_{bal} requires fewer arguments than eq. (4.19), eq. (4.29) provides a useful tool for finding an upper limit on the maximal recovery angle.

4.6. VALIDATION

The relations that were found for both approaches can be compared with results from simulation studies. They can also be compared with other theories, such as the velocity gain.

4.6.1. SIMULATION SETUP

Simulation studies are done in Matlab for approximately 50 different fixed values each of I_2 and n , which provide enough points for calculating accurate cubic splines. Other system parameters are fixed at values listed in table 4.1. These parameters correspond to a particular configuration of Tippy that is used for a series of experiments in chapter 5, and is calculated from CAD data. In Tippy $I_2 = 12.23 \text{ gm}^2$ and $n = 30$. Rotor dynamics are neglected.³

The maximum offset angle $\phi_{\text{max}|s}$ is found using a bisection search for each set of system parameters I_2 and n . The search is done within an interval of $q_1(0) = \phi$ of which it is known that it contains ϕ_{max} and finishes when the interval size is less than 1×10^{-6} rad. At every iteration of the bisection search, a simulation is performed using Matlab's equation solver `ode45` with $q_1(0)$ set to the mid point of the current interval. The simulation stops as soon as it is clear whether q_1 reaches zero or not. The bisection search is then repeated for all combinations of parameters I_2 and n . For each value of I_2 , the corresponding values for $\phi_{\text{opt}|s}$ and $n_{\text{opt}|s}$ are found by finding the maximum of a cubic spline through data points of $\phi_{\text{max}|s}$ for different values of n . Numerical integration is executed with a relative tolerance of 1×10^{-6} .

Actuator properties		
K_T	$18.4 \times 10^{-3} \text{ N mA}^{-1}$	DCX22S-14W-24V (Maxon Motor, 2017)
R	3.69Ω	DCX22S-14W-24V (Maxon Motor, 2017)
J_0	$35.9 \times 10^{-3} \text{ A}$	DCX22S-14W-24V (Maxon Motor, 2017)
U_{max}	30.0 V	8 LiPo cells ⁴
η	0.7	CSF-8-1U-CC-F (Harmonic Drive AG, 2014)
n	(varied)	in Tippy $n = 30$
Dynamic properties		
\tilde{J}_0	0.217 kgm^2	
mc	0.716 kgm	
g	9.81 m s^{-2}	
I_2	(varied)	in Tippy $I_2 = 12.23 \text{ gm}^2$

Table 4.1: Model Parameters, based on Tippy's CAD model, corresponding to the kinematic configuration that is used for experiments in chapter 5.

³Tippy's rotor inertia equals $I_{\text{rot}} = 8.75 \times 10^{-4} \text{ gm}^2$. Its reflected inertia on the reaction wheel is significantly lower than that of the crossbar itself: $n^2 I_{\text{rot}} \ll I_2$.

⁴The nominal voltage of eight LiPo cells is specified as 29.6 V (3.7 V per cell), but is found to be slightly higher than this in practice, up to a maximum (fully charged) voltage of 33.6 V (4.2 V per cell). During most of their discharge cycle, the total battery voltage would range from ca. 30 V to 32 V, the former of which is the assumed voltage for simulations in this chapter.




Figure 4.2: ϕ_{max} plotted in colour as a function of the transmission ratio n for various reaction wheel inertias I_2 . Values of I_2 are written in colour next to corresponding plots, in gm^2 . ϕ_{opt} is plotted in black. Simulation results are solid ($\phi_{\text{max|s}}$, $\phi_{\text{opt|s}}$), algebraic results are dashed ($\phi_{\text{max|I}}$ and $\phi_{\text{opt|I}}$) and dash-dotted ($\phi_{\text{max|II}}$). Algebraic predictions are accurate for low values of I_2 and n .




Figure 4.3: Results of algebraic method I ($\phi_{\text{opt|I}}$) match very well with those of the exact simulations ($\phi_{\text{opt|s}}$) and slightly poorer—but still well—with those of simulations with a linearised component of the gravity force ($\phi_{\text{opt|sl}}$). Results of algebraic method II ($\phi_{\text{opt|II}}$) match much poorer with both $\phi_{\text{opt|s}}$ and $\phi_{\text{opt|sl}}$.

4.6.2. SIMULATION RESULTS

Figure 4.2a plots the maximum offset angle as a function of n for various reaction wheel inertias I_2 in colour. Differently coloured lines correspond to different values of I_2 , where the value of I_2 is increased quadratically. The optimal offset angle is plotted in black. ϕ_{bal} (eq. (4.29)) is plotted in grey. Solid lines correspond to simulation studies ($\phi_{\text{max|s}}$, $\phi_{\text{opt|s}}$). Dashed and dashed-dotted lines correspond to the algebraic approaches I and II respectively.

Figure 4.2a clearly shows that both approaches (I and II) provide good estimates of the maximum offset angle, but that the first approach (using reaction wheel inertia) outperforms the second one (using apparent inertia), i.e. $\phi_{\max|I}$ provides a better estimate than $\phi_{\max|II}$. A possible explanation is the following. The apparent inertia is less than the actual inertia, so it predicts a faster acceleration. In the absence of gravity, this prediction is exact. However, gravity acts to retard the acceleration, with the result that the acceleration predicted by using the actual inertia of the reaction wheel is closer to the acceleration computed in the simulation.

The observation that $\phi_{\max|I}$ matches $\phi_{\max|s}$ so well is at least partly due to a numerical coincidence. Differences between algebra and simulation result from (1) the linearised component of the gravity force and (2) simplifications of the effect of pendulum dynamics on the reaction wheel dynamics. These errors both tend to partly cancel each other out, which becomes evident if simulation studies are repeated with a linearised component of the gravity force such that the difference between algebra and simulations is only due to (2) simplifications of the effect of pendulum dynamics on the reaction wheel dynamics. Figure 4.3 plots variously obtained values of ϕ_{opt} as a function of I_2 . It also plots $\phi_{\max|sl}$, which are values obtained through simulation with linearised EoM at $q_1 = 0$ (eq. (4.10)) but otherwise similar to $\phi_{\max|s}$. It shows that $\phi_{\max|I}$ matches better with $\phi_{\max|s}$ than $\phi_{\max|sl}$, which supports the theory that (1) and (2) partly cancel each other out. Nevertheless, $\phi_{\text{opt}|I}$ provides a much more accurate prediction of both $\phi_{\max|s}$ and $\phi_{\max|sl}$ than $\phi_{\text{opt}|II}$, which strengthens the conclusion that the first approach provides more accurate results than the second.

ϕ_{bal} in fig. 4.2a, plotted as function of n , clearly illustrates an upper limit on the maximal recovery angle ϕ_{\max} . The other curves are also asymptotic to it.

Figure 4.2b shows the same plots as fig. 4.2a but with a relative logarithmic x-axis, where n is scaled by $n_{\text{opt}|I}$ for every value of I_2 . It shows how the maximum offset angle as a function of n is symmetric about its optimum. For all practical purposes, it is useful to know how the system behaves around the optimum, because the transmission ratio (and possibly other system parameters) cannot always be chosen freely. It follows directly from eqs. (4.19) and (4.24) that

$$\phi_{\max|I} = \frac{2z}{z^2 + 1} \phi_{\text{opt}|I} \quad \text{with} \quad z = \left(\frac{n}{n_{\text{opt}|I}} \right)^{\pm 1}.$$

For example, if n is double ($z = 2$) or half ($z = 0.5$) the optimum $n_{\text{opt}|I}$ then $\phi_{\max|I}$ is still 80% the optimum $\phi_{\text{opt}|I}$.

4.6.3. VELOCITY GAIN

The angular velocity gain of the RWP equals

$$G_\omega = \frac{\Delta \dot{q}_1}{\Delta \dot{q}_2} = -\frac{I_2}{I_0} = -\frac{I_a}{\tilde{I}_0}. \quad (4.30)$$

Note that the numerator I_2 respectively I_a also appears in the numerator of eqs. (4.25) and (4.28). Also \tilde{I}_0 (similar to I_0) appears in the denominator of both equations. However, the powers are different (i.e. I_a^1 versus $I_a^{1/2}$ and \tilde{I}_0^1 versus $\tilde{I}_0^{1/4}$). This indicates that the

velocity gain is positively correlated with the maximal recovery angle, but not linearly so. Moreover, eqs. (4.25), (4.28) and (4.30) confirm that the velocity gain is missing significant effects, as it does not include an actuator term in the numerator ($\bar{P}_m^{1/2}$ in eqs. (4.25) and (4.28)) and a gravitational term in the denominator ($(mcg)^{3/4}$ in eqs. (4.25) and (4.28)).

4.7. EFFECTIVE OFFSET ANGLE

The balancing performance has previously been defined as the maximum offset angle from which the robot can recover, according to which eq. (4.16) has been solved with the initial conditions $q_1(0) = \phi$ and zero velocity ($\dot{q}_1(0) = 0$). If instead the case of a non-zero initial velocity $\dot{q}_1(0) = \psi$ and zero offset angle ($q_1(0) = 0$) is considered, then the corresponding solution is

$$q_1(t) = A \left(e^{-t/\tilde{T}_m} - \cosh(t/\tilde{T}_c) + \frac{\tilde{T}_c}{\tilde{T}_m} \sinh(t/\tilde{T}_c) \right) + \underbrace{\psi \tilde{T}_c \sinh(t/\tilde{T}_c)}_{\text{particular solution}}. \quad (4.31)$$

This solution differs from the original solution only in the particular solution: cosh is replaced by sinh and ϕ by $\psi \tilde{T}_c$. Identical to the approach followed in section 4.5, i.e. by solving for the exponential growth terms to cancel out, it follows that

$$\psi_{\max|I} = \frac{\phi_{\max|I}}{\tilde{T}_c}. \quad (4.32)$$

The initial balance error could be defined by a combination of an initial offset angle $q_1(0) = \phi$ and initial velocity $\dot{q}_1(0) = \psi$. Their combined contribution to the balance error can be quantified by taking the sum

$$\boxed{\phi' = \phi + \psi \tilde{T}_c}. \quad (4.33)$$

ϕ' is referred to as the *effective offset angle*. Note that the principle of ϕ' is similar to that of the extrapolated center of mass (XCoM) ξ (Hof, 2008), where

$$\xi = x + \frac{v_x}{\omega_0}$$


with x the horizontal position of the robot's CoM with respect to the support point (joint 1), v_x the horizontal velocity of the CoM and $\omega_0 = \tilde{T}_c^{-1}$ the inverse of the time constant of toppling. The difference between ϕ' and ξ is the use of angular instead of linear coordinates.

A particular combination of ϕ and ψ is of frequent occurrence in practice, namely that resulting from unactuated toppling of the robot from a near-balanced configuration (i.e. from $q_1(0) \approx \dot{q}_1(0) \approx 0$), which would for example occur if the robot is delayed in observing itself losing its balance. For this behaviour $\dot{q}_1 \approx q_1/\tilde{T}_c$. If $\phi = \phi_c$ and $\psi = \phi_c/\tilde{T}_c$ are chosen, then ϕ_c corresponds to the offset angle resulting from unactuated toppling from a near-balanced configuration, for which follows from eq. (4.33) that $\phi' = 2\phi_c$. This implies that the maximum angle to recover from unactuated toppling is

approximately half the maximum recovery angle from a static offset, i.e. $\phi_c = \phi_{\max}/2$. Arguably, the former is of more practical importance than the latter, for which one could alternatively consider using $\phi_{\max}/2$ instead of ϕ_{\max} to describe a robot's physical ability to balance.

Like the derivation of eq. (4.19), eqs. (4.32) and (4.33) are derived from simplified dynamic equations that differ from the true dynamics. As such, it is required to verify their accuracy. To check for the accuracy of eq. (4.32), simulations were done leading to the results shown in fig. 4.4a, equivalent to the simulations done for eq. (4.19) leading to the results shown in fig. 4.2a. In order to compare the two graphs, fig. 4.2a plots the value of ψ_{\max} scaled with T_c . The results shown in fig. 4.4a are very similar to those shown in fig. 4.2a, which indicates that also eq. (4.19) is accurate.

The validity of eq. (4.33) was checked for the case $I_2 = 12.23 \text{ gm}^2$ and $n = 30$, by taking six evenly-spaced values for $q_1(0)$ ranging from 0 to ϕ_{\max} and, for each of these initial offsets, finding (by simulation) the accurate value of ψ . The results are plotted in fig. 4.4b. Observe that the accurate values of ψ are very close to those predicted by the formula. For example, when $q_1(0) = 0.2\phi_{\max}$ the simulation obtains $\psi T_c = 0.805\phi_{\max}$, which is very close to the $0.8\phi_{\max}$ predicted by the formula. Further evidence that eq. (4.33) does indeed represent an effective offset can be seen from the fact that all of the curves in fig. 4.4b converge to the same exponential decay.



(a) Maximum initial velocity $\dot{q}_1(0) = \psi_{\max}$, scaled by T_c for comparison with the previous studies that use a maximum offset angle $q_1(0) = \phi_{\max}$ instead.

(b) Pendulum trajectory as a function of time for various initial conditions $q_1(0) = \phi$ and $\dot{q}_1(0) = \psi$, with $q_1(t)$ drawn solid and $\dot{q}_1(t)$ dashed.

Figure 4.4: Results for an initial non-zero velocity are very similar to those for an initial non-zero position. The effective offset angle remains approximately constant, i.e. $\phi' = \phi + \psi T_c \approx \text{const}$.

4.8. CONCLUSION

This chapter has focused on the problem of designing for a robot to physically be a good balancer, by analysing its ability to recover from a balance offset to a balanced config-

uration by operating the actuators in saturation, for which the study is an example and enabler for designing at the performance envelope. The study assumes a simple balance robot model known as the reaction wheel pendulum with a realistic actuator model, resembling that of a DC motor with transmission mechanism. The model also resembles that of Skippy and Tippy in their frontal planes. The contributions are (1) a formula that describes the robot's largest balance offset angle from which it can still recover as function of its physical parameters (eq. (4.19)), (2) a design criterion to maximize this angle (eq. (4.22)) and (3) a generalisation of this angle as a performance measure to a combined position and velocity offset (eq. (4.33)), similar to the extrapolated centre of mass. The novelty is that this work takes both realistic actuator properties into account and a time trajectory of the full balance recovery motion, including the effect of gravity, whereas previous work on performance measures do not consider actuator characteristics, or are instantaneous, or both. The algebraic formulations of the resulting relations have been the result of model simplifications, but accurate numeric simulations have demonstrated their accuracy. The functions provide insight into how various physical parameters influence balance performance. In particular, the design criterion states that the robot's time constant of toppling should be equal to the actuator's time constant under load, which can be used as guideline for balance robot design initiation or optimisation.

5

TIPPY: A BALANCING-ONLY PROTOTYPE

This chapter introduces, describes and analyses a new balancing machine called Tippy. Tippy is designed and built to precede the balancing and hopping machine Skippy for the purpose of (1) investigating its balancing performance, (2) testing new balancing algorithms for 2D and 3D balancing and (3) being a design exercise for Skippy. The design criterion from chapter 4 has been used to design and confirm Tippy to be a good balancer, and the balancing theory that has been introduced in section 3.3 has been successfully demonstrated on a simplified planar configuration. Important lessons were put forth from Tippy's design and performance that are useful for revisions of Skippy's design.

5.1. INTRODUCTION

Chapter 2 has highlighted the importance of balancing for Skippy and legged robots in general, which has inspired research on a new balancing control theory (section 3.3) and physical balancing performance (chapter 4). However, the new theories have not been brought to practice. Despite having been tested in realistic simulation environments that include electric motor models (chapter 4) or motor saturation, time delays, sensor noise, sensor bias and model errors (Featherstone, 2017), unknown caveats could remain. The new objective is to discover the possibilities and limitations of balancing performance on a *real* system. Since a simple system to fulfil that purpose was not available, such a system would have to be designed and built. In addition, the exercise of designing that system could be useful for designing Skippy.

Acknowledgements — Parts of this chapter have been published in *IEEE/RSJ International Conference on Intelligent Robots and Systems 2018* (Driessen et al., 2018) and *IEEE International Conference on Robotics and Automation 2019* (Driessen et al., 2019), and includes excerpts from those papers. In particular, fellow PhD students Antonios E. Gkikakis and B. Roodra P. Singh have been largely responsible for implementing the balance controller, the control hierarchy, the motor velocity servo and velocity estimators, as described in section 5.5.2, and have contributed to performing the experiments described in section 5.6.2. In addition, Antonios E. Gkikakis has assisted in setting up the recovery angle experiments described in section 5.6.1.

A desire for implementing the balancing theories in practice, and for learning from that implementation for the design of Skippy, has resulted in a purpose-built 3D balancing machine called Tippy, depicted in fig. 5.1. In particular, Tippy serves the following goals:

1. Investigating the balancing performance of itself and – since it shares a similar structure – also of Skippy.
2. Performing balancing experiments for the purpose of discovering whether the new balancing theory described in section 3.3 performs as well in practice as it does in simulation.
3. Being a mechanical and electronic design exercise for Skippy.




Figure 5.1: Photos of Tippy. The right photo shows Tippy's frontal plane, in which it resembles an RWP.

OUTLINE

This chapter focuses both on the design of Tippy and on the experiments that have been carried out in collaboration with colleagues. Tippy's design requirements are established in section 5.2, based on which critical decisions are described in section 5.3 to come to a more tangible design. Tippy's design is then described in more detail in section 5.4. Experimental preparations and results are then described in sections 5.5 and 5.6 respectively.

5.2. REQUIREMENTS

Design requirements for Tippy largely stem from (1) the experiments that it should be used for and (2) its desired resemblance to Skippy. Experiments on Tippy include recovery experiments to test the theory described in chapter 4 and balancing experiments using the balancing controller described in section 3.3 and Featherstone (2015a, 2017). The

balancing theory is applicable to simple and complicated 2D and 3D balancers, none of which have yet been experimented with. Moreover, all previous simulation studies assume that all joint angles can be directly measured, which in practice is not typically the case, also not for Skippy: there is no possibility to mount an encoder between the foot and the ground. The system must instead rely on an IMU in order to be fully observable, which likely introduces drift and more delay than a joint encoder. In order to make safe and steady progress, the following experiments are to be carried out in succession:

1. 2D balancing in a plane in which Tippy represents a RWP, as introduced in chapter 4 and depicted in fig. 4.1, using two joint encoders (fully observed). The advantage of the RWP is that it can be described by fewer model parameters, which facilitates system identification, and that it has no joint limits at the actuated joint, which makes the system overall safer to use. The RWP is also the main example treated in Featherstone (2017).
2. 2D balancing in the same plane, but instead using an IMU to replace the encoder in the support joint.
3. 2D balancing in a plane in which Tippy represents an Acrobot, as introduced in section 3.3 and depicted in fig. 3.3.
4. 3D balancing, according to the theory described in Featherstone (2017).

A list of design requirements can be established by incorporating this expanded list of planned experiments with the goals.

1. Tippy's morphology, size and mass distribution closely have to resemble that of Skippy.
2. It is to be physically capable of high-performance balancing, so that a poor performance can be blamed on the controller.
3. It should be easy to operate and safe to itself and users, in order to quickly progress with experiments.

In addition, Tippy is to be a highly configurable robot. It should include the ability to

4. lock several of its joints to become a 2D RWP and in another way to become an Acrobot;
5. change several kinematic and inertial parameters to become a better or worse balancer, and
6. be usable with and without IMU.

5.3. CRITICAL DESIGN DECISIONS

In line with the design requirements, Tippy is designed to have a shape and size similar to that of Skippy. Design decisions of Tippy can be best explained by comparing it with Skippy. Like Skippy, its morphology is essentially an Acrobot with a crossbar on top, depicted in fig. 5.2. The two actuated joints are the bend and swivel joint, which are both actuated by two low-backlash actuators that use brushed DC motors identical to those used in early Skippy simulation studies (section 3.2). However, in contrast to Skippy,




Figure 5.2: A schematic body diagram of Tippy, showing the kinematic tree with joint numbers and names and structural planes. Not to scale. The axes of the second (roll) and third (pitch) joint intersect in reality, but are drawn apart and connected with a dotted line to illustrate their order in the kinematic tree.

Tippy (1) has simplified actuation mechanisms, (2) is fixed to the ground via a 3-DoF series of joints, called the base joint, and (3) is built to be easily customizable. Each of the three topics are elucidated below.

5.3.1. SIMPLIFIED ACTUATION

Skippy's actuation mechanisms include series elasticity and a linkage mechanism with a non-constant transmission ratio. Whereas these features are crucial for the jumping performance, they are less so for balancing. In addition, the efficiency of the transmission ratio is less important, as the focus is not on maximizing power throughput. As a first step to implementing the new balancing theory on a real system, it was decided to apply the following simplifications to the actuation mechanisms:

- Tippy does not have an ankle joint and thus no ankle spring: its leg is stiff. As explained in section 3.1, the ankle spring is mainly responsible for controlling lift-off momentum, which is not relevant for balancing.
- Tippy's bend joint is stiff and it is revolutive, featuring a strain wave gear and motor directly at the joint, whereas the joint in Skippy features a four-bar linkage, ball-screw and a non-linear series elastic spring. A strain wave gear is chosen for the desire of zero backlash and because of in-house availability. The transmission ratio of the strain wave gear in Tippy is approximately equal to the (non-constant)

transmission ratio of the combined four-bar linkage and ball screw in Skippy.

Note that the omission of series elastic springs (both in the bend joint and ankle) should make balancing easier, as a higher stiffness increases the control bandwidth. However, no significant performance change is expected, since both the ankle and main spring of Skippy are designed to be very stiff at the relatively low loads required for balancing.

5.3.2. BASE JOINT

The base joint consists of three revolute joints in series, with all of their three axes coinciding at one point, so that it essentially represents a spherical joint. The base joint allows for

- measuring Tippy's orientation with respect to the ground with encoders in 3D;
- locking/fixing different DoFs independently from each other, such that the robot can be used for planar experiments in both the sagittal and frontal plane, and
- having greater control over the robot's range of motion, by defining the range with mechanical end stops for each of the DoF's independently. This addresses possible safety issues (i.e. prevention of damaging the environment or itself), and revokes the necessity to cover Tippy with shock-absorbing bumpers.

JOINT ORDER

The order of the three revolute joints in the base is yaw, roll and pitch, which rotate the robot in its transverse, frontal and sagittal plane respectively, as illustrated by fig. 5.2. The planes are defined to be local to the coordinate frame of the roll joint. This implies that the frontal plane is always vertical, but that the sagittal plane can be tilted with respect to the vertical.

All of the three base joints can be locked independently. RWP experiments are carried out in the frontal plane, by actuating the swivel joint and fixing the yaw, pitch and bend joint. Acrobot experiments are carried out in the sagittal plane, by actuating the bend joint and fixing the yaw, roll and swivel joint.

The choice behind the order of joints can be explained as follows. The roll joint is aligned before the pitch joint as this allows for performing 2D low-gravity experiments in the sagittal plane, by locking the roll joint under an angle (e.g. a roll angle of 60° with respect to the vertical would imply half the gravity). In addition, a successive alignment of the bend and pitch joint allows for many possibilities for aligning the swivel joint horizontally for RWP experiments, by choosing the right combination of locked joint angles for the pitch and bend joint. Lastly, the yaw joint is the first joint because it is the only base joint with infinite joint range, and it is only used for 3D experiments.

5.3.3. CUSTOMIZABILITY

Tippy is designed to be customizable and versatile such that it can be used for many different experiments. This includes not only the ability to lock various of the joints, but also the ability to easily change body lengths (leg and torso) and joint ranges (the base joint and bend joint), and the ability to add mass. The structural design style of Tippy is inspired by the model construction system Meccano, featuring cheap off-the-

shelf beams with flat surfaces, modified to have many mounting holes at fixed distances. This introduces flexibility for placement of additional mass and changing body lengths and mounting angles.

In addition, Tippy need not be optimised for weight reduction, as it does not need to be able to jump and survive crashes from its jumps. As such, most of Tippy's components are selected or designed with economic shapes and sufficient margin regarding size and strength. Two parts of the robots that require extra attention are the base joint and bend joint:

- The base joint is ideally a real spherical joint ('point' joint). As such, the effect on robot dynamics caused by motion of the various bodies interconnecting the base joint should be kept to a minimum. This implies that base joint should be as compact as possible, and that it should not approach a kinematic singularity.
- The motor of the bend joint is located at the joint, whereas in Skippy it was placed close to the swivel bar. This has an effect on the weight distribution, which is now more favourable for balancing but less for jumping. Since Tippy is a balancing-only robot, this is not a major issue. However, owing to the desire for Tippy to resemble Skippy, the mass of Tippy's bend joint should be kept to a minimum, and it should be possible to add dead weights to Tippy close to the swivel bar.

5.4. DETAILED DESIGN DESCRIPTION

This section describes Tippy's design in more detail, including both its mechanical and electrical design. An introduction of Tippy's kinematic structure has already been given in section 5.3, and is explained in more detail in section 5.4.1.

5.4.1. KINEMATIC STRUCTURE

Tippy is designed to have various easily changeable kinematic properties. Tippy's kinematic structure is defined in such a way to facilitate recalculating its dynamic model if its kinematic properties are changed. For this purpose, the kinematic structure is defined to consist of main bodies and sub-bodies.

Tippy's main bodies are its six rigid bodies (including the base to which it is mounted to the ground), all connected in one serial chain via the five joints that are depicted in fig. 5.2. Joints and bodies are numbered and named, listed in table 5.1. The five joint positions are described by joint variables and are elements of the vector \mathbf{q} :

$$\mathbf{q} = [q_1 \quad q_2 \quad q_3 \quad q_4 \quad q_5]$$

Joint i connects main body i and $i - 1$. Joint 1 connects to a fixed ground, which can be treated as body 0.

The main bodies are divided into sub-bodies. Sub-bodies are also numbered and named, listed in table 5.1 and shown in fig. 5.4. All sub-bodies individually have fixed dynamic properties, but their kinematic relations are defined by physically tunable¹ kinematic parameters a to e , which are listed in table 5.2 and depicted in fig. 5.3 for

¹By tunable is meant the ability to easily change parameters in between – but not during – individual experiments and simulations

an arbitrary selection of a to e , which shows Tippy to scale in its sagittal plane. In this figure $\mathbf{q} = 0$, and it also shows coordinate systems of each of the bodies (i) and sub-bodies (i_j). The coordinate frame of a body's first sub-body (i_1) coincides with that of the corresponding body (i).

The sub-bodies correspond to the (sub-)assembly structure of Tippy's CAD model, from which inertial properties can be easily extracted. The dynamic and kinematic properties of each main body are essentially a function of inertial properties of sub-bodies and the (tunable) kinematic relations. Having properties of all sub-bodies readily available facilitates recalculation of kinematics and dynamic properties when the system is physically altered by changing any of the tunable kinematic parameters. These functions take the shape of series of coordinate transforms, and can be found in Tippy's Mechanical Reference manual (Driessen, 2017).

A 5-DoF rigid-body dynamic model can be constructed from the main bodies. Two extra bodies can be added to the model to include drive axle inertia of the bend and swivel actuators, which is required to perform accurate 3D simulations. In addition, these bodies can be used to include transmission models that take finite stiffness and backlash into account.

i	joint (axis)	body	j	sub-body
0	–	base	1	table
			2	yaw base
1	yaw (z)	roundabout	1	roundabout
2	roll (y)	cross axle	1	cross axle
3	pitch (x)	leg	1	output clevis
			2	lower leg extensor
			3	upper leg extensor
			4	hip adapter
4	bend (x)	torso	1	hip housing
			2	torso extensor
			3	head housing
5	swivel (y)	crossbar	1	crossbar

Table 5.1: Tippy's joints and corresponding bodies, and sub-bodies.

5.4.2. WORK SPACE

The work space of Tippy is limited by the finite range of motion of joints, and can be further limited by fixing joints. The joints' ranges of motion are determined by end stops, and joints can be fixed by installing joint locks. Ideally, the range of motion is set in such a way that the robot cannot collide with itself or the environment. Depending on the experiment (swivel, bend, 3D), this can be achieved by choosing various combinations of end stop positions and joint locks.

END STOPS

Joints 2 to 4 (roll, pitch and bend) have end stops, and thus a finite range of motion, whereas joints 1 and 5 (yaw and swivel) provide infinite rotation in both directions

KP	acceptable values	notes
<i>a</i>	$[-37.5 \ 37.5]^\circ$	These values hold for currently installed pitch stop brackets, but they can be changed to other values by laser-cutting new brackets.
<i>b</i>	$[36 \ 48 \ 60 \ \dots \ 156] \text{ mm}$	The leg can be made longer by replacing the leg beam.
<i>c</i>	$[-45 \ -22.5 \ 0 \ \dots \ 45]^\circ$	
<i>d</i>	$[-45 \ -22.5 \ 0 \ \dots \ 45]^\circ$	It is recommended that $c+d \geq -45^\circ$ such that the leg and torso cannot hit each other (assuming the bend end stop to be at -30°).
<i>e</i>	$[0 \ 10 \ 20 \ 30]^\circ$	

Table 5.2: Kinematic parameters.

(cables do not cross these joints either). The range of motion of all five joints are defined in table 5.3, and can furthermore be viewed in fig. 5.5.

The joint range of both the pitch and roll joint is symmetric about zero, and both do not allow either joints to approach their singular positions at $\pm 90^\circ$. The base joint is designed so that the total joint range of pitch can exceed 90° without internal collisions, so that Tippy's leg can be completely horizontal (crouched) and vertical (stretched). This is realised by mounting the leg asymmetrically to the pitch joint, by choosing a non-zero value for *a*, as listed in table 5.2 and shown in fig. 5.3, where $a = 37.5^\circ$ for Tippy's default mounting brackets.

<i>i</i>	joint	$q_{i,\min} (^\circ)$	$q_{i,\max} (^\circ)$	notes
1	yaw	$-\infty$	∞	
2	roll	-22.5	22.5	Can be increased or decreased by printing different end stops.
3	pitch	-52.5	52.5	Can be increased or decreased by printing different end stops.
4	bend	-30	240	The range of motion can be decreased in steps of 30° on both sides by displacing the end stops, up to a minimum range of motion of 90°
5	swivel	$-\infty$	∞	

Table 5.3: Range of motion of all joints.

JOINT LOCKS

Joints 1 to 4 (yaw, roll, pitch and bend) can be locked in a number of positions, effectively decreasing the number of DoF's to allow for various types of planar experiments. Joint 5 (swivel) does not have a joint lock, because the crossbar is symmetric and the transmission offers sufficient friction to consider it locked if not actuated. See table 5.4 for information regarding joint locks positions.




Figure 5.3: Right view of Tippy, scale 1:6.25. Fixed dimensions displayed in blue and tunable dimensions in green with parameters a to e , with lengths in mm and angles in $^\circ$. All joint angles are zero ($q = 0$). Right-handed yz coordinate frames of (sub-)bodies in red, with x pointing out of the drawing.

Fixing a joint lock is done using several M3 screws. The mechanical reference manual contains more information regarding the involved parts and assembly instructions (Driessen, 2017).

SWIVEL CONFIGURATION

For planar swivel experiments, joints 1, 2 and 4 (yaw, pitch and bend) are locked. To obtain a true planar set-up, the swivel axis should be horizontal, which can be achieved by choosing a particular set of kinematic parameters and joint locks for which the following relation holds:

$$q_{13} + a + c + q_{14} + d + e = 0 \quad (5.1)$$

5.4.3. ACTUATORS

Both the bend and swivel joint are actuated by the motors that were previously selected in the study conducted by Driessen (2015) for Skippy. Each motor is connected to a strain




Figure 5.4: All sub-bodies and their numbers (i_j). The bodies are not all to the same scale.




Figure 5.5: Joint ranges as defined by end stops.

joint	KP	acceptable values	notes
yaw	q_{11}	(0)	The value does not matter. Aluminium part.
roll	q_{12}	0 and more	The roll and pitch lock use exactly the same method of locking the joint.
pitch	q_{13}		The part used for locking either joint is printed and can be printed for any and multiple locking angles.
bend	q_{14}	$[-30 \ -15 \ 0 \ \dots \ 240]^\circ$	Aluminium part.

Table 5.4: Joint locks positions.

wave gear from Harmonic Drive and each is driven by a Pololu motor driver.

BEND JOINT ACTUATION

The bend joint is actuated by Maxon's brushed RE35-24V motor (Maxon Motor, 2017) and Harmonic Drive's CSD-14-2A strain wave gear (Harmonic Drive AG, 2014), as depicted in fig. 5.6. The strain wave gear has a reduction gear ratio of 100:1, which corresponds approximately to the average reduction ratio by the linkage mechanism and ball screw. It is computed as the product of the following two transmission ratios:

- The reduction ratio for a ball screw with a pitch of p is $2\pi/p \text{ rad m}^{-1}$.
- The average reduction ratio of the four bar linkage including the lever with a length of l_1 is approximately $\sqrt{2}l_1/\pi \text{ m rad}^{-1}$, corresponding to the simplified mechanism shown in fig. 3.1b with a lever range of 90° .

The desired transmission ratio can thus be approximated by

$$n = \frac{2\sqrt{2}l_1}{p}.$$

For a lever length of $l_1 = 7 \text{ cm}$ and a pitch of $p = 2 \text{ mm}$ this implies $n = 99$, which is closest to the available gear reduction ratio of 100:1 for strain wave gears. The CSD series

has been chosen as it is the short series, which helps to reduce protrusion of Tippy from its sagittal plane (as caused by the long RE35 motor). Furthermore, in order to reduce Tippy's mass at the bend joint, it has been decided to use a component set rather than a complete housed gear unit and to choose a relatively small size (size 14).

- A component set consists of only the core components of a strain wave gear, consisting of the flex spline, circular spline and wave generator. Using a component set requires a custom gear unit housing, which can be designed to save significant weight. The component set of a strain wave gear size 14 weighs in at 60 g, whereas the lightest COTS housed gear unit of similar size (CSF-14) has a mass of 290 g. The custom housed gear unit depicted in figs. 5.4i and 5.6 (excluding the motor, mounts and end stops) has approximately the same mass, but is designed so that the encoder, end stops and torso can be mounted directly on or in the housing, leading to an overall weight reduction of at least 0.1 kg. The output bearings of the housing are furthermore stronger than that typically found in a COTS gear unit, as required for withstanding torsional loads that could result from crashes at the swivel end stop in the base joint. The designing of a custom housing furthermore serves as an exercise for designing Skippy, which is planned to have a custom housing for the strain wave gear of the swivel actuator.
- Size 14 can withstand momentary peak torques of 31 Nm. This is significantly lower than the stall torque of the RE35 motor, which – for an expected gear efficiency of $\eta = 0.7$ and voltage of $U = 30\text{V}$ – is 105 Nm at the output shaft of the strain wave gear. However, recall that the choice behind the motor stems from the required jumping performance of Skippy; it is significantly overpowered for balancing and as such need not be used to its fullest capability on Tippy.²

The output bearings are 7302B angular contact ball bearings, mounted back-to-back, which have a 6.7 kN static load rating and an effective moment arm of 18.7 mm. The corresponding torsional load rating is approximately 250 Nm, which is sufficiently high to endure crashes on the swivel end stops.³

SWIVEL JOINT ACTUATION

The swivel joint is actuated by Maxon's brushed DCX-22-S-GB-24V motor (Maxon Motor, 2017) and Harmonic Drive's CSF-8-1U-CC-F strain wave gear with a gear reduction ratio of 30:1 (Harmonic Drive AG, 2014). The choice behind designing the swivel joint to have a gear reduction ratio of 30:1 (the lowest available), together with the crossbar to have a (default) inertia of 12.23 gm^2 , is based on maximizing the balancing performance, using eq. (4.23). The crossbar inertia also approximately resembles that of having the batteries placed at the tip of Skippy's crossbar, at a radius of 0.25 m, as in fig. 3.1a. Using eq. (4.23) requires knowledge of various of Tippy's inertial and kinematic parameters. However,

²If instead the strain wave gear would be chosen based on stall capabilities of the motor, a size 25 would have to be chosen instead, which is capable of withstanding momentary peak torques of 152 Nm, but it is 4 times heavier than a size 14.

³In a worst case crash scenario, Tippy is fully stretched and the motor accelerates the robot to a velocity higher than it would obtain from unactuated toppling. It can be calculated that, for end stop compression of several millimetres, the moment in the output bearing of the hip is as large as approximately 100 Nm.




Figure 5.6: Tippy's bend joint actuator assembly, comprising a custom housing for the strain wave gear, strong output bearing, encoder and customizable mounting options for the leg, torso and bend joint end stops with limit switches. This assembly depicts the hip adapter (fig. 5.4h) and hip housing (fig. 5.4i) sub-bodies.

these are a function of Tippy's configuration and furthermore assume the crossbar axis to be aligned horizontally.

To maximize Tippy's balancing performance and to reduce impacts at end stops during swivel experiments, Tippy has been configured to be close to the ground (reducing c). Selection of the transmission ratio and crossbar inertia are justified accordingly using eq. (4.23).

Tippy's kinematic parameters are assumed to take their default values $a = 37.5^\circ$, $c = -22.5^\circ$, $d = -22.5^\circ$ and $e = 30^\circ$, as depicted in fig. 5.3. The bend and pitch joint angles are then assumed to be $q_{13} = -7.5^\circ$ and $q_{14} = -15^\circ$ respectively: the former chosen as low as feasible so that the leg and torso do not collide, and the latter chosen so as to align the crossbar with the horizontal, using eq. (5.1).

This configuration is assumed for simulation studies in section 4.6 and used for swivel experiments in section 5.6. It is also the configuration of Tippy depicted in figs. 5.1 and 5.7.

The gear reduction ratio and crossbar inertia are designed to be $n = 30$ and $I_2 = 12.23\text{gm}^2$. $n = 30$ is the lowest gear reduction available gear ratio for the selected series of strain wave gears. For $n = 30$, $I_2 = 12.23\text{gm}^2$ and the parameters listed in table 4.1, the calculated maximum balance recovery angle is 13.9° . The transmission ratio corresponds close to the optimal $n_{\text{opt}} = 32.9$, see also the green line in fig. 4.2a. This recovery angle is quite large, despite the relatively light motor, which takes up only 2.3% of the robot's total mass: the motor weighs 0.066 kg and Tippy (bodies 2 to 5) weighs 2.88 kg according to CAD data.

In contrast to the use of a component set and custom housing like in the bend joint, this unit is purchased as a complete gear unit housing, as there is no strict weight limit on Tippy's head. In fact, as discussed in section 5.3, to more closely resemble the weight distribution of Skippy, it is desired this part of Tippy to be slightly heavier to compensate for the heavy bend joint.

5.4.4. BRACE

In early stages of swivel balancing experiments it was discovered that the robot's low structural stiffness limits the balance performance. This is an issue especially for swivel experiments, where the leg, torso and various base joints are locked together, creating a single (big) rigid body that has a high natural frequency. In response, a brace has been developed and installed on Tippy to stiffen the robot for swivel experiments. The brace consists of a beam (of identical sizing as beams of leg extensors) and an adapter. The beam is directly mounted on the output clevis of the base joint and the adapter is clamped on the torso close to the crossbar, as depicted in fig. 5.7.

The experiments reported in section 5.6 are done on Tippy including the brace. The brace adds 217 g to the robot's mass, increasing the robot's mass from 2.88 kg to 3.10 kg (+7.5%) according to CAD data. Corresponding updated parameters are $mc = 0.747\text{kgm}$ and $\check{I}_0 = 0.223\text{kgm}^2$ (previously $mc = 0.761\text{kgm}$ and $\check{I}_0 = 0.223\text{kgm}^2$, see table 4.1). Since these are only marginal changes, they do not significantly affect the robot's performance.




Figure 5.7: Tippy depicted with its brace, consisting of a brace mount and beam, designed and installed to improve its structural stiffness for swivel balancing experiments (in the frontal plane).

5.4.5. ELECTRONICS

Tippy is not tethered, and all computations are performed real-time on board. Also the electronics are intended to be a design exercise for Skippy, and as such should be capable of handling the high power required for jumping, even if this is not required for balancing. The strategy behind choosing Tippy's electronics is furthermore to keep everything as simple as possible, whilst excelling in real-time computations. This is achieved by using a single microcontroller unit (MCU) that takes care of both lower and higher levels of control. Having a single MCU allows for the use of simple synchronized communication

protocols and interrupt routines, and high-frequency sampling, subsequently reducing noise and delays.

All computations are performed on Texas Instrument's C2000 Delfino MCU F28377S LaunchPad Development Board, which is specialised in real-time computations and offers sufficient peripherals to directly connect all electronics, including all encoders and motor drivers. Other electronics include step-down voltage regulators, a custom printed circuit board (PCB), called TP-PCB (Tippy PCB), a user control interface and a Raspberry PI Zero. Figure 5.8 shows a schematic diagram that connects all of the electronics, and additional explanations for each of the electronics follow below.




Figure 5.8: Electronic connectivity diagram of Tippy. Texas Instrument's C2000 Delfino MCU F28377S Launch-Pad Development Board plays a dominant role in Tippy's electronics, as it takes care of both low level (e.g. motor servos) and higher level routines (e.g. the balancing controller).

BATTERIES

Tippy is equipped with two 4-cell LiPo batteries from LiPoTech⁴ that power the robot with a nominal $8 \times 3.7 = 29.6\text{V}$. The batteries have a capacity of 1300 mAh and a peak discharge rate of 60C, implying a maximum drawing current of 78 A. The experienced voltage is measured to be slightly higher than its nominal value for healthy cells, and peaks at $8 \times 4.2 = 33.6\text{V}$ when fully charged. For most experiments the battery voltage has been drained to 30.0 V.

The batteries power the motor drivers and step-down voltage regulators, which in turn power all other electronics. In addition, the batteries are connected to standalone cell monitoring circuitry, which provides visual feedback on per-cell voltage and issues audible warnings if a cell's voltage drops below a certain critical level of 3.5 V.

⁴<http://www.lipottech.net/newstore/batterie-lipo-4s-148v/772-ltv30-1300mah-148v-volt-xt60.html>

STEP-DOWN VOLTAGE REGULATORS

Two Pololu step-down voltage regulators are used to power low-power electronics, one each of 3.3 V (D24V25F3) and 5 V (D24V25F5), each rated up to a maximum output current of 2.5 A. All electronics, including the two step-down voltage regulators, share the same ground.

Originally the plan was to have only one 3.3 V step-down voltage regulator on Tippy, which would power the Launchpad and Raspberry Pi. All electronics working on 5 V would then be powered by the Launchpad. However, it was found later that the Raspberry Pi works better on 5 V and that the Launchpad is not able to supply sufficient current to power all 5 V electronics. The step-down voltage regulators now power most of the lower-power electronics directly, except for the encoders and IMU, which are still powered by the Launchpad, as can be seen in fig. 5.8.

DRIVERS

Each motor is driven by a G2 24v21 brushed motor driver from Pololu (2016). Drivers are used for power amplification, and do not perform low-level feedback control – in contrast to motor controllers, which typically have built-in position, velocity or acceleration servos. The use of drivers allows for full transparency, as it enables the MCU to directly control the motor through amplification of pulse-width modulation (PWM) signals.

The drivers each can handle up to an input voltage of 40 V and current of 50 A. These drivers are more powerful than required for balancing, but have been selected based on requirements for Skippy. For safety reasons, the maximum current has been limited to approximately 10 A per driver, by physical placement of a resistor of 20 k Ω according to the instructions on Pololu (2016). This current limits the maximum output torque of the joint to 25 N m, for $n = 100$ and $\eta = 0.7$ (like in section 5.4.3), so that the maximum torque remains below the specified peak torque of the strain wave gear of 31 N m.

The drivers also have an inverted sleep pin, which is used to disable each motor if desired, either through (1) limit switches in the bend joint, (2) commands from the MCU or (3) physical toggle switches on Tippy's user interface. This threefold feature of disabling the motor has been designed mainly for safety, and its required circuitry is put on TP-PCB.

TP-PCB

⁵ TP-PCB is the only custom-built electronic board present on Tippy. It contains the following simple circuits:

- A programmable RGB (red-green-blue) LED (light emitting diode) and two programmable ordinary LEDs. They are used for visual status feedback, and protrude from Tippy's body at its user interface.
- Circuitry for two programmable push buttons on the user interface.
- Circuitry for the aforementioned safety features for disabling the motor.
- Various low power outlets and fuses.

⁵ *Acknowledgements* — The electronic circuit design of TP-PCB has been done by the author and has been implemented in electronic CAD software by Marco Migliorini.

USER INTERFACE

Tippy's user interface contains the two aforementioned toggle switches and two programmable push buttons. It also contains a power switch (P-switch), used to prevent any power leakage if the robot is not in use, and an emergency switch (E-switch), used to disable the motors by cutting their power. Both the power and emergency switch are rated for 10A.

ABSOLUTE ENCODERS

All five joints in the robot are equipped with iC-Haus' iC-MU150 18-bit magnetic absolute encoders (iC-Haus, 2018). The encoders themselves are sold as chips, and their boards are an IIT in-house design. The encoders are sampled at a frequency of 5 kHz. At least 3 least-significant bits were found to correspond to noise (effectively rendering the encoder 15-bit or lower). For the actuated joints, these encoders are placed at the output shaft of the actuator (after the strain wave gears).

All five absolute encoders are daisy chained and connected to the MCU through a single serial peripheral interface (SPI) module. Originally it was planned to split the daisy chain between two SPI modules, but the other SPI module was later claimed for connecting the Raspberry Pi Zero. The cables interconnecting the encoders use Molex' Pico Clasp connectors.

RELATIVE ENCODERS

In addition to absolute encoders at the actuated joints, both actuators are equipped with a relative encoder placed directly on the respective motor's output shaft. The motor shaft positions are read by Maxon's ENX 16 12-bit incremental encoder (Maxon Motor, 2017), sampled at a frequency of 20 kHz. Despite having fewer bits than the iC-Haus encoders, they are mounted before the transmissions, effectively increasing their resolution. For the swivel joint, the resolution approximately corresponds to that of a 17-bit encoder at the joint, because $12 + \log_2(30) \approx 17$. The relative encoders are connected to the MCU through dedicated QEP (quadrature encoder pulse) modules.

The absolute iC-Haus and relative ENX encoders have a different purpose.

- Absolute encoder readings are not affected by finite stiffness or backlash from the transmission for measuring moving inertia in the system, for which they are more suitable to be used by the balance controller.
- The relative encoders show less noise than the absolute encoders and can be sampled at higher frequencies. Moreover, since they are directly connected to the motor shaft, they are more suitable for motor servos (programmed on the MCU) that directly control the motor position or velocity.

In addition, the readings from both encoders can be combined for the implementation of a transmission model with finite stiffness, which adds an extra DoF to the corresponding actuator.

RASPBERRY PI ZERO

A Raspberry Pi Zero is used for data logging and highest-level communication. The MCU has limited storage (1 MiB flash, 164 KiB RAM), which is not intended nor sufficient for

storing raw data structures from balancing experiments.⁶ The Raspberry Pi Zero offers a simple solution for data storage as it comes with an SD-card slot. Moreover, it includes a WiFi module, which can be used to remotely start and stop experiments and request logs.

IMU

Tippy is equipped with a Vectornav VN-100 IMU (Vectornav, 2009). It is selected for its low drift and high survivability for high-deceleration impacts, which are important properties for Skippy.

The UART (universal asynchronous receiver-transmitter) interface of the IMU is connected to the MCU through a dedicated SCI (serial communications interface) module. UART/SCI is not as fast as SPI, but it is used nonetheless because (1) the COTS rugged version of the IMU is not available with SPI and (2) both of the SPI modules that are physically present on the MCU's Development Board are already in use (one for the absolute encoders and one for the Raspberry Pi Zero).

5

5.5. SWIVEL EXPERIMENTAL SETUP

This section explains the particular experimental set-up that is used for swivel experiments, both in hardware and software, other than general properties of the robot that are already explained in section 5.4. For this set-up, Tippy is locked according to the description in section 5.4.3. The joint positions of the unactuated roll and actuated swivel joint are redefined to be q_1 and q_2 , corresponding to the models used in section 3.3.2 and chapter 4.⁷

5.5.1. CONTROL LAW SIMPLIFICATIONS

The control theory described in section 3.3 is used for balancing Tippy in the swivel plane. Tippy is a RWP in the swivel plane, which is a special case of the 2R robot described in section 3.3. This allows for further simplification of the balancing theory. In addition, practical implementation allows for further simplifications like applying small-angle approximations.

CONSTANT G_ω AND T_c

As mentioned in section 4.2 the angular velocity gain G_ω of an RWP is constant. It is calculated according to eq. (4.1), and related to the linear velocity gain G_v according to eq. (4.2).

The natural time constant of toppling defined in eq. (3.8) is valid for any configuration, but it is not constant. By altering the definition for T_c to be valid only for the balanced configuration, requiring the substitution $c_y = c$, i.e.

$$T_c = \frac{H_{11}}{mcg}, \quad (5.2)$$

⁶For reference, logging 50 floating-point numbers of 4 bytes each at a sampling rate of 1 kHz requires 0.2 MiBs^{-1} of storage. For an experiment that lasts 1 min this equates to 12 MiB.

⁷In section 5.4.1 the roll and swivel joint were defined as q_3 and q_5 respectively.

it becomes constant for the RWP, like \tilde{T}_c in eq. (4.18).⁸

SUBSTITUTION

The EoM of the RWP are defined in eqs. (4.3–4.5). By substituting these EoM and our definitions of G_ω and T_c in the controller state variables, defined in eqs. (3.3–3.6), we obtain

$$L = H_{11}\dot{q}_1 + H_{12}\dot{q}_2 = mcgT_c^2(\dot{q}_1 - G_\omega\dot{q}_2) \quad (5.3)$$

$$\dot{L} = -mgc_x = mcg\sin(q_1) \quad (5.4)$$

$$\ddot{L} = -mg\dot{c}_x = mcg\cos(q_1)\dot{q}_1 \quad (5.5)$$

$$\ddot{\tilde{L}} = -mg\ddot{c}_x = mcg(\cos(q_1)\ddot{q}_1 - \sin(q_1)\dot{q}_1^2), \quad (5.6)$$

and for the controller gains Y_1 and Y_2 , defined in eq. (3.7), we find

$$Y_1 = \frac{-1}{mcgT_c^2G_\omega} \quad \text{and} \quad Y_2 = \frac{1}{mcgG_\omega\cos(q_1)}. \quad (5.7)$$

Note how the constant mcg appears in all of the above equations. The equations can be shortened by making the substitution $M = L/(mcg)$, which represents a scaled angular momentum variable. In addition, by applying the small-angle approximation to q_1 and by ignoring the term containing \dot{q}_1^2 , the equations can be further simplified. For the particular case of the RWP these simplifications imply that Y_1 and Y_2 become constant, so that the plant becomes linear. The simplifications have been implemented only for reasons of convenience; it should be noted that they are not required in order for the control system to work in practice. The result is

$$M = T_c^2(\dot{q}_1 - G_\omega\dot{q}_2), \quad (5.8)$$

$$\dot{M} = q_1, \quad \ddot{M} = \dot{q}_1, \quad \ddot{\tilde{M}} = \ddot{q}_1, \quad (5.9)$$

$$Y_1 = \frac{-1}{T_c^2G_\omega} \quad \text{and} \quad Y_2 = \frac{1}{G_\omega}. \quad (5.10)$$

By applying these simplifications, the control law in eq. (3.10) becomes

$$\ddot{\tilde{M}} = k_{dd}\ddot{\tilde{M}} + k_d\dot{\tilde{M}} + k_M M + k_q(q_2 - u). \quad (5.11)$$

The simplifications also simplify the balance offset observer (introduced in section 3.3.6), for the fact that \dot{L} is no longer a function of q_2 . By substitution of L for M , eqs. (3.18–3.20) respectively become

$$\dot{M} = \frac{dM}{dt} = \frac{d}{dt}(T_c^2(\dot{q}_1 - G_\omega\dot{q}_2)) \quad (5.12)$$

$$\dot{M} = \hat{q}_1 - q_o \quad (5.13)$$

$$q_o = \text{LPF}\left(\hat{q}_1 - \frac{dM}{dt}\right). \quad (5.14)$$

⁸Note that T_c and \tilde{T}_c are still different from each other. The difference resembles that between a composite rigid-body inertia and an articulated-body inertia.

CONTROLLER OUTPUT

As mentioned in section 3.3.3, the output of the balance controller has to be either a torque or an acceleration command for joint 2. It can be computed from $\ddot{M} = \ddot{q}_1$ and the EoM in eq. (4.3):

$$\ddot{q}_a = (mcg \sin(q_1) - H_{11}\ddot{M})/H_{12} \quad (5.15)$$

$$\tau_a = H_{21}\ddot{M} + H_{22}\ddot{q}_a, \quad (5.16)$$

where τ_a is the command sent to the low-level servo if joint 2 is force-controlled, and \ddot{q}_a if it is acceleration-controlled.

REDUCED NUMBER OF MODEL PARAMETERS

Observe that the RWP needs to know only three model parameters: T_c , G_ω and the value of the product mc . Furthermore, simulation studies have shown that the balancing performance is insensitive to errors in the value of mc , so only T_c and G_ω need to be known with any degree of accuracy.

5

5.5.2. IMPLEMENTATIONS

This section describes various of the decisions and procedures on the implementation and implemented modules to make the balance controller work on real Tippy.

SYSTEM IDENTIFICATION

The parameters that fully define the dynamic model defined by the EoM are $g = 9.805 \text{ ms}^{-2}$, $mc = 0.7471 \text{ kgm}$, $G_\omega = -0.050$ and $T_c = 0.18 \text{ s}$.

- T_c was obtained experimentally and was found to be very close to the value 0.1790 s calculated from CAD data, which can be calculated from body inertias using eq. (5.2).
- The velocity gain G_ω was initially calculated from CAD data using eq. (4.1), equalling $G_\omega = -0.0521$, but it was found experimentally that this value was slightly overestimated. The corresponding behaviour of an overestimated G_ω is reported in Featherstone (2017), and mainly leads to overshoot of q_2 for commanded motion.
- The first moment of mass mc can also be obtained experimentally. However, since (1) simulation results have shown that errors in mc as large as 10% do not significantly affect the performance, (2) measuring it accurately requires a significant amount of time and (3) CAD data had proven to provide reasonably accurate values for other parameters previously, it was decided to assign the value $mc = 0.7471 \text{ kgm}$ to it obtained from CAD (corresponding to that of Tippy including the brace, as described in section 5.4.4).

SENSING

The IMU is not used for the balancing experiments presented below. This implies that the robot has no direct feedback about the absolute vertical ($q_1 = \hat{q}_1 + q_0$), and must

instead rely on the magnetic encoder measuring q_1 and the balance offset observer. The joint angle readings \hat{q}_1 and \hat{q}_2 (used for the balance controller) are measured by the absolute iC-Haus encoders and the motor shaft reading $\hat{q}_m \approx 30\hat{q}_2$ (used for the motor servo) is measured by Maxon's ENX 16 encoder, sampled at 5 kHz and 20 kHz respectively, as described in section 5.4.5.

CONTROL HIERARCHY

The controller consists of a two-layered system: the balance controller and a lower level motor servo for tracking the controller's output. The balance controller, including the balance offset observer, is executed at 1 kHz. The output of the balance controller is the acceleration command \ddot{q}_a of the reaction wheel (eq. (5.15)). This output is integrated to obtain a velocity command, and multiplied by 30 to account for the gear reduction of the strain wave gear. It is then passed over to the motor velocity servo, which runs at a frequency of 5 kHz.

MOTOR VELOCITY SERVO

The motor velocity servo is a proportional controller based on the velocity estimate from the Maxon ENX 16 encoder. The proportional gain equals 0.003. The output is the normalized duty cycle of the PWM output from the microcontroller to the motor driver, which is a number between -1 and +1.

VELOCITY ESTIMATION

Both the balance controller and motor servo require velocity as input. Velocity estimation of both absolute and incremental encoders is done using a simple numerical differentiation combined with a second-order Butterworth filter having a cut-off frequency of 3500 rad/s for the incremental encoder and 500 rad/s for the absolute encoders. The filter was converted from continuous to discrete time using the bilinear (i.e., Tustin) transform.

BALANCE OFFSET OBSERVER

The LPF of the balance offset observer is implemented as a second-order Butterworth filter with a cut-off frequency of 1 rad/s.

INPUT SIGNAL GENERATION

It follows from eq. (3.13) and section 3.3.5 that the input signal is calculated as $u = q_f + \alpha_1 \dot{q}_f + \alpha_2 \ddot{q}_f$. The controller thus requires the modified command signal for leaning in anticipation q_f and its two derivatives \dot{q}_f and \ddot{q}_f .⁹

In theory, q_f , \dot{q}_f and \ddot{q}_f can be calculated analytically from a well defined command signal q_c , its derivatives \dot{q}_c and \ddot{q}_c and a first-order LPF in reverse time. However, in reality, execution of these analytic calculations is not practical for performing experiments with any arbitrary command signal. To prevent loss of practicality, additional calculations have been implemented so that the operator is required to feed only q_c as a piecewise continuously differentiable function in discrete time, with a sampling time of 1 ms.

⁹'Required' here refers to achieving an optimal tracking performance. It is permitted to set \dot{q}_f and \ddot{q}_f to zero for the penalty of reduced tracking performance, but it generally does not lead to the system losing its balance.

First, the acausal filter for leaning in anticipation is implemented by preprocessing the command signal q_c by (1) reversing it in time, (2) passing it through the (stable) filter $1/(1 + T_c s)$ (in discrete time) and (3) reversing it again, to obtain q_f . \dot{q}_f and \ddot{q}_f are then directly approximated by numerical differentiation using a one-dimensional five-point stencil.

Since two derivatives are calculated, q_f – and thus q_c – ideally is a continuously differentiable function, i.e. of continuity class C^1 or higher. If instead q_c is piecewise continuously differentiable, by having points with class C^0 continuity (e.g. a ramp function) or no continuity (e.g. a step function), then corresponding impulses in the derivatives are ignored by being set to zero. For C^0 continuity this is required only for the second derivative, but for non-continuity this is required for both derivatives. For improved performance, it is recommended to avoid at least the latter, and have a continuity of at least C^0 , which can be achieved by replacing steps for steep ramps, following the suggestion from Featherstone (2017).

CLOSED-LOOP SPEED

The choice of the fastest pole λ_4 influences the speed of the response of the closed-loop system, as explained in section 3.3.4. In the experiments reported below, we chose $\lambda_2 = \lambda_3 = \mu_1 = \mu_2 = -1/T_c$ and $\lambda_4 = -10$. In simulation our system could work with a value of $\lambda_4 = -20$. However, in practice such a fast pole would lead to structural vibrations due to the low structural stiffness of the robot's body and the imperfect servo. With these choices, the transfer function from eq. (5.17) becomes

$$q_2(s) = \frac{1}{1 + 0.1s} q_c(s) \quad (5.17)$$

5.6. EXPERIMENTAL RESULTS

This section presents results on swivel position balancing and on finding the maximum offset angle of Tippy, using the experimental set-up described in section 5.5.

5.6.1. MAXIMUM RECOVERY ANGLE EXPERIMENTS

The maximum recovery angle ϕ_{\max} (eq. (4.19)) was found using the same bisection search strategy as was used in the simulation study in section 4.6.

Tippy is first commanded to statically balance with the command signal set to zero ($q_c = q_f = 0$) for the first 10 s. The purpose of this is finding q_o , using the balance offset observer, so that Tippy is truly balanced when $q_1 = 0$. It is then commanded to stop balancing after a few seconds and is caught by a human to prevent it from tipping over. It is then slowly (quasi-statically) lowered from its balanced configuration with a string until it reaches the preprogrammed angle $q_1(0) = \phi$, from which it tries to reach $q_1 = 0$ by commanding the motor to operate in saturation. The experiment is then repeated for different values of $q_1(0)$ until ϕ_{\max} is found within a margin of 0.1° .

Upon finding ϕ_{\max} , the experiments are repeated several times and for both tipping directions. In between experiments pauses of several minutes are introduced to prevent significant heating of the strain wave gear and motor, which could influence experimental results. Prior to starting the experiments, the batteries were drained to match a volt-

age of approximately $U = 31\text{V}$. This corresponds to the experienced average nominal battery voltage for the used 8 LiPo cells.

It was found repetitively that $\phi_{\max} = 14.2^\circ$: Tippy was able to recover from 14.2° but not from 14.3° . This result agrees well with the theory and simulation, for which we found $\phi_{\max|I} = 13.9^\circ$.¹⁰ Encoder readings and discrepancies between real and estimated values of physical parameters could explain the difference between measured and simulated performance. Physical parameters with inexact estimates plausibly include those describing actuator properties and input voltage. The least accurate physical parameter is likely the transmission efficiency $\eta = 0.7$, as the theory assumes it is constant, whereas that of the strain wave gear is in reality a function of temperature, torque and velocity. It was estimated to have a constant value $\eta = 0.7$ from look-up tables of the datasheet (Harmonic Drive AG, 2014) for initial conditions of high torque, low velocity and room temperature (ca. 20°C). It is possible that the temperature in the strain wave gear was higher, which would have increased the efficiency. The maximum offset angle would be calculated as $\phi_{\max|I} = 14.2^\circ$ for $\eta = 0.73$.

5.6.2. BALANCING EXPERIMENTS

This section presents experimental results aimed at evaluating the tracking performance of the control system, and compares them with the theoretical result and previous work.

The experiment begins when the balance controller is activated. The command signal is set to zero for the first 10s; then follows the motion command described below; and then the command signal remains at zero until the robot is switched off.




Figure 5.9: Tippy's command profiles, preprocessed with the acausal filter and numerically differentiated using a 5-point stencil. Impulses that would result from differentiating non-continuous intervals do not appear in any of the command profile's derivatives.

¹⁰The value $\phi_{\max|I} = 13.9^\circ$ calculated for Tippy including the brace and $U = 31\text{V}$ corresponds to that of Tippy excluding the brace and $U = 30\text{V}$, as assumed in chapter 4 (table 4.1).




Figure 5.10: Tippy accurately follows a motion command while remaining balanced.

5

MOTION COMMAND TRACKING

A predefined motion command signal q_c with a duration of $t_{\text{end}} = 10$ s is used to perform the balancing experiments. The command signal is taken from Featherstone (2017), with the exception that steps are replaced with steep ramps, following the explanation in section 5.5.2 and suggestion from Featherstone (2017). The absence of steps allows for faster tracking and elimination of overshoot when defining two finite zeros μ_1 and μ_2 in the control system. The ramps that replace the steps have a duration of 0.25 s. The command signal now consists of various ramps and sinusoids, as displayed in fig. 5.9. The amplitude of the signal has been scaled by a factor of $3/4\pi$ rad relative to the command signal in figure 5 of Featherstone (2017) because the inertia of Tippy’s reaction wheel is lower than that of the hypothetical robot from Featherstone (2017). This scaling allows for leaning in anticipation to be observed more clearly.

As described in section 5.5.2, leaning in anticipation has been implemented by preprocessing the command signal for the full duration with the acausal filter; i.e. we have $q_t(t_{\text{end}} - t) = \text{LPF}(q_c(t_{\text{end}} - t))$, where LPF is a first-order low pass filter with cut-off frequency $1/T_c$. The preprocessed command signal and its derivatives are plotted in fig. 5.9.

The tracking of the command signal is displayed in fig. 5.10, proceeding from $t = 10$ s to $t = 20$ s. The theoretical motion response q_t corresponds to passing q_c through a first order low-pass filter with transfer function $1/(1 + 0.1s)$, which is also plotted in fig. 5.10.

It can be observed that q_2 follows q_t very closely, which demonstrates that the performance of the balance controller in practice is very close to its performance in theory. q_t , in turn, follows q_c with a delay of 0.1 s, which is a consequence of placing the fast pole at -10 . In Featherstone (2017), the fast pole was set at -20 ; but the robot used in those simulations was rigid, whereas Tippy’s mechanism has significant flexibility that limits the maximum bandwidth that can be achieved. Nevertheless, -10 is still a fast pole for a robot of the size of Tippy.

q_2 even seems to lead q_t during parts of the trajectory, which is likely due to the

mechanism's flexibility and imperfections of the velocity observers and motor servo.

Figure 5.10 also plots q_1 , scaled up by a factor of 10. The effect of leaning in anticipation can be observed in the change of q_1 preceding motion trajectories. It can also be observed that leaning in anticipation leads to elimination of non-minimum phase behaviour, as q_2 does not move in the opposite direction of the commanded motion.

The duty cycle of the motor reached up to 95% while tracking the commanded motion. Faster poles, a more aggressive motion command or omission of leaning in anticipation are likely to lead to saturation (where the duty cycle equals 100%), and will significantly reduce tracking performance. Note however that the selected (brushed) DC motor is relatively weak for the robot's size (2.1% of the total mass including the brace).




Figure 5.11: This graph shows two different phenomena. First, for $0 < t < 4$ s, it can be observed that the balance offset observer successfully corrects for the balance offset error. Second, for $6 < t < 10$ s, hunting can be observed when the robot attempts to maintain statical balance, as caused by non-modelled stiction in the strain wave gear.

BALANCE OFFSET CORRECTION

In the experiment, the motion command is preceded by 10 s of balancing at $q_c = 0$ in order for the balance offset observer to converge to the correct value of q_o . Experimental results from these first ten seconds are displayed in fig. 5.11, from which the effect of the offset observer can be well observed during the first few seconds. It shows convergence of q_1 to zero with a settling time of approximately 3 s, corresponding well to the selected cut-off frequency of 1 rad s^{-1} of the observer (for which a settling time within 5% corresponds to 3 s).

The initial movement that is visible in fig. 5.11 is due to corrective actions of the balance controller to the initial errors in the values of q_1 and q_2 . These errors are caused by a human who holds the robot approximately upright with non-zero q_2 , prior to starting the experiment. Corrections from the balance offset observer do not significantly affect the actions of the controller, because the controller acts much faster.

A properly working balance offset observer furthermore suggests that q_2 should converge to its commanded value $q_c = 0$, which it almost does. The reason it does not

completely is due to stiction in the strain wave gear, which leads to hunting, explained in the subsequent subsection.

HUNTING

While the robot is supposed to be stationary ($q_c = \text{const}$) for a prolonged time, it enters a small limit cycle. This is caused by non-linear friction (stiction) in the transmission mechanism and our simplified servo, which fails to model this friction, and subsequently leads to *hunting* around the commanded position: very slow, small movements are hard to achieve, so the joint tends to under-respond until there is enough of a balance error to need a bigger corrective motion. It is believed that optimisation of the servo will reduce this phenomenon.

In the presented experiment, hunting can be observed from $t = 6\text{s}$ to $t = 10\text{s}$, as can be seen in fig. 5.11, during which it is commanded to $q_c = 0$. The observed hunting frequency approximately equals 1 Hz and the amplitude of q_2 does not exceed 2° . However, hunting likely does not significantly affect the balancing performance as it does not influence commanded non-zero velocities.

5.7. REFLECTION

Tippy has been designed to be a good balancer. Its high *physical* balancing capability has been confirmed by the experiments conducted in section 5.6.1. The balance experiments have furthermore proven that the balance control theory can be successfully applied in reality, and that its performance closely matches that of the theoretical one, as can be seen in fig. 5.10. However, there are a few indications that clearly limit Tippy's performance, such as caused by selection of the fastest pole at $\lambda_4 = -10$ instead of a higher in magnitude value. In addition, most of the planned experiments for Tippy (i.e. those listed in section 5.2) have not yet been executed. The main reason for this are delays, but also because priority of the project had shifted to the design of Skippy, which is the core focus of this thesis project. The results that have been obtained from Tippy are sufficient to proceed to the design of a first version of Skippy, but the proposed experiments in section 5.2 will still have to be carried out to (1) fully demonstrate the capabilities of the newly proposed controller by Featherstone (2017) in 3D and (2) design a future version of Skippy that is the result of a multi-objective optimisation for both balancing and jumping in 3D.

A reflection can be made on the causes of delays and on Tippy's weakest links, which have limited its performance.

DEFECT ELECTRONICS

Most of the delays were introduced by electronics, in particular by the daisy chain of iC-Haus encoders. A combination of (1) an unshielded encoder board design that uses fragile connectors and provides no status feedback, and (2) a long daisy chain, have lead to many and long physical debugging sessions. The connectors used are MOLEX Pico-Clasp™¹¹ connectors, which regularly broke or had (temporal) loose contacts, likely caused by cable bending and shocks. As a result, it is a tedious task to track down the

¹¹<https://www.molex.com/molex/products/family?key=picoclasp>

defect in the daisy chain of the five encoders. Moreover, the boards do not have reverse-voltage protection, which once lead to a few boards being destroyed due to a short circuit, causing a few months of delays. In some other cases, the encoders issued random readings that would disturb the balance controller. Remarkably, these random readings were also experienced by other robotic teams within IIT, and no certain cause was ever found.

Most other issues related to electronics in Tippy also stem from cabling. The TI's Launchpad is a development kit, intended to be a prototype board and so it only has pin headers. These connectors are large and protrude from Tippy's torso, introducing a feeble spot, and are furthermore not meant to be used in a dynamic environment, resulting in them easily coming loose.

LOOSENING PARTS

Shock and vibration have in some cases lead to loosening of parts, though in most cases this was quickly solved by applying threadlocking adhesives. In one particular case, the swivel motor shaft loosened, as caused by a loose grub screw of the strain wave gear's mounting hub, despite having been fixed with a threadlocking adhesive. This was not discovered initially as the motor shaft had a flat that prevented free rotation, yet introduced backlash, subsequently deteriorating the balancing performance. Shaft loosening is a known problem with mounting hubs and grub screws. However, since the strain wave gear is not available with other shaft mounts, the solution here is either or a combination of (1) using round motor shafts and glue (2) using a larger strain wave gear that can handle higher peak input torques (3) do not use strain wave gears.

STRAIN WAVE GEARS

Other than the loosened shaft, the strain wave gear itself has been proven difficult to model. A few attempts have been done to implement a more realistic model of its behaviour to improve the motor servo, focussing in particular on zero crossing behaviour. Despite the device showing practically no backlash, it does show significant effects of stiction. The stiction seems—just like efficiency—to be a function of other variables, such as temperature ([Harmonic Drive AG, 2014](#)).

None of the attempts were successful and as such it was decided to settle on the proportional motor servo described in section 5.5, subsequently leading to the aforementioned hunting.

Other difficulties in modelling the strain wave gear are caused by its finite stiffness and dependencies of temperature, velocity and torque on the efficiency.

STRUCTURAL VIBRATIONS

The swivel balance performance is partly limited by the choice of the fastest pole. Increasing the pole to a higher (in magnitude) value than -10 appeared to render the system unstable, and strong indications have shown that finite structural stiffness of the machine are the main cause, as physical placement of an additional structural beam to increase the stiffness has allowed for increasing the pole.

TIPPY'S MASS

Other than learning from causes of delays and performance change, a reflection can be made on Tippy's mass, which—excluding the base joint and brace—totals 2.76 kg. This figure provides comparative material for estimating Skippy's mass, which was initially budgeted to be 2 kg, as mentioned in section 3.1. Tippy is not optimised for mass, and its customizability makes it heavier, yet it is slightly smaller, it does not contain a foot, protective bumpers, and a linkage mechanism to drive the bend joint. Taking these considerations into account, it is expected that Skippy will have approximately the same mass, implying that Skippy's initial mass figure of 2 kg will likely not be realisable.

5.8. CONCLUSION

Tippy was introduced for the purpose of analysing balance performance on a real system and as a design exercise for Skippy. Its high balancing performance has been confirmed experimentally and it has demonstrated a recently proposed new method of balance control that aims to achieve accurate tracking of large and fast motion commands while maintaining the robot's balance, in 2D. The experimental results show that the new balance controller's performance closely resembles that of the theoretical performance, which was furthermore experienced to be limited by structural vibrations and imperfect drive train characteristics. Although the experiments presented in this chapter are limited to the application on a RWP in 2D, the robot and controller are also operable in and extensible to 3D. Causes of delays include defect electronics and loosening of parts. The causes of delay, physical performance limits and mass figures are taken into consideration for Skippy's design.

6

SKIPPY MECHANICAL DESIGN

This chapter describes Skippy's mechanical design in detail, which is the result of considerations for weight reduction, physical robustness and manufacturability. Several design revisions and decisions have been influenced by reflecting on Tippy and early simulation studies. Various parts of Skippy are explained in detail, including its hip, the actuators, the main spring and the ankle with its spring. The presented design leaves a few parameters for optimisation. This final step of iterative design is presented in chapter 8. More generally, the design has resulted in the development of various theories that are more widely applicable than to Skippy alone, including (1) a design strategy for physically robust design, (2) an approach for designing stiff, lightweight and manufacturable structural components, (3) a generalization of the center of percussion, for placement of shock-sensitive components such as IMUs and (4) a model to describe the behaviour and failure of compression leaf springs, the latter of which is treated in chapter 7.

6.1. DESIGN OVERVIEW

The establishment of Skippy's original design is explained in chapter 3. The design has undergone several revisions and has been worked out in greater detail. The strategy towards completing the design has largely been influenced by the proposed philosophy in chapter 2, which aims to proceed through trial-and-error and to settle for a design that is 'good enough'. For this purpose, most design decisions are made through educated-risk considerations and analyses rather than by performing complex simulations and global optimisations, generating Pareto fronts or constructing morphological overviews. For the same reason, earlier design decisions have not been overhauled unless there were sufficient indications for the necessity of doing so.

Several of the design decisions and revisions were influenced by reflecting on Tippy, such as its high mass, low structural stiffness and feeble electronic connections, and others by reflecting on results of previous simulation studies, such as those in [Driessen \(2015\)](#). The revisions mainly involve a change of the main motor and moving the batteries from the crossbar to the torso. Most design novelties and decisions are a product

of the challenging requirements for being lightweight, affordable and physically robust. Figure 6.1 provides an overview of the design's aspects of interest, each of which will be treated in this chapter. Roughly, the chapter can be divided in four parts:

- Section 6.2 discusses design revisions of Skippy's actuation and battery placement.
- Sections 6.3 and 6.4 explain the design strategies and decisions towards making the robot lightweight, affordable and physically robust. In particular, section 6.3 focuses on physical robustness, and section 6.4 on the design style employed for structural components to realize a lightweight and stiff but manufacturable design.
- Sections 6.5 to 6.9 explain in more detail the mechanical design of Skippy's specific parts, namely the hip (including the four-bar linkage), the ankle and foot (including the ankle spring), the main actuator (including the main spring), the swivel actuator, and the head, respectively.
- Section 6.10 is a more theoretical section on IMU placement, and introduces a generalized definition of the *Center of Percussion* and a method to calculate it.

The main focus in most of these sections is on practical realisation of the mechanical design, and involves Skippy's detailed creation in CAD software (Siemens NX 11). Most of the related design decisions have been finalized in this chapter (or had already been finalized in a previous stage of Skippy's iterative design process). However, some decisions on physical parameters, such as those of some link lengths and spring strength are finalized in chapter 8 through optimisation studies and a subsequent (final) design iteration. Furthermore, some parameter choices are explained in more detail in chapter 8, as that chapter has a stronger focus on modelling and computation, whereas this chapter focuses more on CAD.

The detailed CAD data of the nearly finished design that is presented in this chapter provides more realistic values of Skippy's dynamic properties (mass and its distribution) and of physical constraints, which are required for performing realistic optimisation studies. By physical constraints are meant limits (minima and maxima) of for example link lengths and a joint's work space, to prevent internal collisions of bodies.

6.2. DESIGN REVISIONS

The most prominent design revisions that have been done since completion of Tippy experiments reported in chapter 5 are those regarding battery placement and of both main and swivel actuator selection, which are discussed in this section in more detail. Battery placement is revised upon discovery that the implementation of its originally intended placement in the crossbar is not as feasible nor rewarding as initially conceived. Revisions regarding actuation have various causes. For the main actuator, the main reasons are that of an increase of the expected robot mass and limited maximum actuation velocity. The main reasons for revising the swivel actuation is dissatisfaction of the strain wave gear transmission mechanism and the necessity for placing dead weights in the crossbar (as a consequence of moving the batteries out of the crossbar). Skippy's increased anticipated mass also influences other design decisions.




Figure 6.1: Right side view of the Skippy's design. Various design points of interest are indicated.

INCREASED MASS

As mentioned in section 5.7, Tippy's mass of 2.8 kg has raised awareness about Skippy's original target mass of 2 kg (see section 3.1.2) being too tight. Skippy is expected to have a similar mass, as it is anticipated that extra mass from new features (e.g. four-bar linkage, springs, protective bumpers, foot) and reduced mass from better design approximately cancel each other out. Skippy is now also more comparable to Tippy regarding battery placement, since both (will) have the batteries placed directly on the torso rather than on the crossbar, as explained further below in section 6.2.1. Skippy's new anticipated mass of up to 3 kg is combined with a reduced target jump height of at least 3 m (instead of 4 m). This still requires a more powerful main actuator, but makes the problem more tractable. It also influences other choices in the robot design, such as those related to protective bumpers, and in optimisation studies.

FURTHER REVISIONS

Other revisions include the addition of an end stop in the hip to act as parallel spring and those related to electronics. The hip end stop has been added to boost Skippy's jump height, as a response to the anticipated increase of its mass. It is further explained in section 6.5. Revisions regarding electronics are a response to encountered issues in Tippy, which mainly concern cabling and connectivity, but do not concern many changes in the actual circuitry. This is further discussed in section 6.9.

6.2.1. BATTERY PLACEMENT

As indicated by eq. (4.25), the maximum balance recovery angle is proportional to the square root of the crossbar's inertia (assuming an optimal transmission ratio). As such, the crossbar must have a significant mass if the balance performance is not to be sacrificed. With the intention of lowering the total mass, the original plan was to place Skippy's batteries in the crossbar, as depicted in fig. 3.1a, which would prevent the necessity of using *dead weights* (objects that have no other purpose than having mass).

Placing the batteries in the crossbar requires the use of a slip ring to transfer power from the rotating crossbar to the torso. The slip ring has to be able to transfer up to ca. 60 A peak current. Most of the industrial COTS slip rings with this rating have a mass of several hundred grams, and thus offer no feasible solution. Their axial size furthermore pushes the swivel motor further back from the crossbar, which is unfavourable for the mass distribution. A possible solution is the use of a pancake-style slip ring, which consists of two spring-loaded PCBs with cantilever springs as sliding-contact points, of which a custom design for our current rating is depicted in fig. 6.2. The slip ring takes up little axial space and is lighter and simpler than conventional industrial slip rings. Nevertheless, the added mass of the slip ring in combination with that of extra high-current cable length and protection to shield the batteries from crashes (the tips of the crossbar are a vulnerable location) makes the plan less attractive than initially anticipated. In addition, the pancake-style slip ring's big diameter (compare to the pen in fig. 6.2) and its sliding contact principle introduce significant friction, of which mainly stiction, which had previously shown to negatively affect Tippy's balancing performance. These disadvantages, and possible yet undiscovered problems (e.g. those related to noisy power lines), were sufficient reasons to step away from the original plan and to redesign for the batteries to be placed directly in the torso.

6




Figure 6.2: The pancake-style slipring that would have been mounted on Skippy if batteries were to be placed in the crossbar.

6.2.2. MAIN ACTUATOR

The studies conducted in [Driessen \(2015\)](#) assume a a Maxon RE35 24 V motor, which has a maximum mechanical power output of 383 W with an assumed supply voltage of 30 V. This limit was reached during most of the thrust phase for the (lighter) 2 kg robot. A more powerful motor is required for a heavier robot, as well as a transmission with a higher power throughput in order to use the motor to its fullest potential. In addition, the same studies have shown that the ball screw speed limit further limits the ability to steer at lift-off, and as such it is desired to increase both the speed and power limit.

Possible considerations are reducing the robot overall size, or replacing the motor in combination with a different transmission mechanism or a ball screw pitch. Each of the considerations are further explained below. The latter of these solutions has been applied for the presented design, by replacing the RE35 motor for the more powerful DCX32L in combination with a ball screw pitch of 4 mm instead of 2 mm.

CHANGING ROBOT SIZE

Reducing the robot overall size and thus mass eliminates the requirement for extra power. However, one would quickly find that downscaling would severely impede the jumping performance if speed is already a limiting factor. As explained by appendix A, speed limit scaling is governed by a scaling factor of S^2 , which would imply that Skippy half the original size would be able to jump approximately a quarter of the height. In addition, we find that the availability and performance of technology that is chosen for Skippy further shrinks for reduced size. For example, we find that high precision ball screws with a nominal diameter below 8 mm have limited availability. We also find that smaller motors generally have a lower power density and efficiency than basic scaling laws (as assumed in appendix A) would suggest. In summary, a smaller version of Skippy would require the use of different technologies. Since this would require major preliminary design revisions, it was decided not to go this route.

CHANGING TRANSMISSION MECHANISM

The studies conducted in [Driessen \(2015\)](#) assume a Steinmeyer ball screw with a pitch of 2 mm, with a specified maximum thrust force of approximately 2 kN and an assumed maximum rotational velocity of 6000 rpm. These limits limit the power throughput to $6000/60 \times 0.002 = 400$ W, similar to the motor's mechanical power output of 383 W. Replacing the ball screw for an other type of rotation-to-linear transmission mechanism that has higher limits on the speed (and thrust force) and thus power may offer a solution in combination with the use of a more powerful motor.

Roller screw The roller screw offers significantly higher thrust force and slightly higher speed than a ball screw. However, it is also substantially more expensive (ca. 2000 instead of 300 Euro), it is slightly less efficient (ca. 80% instead of 90%) and it has a bulkier nut. In addition, at present, a roller screw of the required (small) size is machined per request and can thus not be classified as a COTS component. Its high cost is the main discouraging factor, and it does not scale for bigger purchases, for which it does not offer a long-term solution.

Ring screw A more recent invention is the ring screw ([Featherstone, 2016a](#)), which offers significantly higher speed than the ball screw and approximately the same thrust force, for which it could be paired well with for example Maxon's faster and more powerful EC-4 pole brushless motors. The invention of the ring screw has in fact been a direct consequence of the desire for a higher speed limit in Skippy. It has been tested up to speeds of 16krpm ([Heijmink, 2018](#)), which is several factors higher than that of the ball screw. However, it is still in development. One of the remaining challenges is shrinking the size and weight of the nut whilst ensuring the right pre-load. As the ring screw is still in development, it is not a preferred or viable solution at this stage.

BALL SCREW PITCH

A higher ball screw pitch allows for more power throughput and an increased maximum linear velocity. The maximum permissible rotational speed and thrust force do not change significantly for a different pitch, which implies that a ball screw with double the pitch allows for approximately double the power throughput and double the linear speed. In addition, it was discovered that the Steinmeyer ballscrew with a pitch of 2 mm was discontinued, which further triggered looking in this direction. Nevertheless, the jumping performance is largely determined by the total transmission ratio of the hip, so that doubling the pitch requires doubling the lever length (14 cm instead of 7 cm as used in [Driessen \(2015\)](#)), effectively negating the benefit of the improved linear speed limit of a ball screw with a higher pitch. In addition, the longer lever would make the robot bulkier and more vulnerable.

Changing the ball screw pitch can offer a solution in combination with replacing the motor for one that is more powerful and that is able to deliver more torque. A motor that is able to deliver double the torque would require only half the transmission ratio, whose implementation (e.g. by halving the lever length or doubling the ball screw pitch) would double the speed limit.

The ability of a motor to deliver torque can be measured by the motor constant $K_m = K_\tau / \sqrt{R}$ or stall torque $\tau_s = 2\tau_{\text{opt}} \propto K_\tau / R$, in which K_τ and R are the motor's torque constant and winding resistance respectively. These measures are approximately proportional to the torque a motor is able to deliver under the assumption of either identical Joule heating or identical voltage respectively. The difference in use of the motor constant and stall torque as measure typically resembles that between continuous and peak performance applications. Like stall torque, a motor's peak power output P_{max} is also inversely proportional to a motor's winding resistance, and can be approximated by $P_{\text{max}} \approx 1/4 U^2 / R$, with U the supply voltage.

From available Maxon motors that are compatible with electronics used in Tippy (i.e. brushed, a voltage supply of 30 V and a maximum current of 50 A as limited by the batteries and drivers), the Maxon DCX32L 24 V motor is found to be most suitable. At a mass of 325 g, it is ca. 5% lighter than the RE35 of 340 g. Nevertheless, it is able to deliver substantially more peak power and torque for an identical voltage. Its torque constant and winding resistance are a factor of 1.76 and 1.07 lower respectively, and correspondingly, the motor constant and stall torque are a factor of 1.24 and 1.65 higher. These numbers suggest improved peak torque and power output for the cost of increased joule heating. For a supply voltage of 30 V, the peak power output and corresponding properties are compared in table 6.1. P_{max} is the peak power output, and τ_{opt} , ω_{opt} , I_{opt} are the corresponding motor torque, velocity and current respectively. The peak power output – and therefore also the Joule heating¹ – are approximately 77% higher. The current for maximum peak power has increased from 25.8 A to 45.4 A, which is still within the 50 A limit of the 24v21 Pololu drivers, as used for Tippy (see section 5.4). Note furthermore how ω_{opt} is within the ball screw speed limit of 6000 rpm.

¹Joule heating $P_{J,\text{opt}}$ for peak power output is approximately equal to the peak power output, i.e. $P_{J,\text{opt}} \approx P_{\text{max}}$, corresponding to an efficiency of 50%

	P_{\max} [W]	τ_{opt} [Nm]	ω_{opt} [rpm]	I_{opt} [A]
RE35 24 V	383	0.75	4889	25.8
DCX32L 24 V	677	1.23	5237	45.4

Table 6.1: Comparing Skippy's old (RE35) and new (DCX32L) main motor for maximum power output at 30 V.

Heat Generation With the required speed and current within bounds, the major drawback of the new motor is an increase of heat generation. This is problematic for achieving a lengthy sequence of making high hops. However, it is not Skippy's requirement to do so. The following ballpark figures show that Skippy could spend at least a quarter of its time making its highest hops.

The heat that is dissipated from the motor's windings to the ambient air can be estimated by $P_{\text{wa}} = \frac{T_{\text{w}} - T_{\text{a}}}{R_{\text{wa}}}$ with T_{w} and T_{a} the winding and ambient temperature respectively, and R_{wa} the winding-ambient thermal resistance. Maxon typically provides the winding-housing thermal resistance R_{wh} and housing-ambient thermal resistance R_{ha} , so that $R_{\text{wa}} = R_{\text{wh}} + R_{\text{ha}}$, the latter of which can be reduced by 80% by coupling the motor to a heat-conducting (e.g. metallic) housing (Maxon Motor, 2017). This reduction is assumed in calculations for Skippy, and is justified by (1) coupling the main motor on an aluminium mounting (the rocker) and (2) actively cooling it with a fan and heat sink. We have $R_{\text{wa}} = 2.3 + 7.28 \times 0.2 = 3.76 \text{ KW}^{-1}$ for the DCX32L motor. With a maximum permissible winding temperature of $T_{\text{w}} = 155^\circ\text{C}$ and an ambient temperature of $T_{\text{a}} = 30^\circ\text{C}$, we find $P_{\text{wa}} = 33.3 \text{ W}$. The effective duty cycle of Skippy's motor is conservatively estimated to be 20% of maximum power output for preparing and making its highest hop,² which corresponds to an average power consumption of $0.2 \times 667 = 135 \text{ W}$. These numbers suggest that Skippy can spend approximately a quarter of its time ($33.3/135$) making its highest hops. By applying the same calculations to the RE35 motor, we find that Skippy would be able to spend approximately half of its time making its highest hops.

The maximum duration of each high hop session is furthermore limited by the thermal time constants. As further presented in chapter 8, new simulation studies of the stance phase include thermal calculations, which take both thermal resistances and time constants into account, and have shown that during a single stance phase the coils do not heat up more than 6°C .

Anticipated lever length The motor is paired with a 4 mm pitch ball screw from SBN, which is able to transfer nearly double the power that the motor is able to deliver, compared to the 2 mm pitch ball screw. The studies in Driessen (2015) suggested a lever length of 7 cm to 8 cm for a compromise between jumping and steering performance, where a shorter lever (more speed) favours steering and a longer lever (more torque and power) favours high jumping. Based on these figures and the new pitch (doubled) and stall torque ($1.65 \times$ higher), the new lever is expected to be maximally $8 \times 2/1.65 \approx 10 \text{ cm}$ long. The final lever length is decided based on optimisations in chapter 8.

²Note that Skippy spends approximately 10% of its time in stance while making its highest hop according to the figures in section 3.1, but some energy will also be spent on in-flight control and balancing.

FUTURE CONSIDERATIONS: BRUSHLESS MOTORS

It has been considered to upgrade the main motor to a brushless motor, as these are available with a better power-to-weight ratio, especially if purchased as a component set (only the rotor and stator), so as to eliminate the need for a coupling by designing the screw spindle to also be the motor shaft, and by integrating the motor's housing with the rocker frame and cooling unit. However, this solution is more expensive and more time consuming. To save costs and time it was decided to refrain from this idea and instead save it as consideration for a future version of Skippy.

6.2.3. SWIVEL ACTUATOR

The main reasons for revising the swivel actuation is dissatisfaction of the strain wave gear transmission mechanism and the introduced necessity for placing dead weights in the crossbar. However, no convincing alternative for a strain wave gear was found that makes use of a brushed motor, which is further elucidated below. As such, it was instead decided to stick with the DCX22S motor and the strain wave gear, albeit one size larger and a custom housing to save weight, like in Tippy's bend joint design. The bigger strain wave gear solves issues regarding shaft loosening at the input hub, and furthermore offers more flexibility for designing a slip clutch because of the higher permissible peak torque. The corresponding component set is Harmonic Drive's HFUC-11-2A, which offers a peak torque of 8.5 Nm instead of 3.3 Nm. Slip clutch design is further treated in section 6.3.2.

6

FUTURE CONSIDERATIONS

Referring to equations in chapter 4, the amount of dead weight in the crossbar could be reduced by choosing a slightly more powerful motor and a lower gear ratio, in order to keep the same balancing performance. It could be considered to replace the DCX22S for another brushed motor, but this has other drawbacks, as illustrated below by the exemplary case study of using the DCX22L motor.

The motor power P_m (eq. (4.8)) nearly doubles for using the DCX22L instead of the DCX22S motor. By referring to eq. (4.24), it follows then that only half the crossbar inertia I_2 is required to obtain the same balancing performance (note how both $I_2^{1/2}$ and $P_m^{1/2}$ appear in the nominator), assuming a corresponding optimal transmission ratio n_{opt} . For dynamic parameters identical to those of Tippy, n_{opt} can be calculated as $n_{opt} \approx 14$ using eq. (4.23). Since the crossbar weighs approximately 0.3 kg, this allows for a reduction of crossbar weight of up to 0.15 kg, whereas the DCX22L motor only weighs 0.03 kg more than the DCX22S motor, consequently leading to an overall weight reduction of up to 0.12 kg.

However, there are two disadvantageous consequences. Firstly, due to the reduced transmission ratio, the no-load speed of the crossbar has nearly doubled, which makes it more dangerous. Secondly, the transmission ratio of $n_{opt} \approx 14$ falls in an unfortunate range of transmission ratios: the lowest strain wave gear transmission ratio available is 1:30, whereas the highest gear reduction ratio of feasible alternatives is approximately 1:7. These alternatives are single-stage planetary gearboxes, capstan drives and belt drives (multi-stage drives are avoided for reasons of increased backlash, loss of transparency or bulky and complicated design). The use of such a drive would require a

motor with an even higher motor constant, which was found to be only achievable by switching to flat brushless outrunners. As discussed in section 6.2.2, it was decided not to make the switch to brushless motors due to lack of time, and that this option is instead kept as a consideration for a future version of Skippy.

6.3. PHYSICAL ROBUSTNESS

In line with the design philosophy, Skippy is to be physically robust to survive its crashes from falls. This is challenged by the fact that Skippy will have to survive crashes also from its highest jump height of at least 3 m high. As the result of a crash, all the parts in the robot are subject to a loading, which is the direct result of their inertia (deceleration) and connectivity (internal load propagation), and a crash is survived if none of the loadings exceed their respective permitted values that would lead to a part's failure. A strategy for improving robustness can therefore be established by focussing on the following three core targets.

- Directly reducing *peak decelerations*
- Reducing resulting internal *peak loads*.
- Increasing *permitted loadings*.

Each of the three targets are tackled subsequently.

6.3.1. REDUCING PEAK DECELERATIONS: BUMPERS

Peak accelerations are reduced by externally protecting the robot with bumpers. The bumpers are designed to reduce impact deceleration to a specific minimum, which inevitably requires a certain minimal bumper thickness and energy capacity. In addition, attention is paid to making the bumpers weight efficient, which can be accomplished as follows:

- The main energy absorbing or dissipating material of a bumper should have a high gravimetric energy density.
- As much as possible of the material should undergo deformation during a crash.
- The energy capacity of a bumper should not be more than necessary.

HIGH GRAVIMETRIC ENERGY DENSITY

A material with a high gravimetric energy density is able to store or dissipate a lot of energy per unit of mass. Materials with this property include some types of polymers such as rubbers and glass-fibre reinforced plastic (GFRP). These are also the materials used for Skippy's springs, as further explained in section 6.6.1 (ankle spring, GFRP) and section 6.7 (main spring, rubber). However, both of the materials typically have a density that is too high to be used as a solid bumper. As for the use of GFRP, this has led to the design of (hollow) polyhedral structures constructed from GFRP rods, as observed in various quadcopter bumpers. This solution is feasible for robots of simple shapes that consist of a single body, but is less feasible for multi-body systems such as Skippy due to difficulty of design and construction. Rubbers and some other types of polymers instead are also available as low density foams and sponges, which can be used directly as solid

bumper. Most foams consist of rigid bubbles that deform permanently, whereas sponges typically retain their original shape after deformation. The latter is of interest for Skippy, since it has to be reusable for multiple crashes.

MAXIMIZING MATERIAL DEFORMATION

In lowering the required bumper mass, it is also important that as much as possible of the sponge is utilized during a crash. This can be achieved by (1) lowering the number of bumpers and by (2) maximizing deformation per bumper.

Skippy is to be operational in a concave environment, such as an empty room; it cannot receive a ballistic impact. This allows for placement of bumpers at only a few tactical places on the robot, which are the robot's protruding parts. The bumpers are located at Skippy's hip, head and swivel axle, as visible in fig. 6.1. Note that the heel and crossbar do not have bumpers. The crossbar is designed to be flexible, as explained in the next section (section 6.3.2). The heel instead does not have a bumper because Skippy will be programmed with a reflex causing the leg to flex when it detects that it is going to crash, so that the heel will hit the end-stop on the head instead.

To maximize utilization of sponge within a single bumper, a sandwich construction is made similar to that observed in conventional (e.g. motor cycle) helmets (Mills & Gilchrist, 1991; Viana, 2006), where the sponge is encapsulated by a stiff shell. In fig. 6.1, this shell is visible on the head's toroidal bumper and the frontal swivel axle bumper. The shell redistributes the load to maximize the sponge's deformation, and is made of carbon-fibre reinforced 3D printed plastic. The sponge torus is mounted on a flexible frame of the same material.

6

BUMPER THICKNESS

A bumper has to dissipate most of the energy if the direction of the impact passes through the robot's CoM. If it is then assumed that (1) the robot is rigid, (2) that the bumper has a linear force-displacement profile and (3) that the bumper deforms over its full thickness, then it follows that the robot's peak deceleration is $a_{\max} = 2gh/t$, and the bumper peak force $F_{\max} = 2Gh/t$, where G is the gravitational force $G = gm$, m the robot's mass, h the jump height and t the thickness of the bumper. The assumption of full compression and a linear force-displacement profile is rather optimistic. In reality the force-displacement profile is likely non-linear and progressive, as partly caused by hyperelastic compression, which implies an increase of the peak force. However, the assumption of the robot being rigid is pessimistic: a significant part of the impact energy will be dissipated due to deformation of the robot itself. Accurately modelling these effects is time-consuming and requires accurate knowledge of all material properties. Instead, it is assumed that these effects approximately cancel each other out and that the linear force-displacement profile provides a reasonable approximation of the shock that the robot endures during a crash. Actual building of the robot and experiments should clarify whether these assumptions are to be revised, according to the iterative trial-and-error design philosophy.

The bumpers on Skippy are designed with a sponge thickness of 3 cm, which for a robot with a mass of 3 kg falling from a height of 3 m implies a peak force of ca. 6 kN under the above assumptions. Bearing selection and finite element analyses (FEAs) of structural parts are based on these figures.

ENERGY CAPACITY

Last, the bumper should be designed to have sufficient but not superfluous energy capacity. The required energy capacity influences the choice of the sponge (considering its volumetric energy density), the bumper's dimensions other than the thickness (e.g. the torus width) and the filling density of the bumper. A sponge that is expected to have an energy density in the right range is Misumi's 'low rebound' sponge³. The sponge has a density of $\rho = 97 \text{ kgm}^{-3}$ and a hardness of Asker C27, with a corresponding modulus of ca. $E = 0.5 \text{ MPa}$. The 'low rebound' type sponge furthermore stiffens at impact (due to its viscosity), further improving its gravimetric energy density.

The toroidal bumper is placed close to the robot's CoM, at which also most of Skippy's mass is located. This implies that the robot's inertial (deceleration) force at a crash passes through or close to the bumper, initiating little rotation, which makes the bumper responsible for dissipating most of the robot's kinetic energy. A rough estimate can be made of the torus' ability to dissipate this energy. As a solid body (i.e. a filling density of 1), the sponge torus as depicted in fig. 6.1 has a volume of $V = 661 \text{ cm}^3$. If we assume a deformation of $\epsilon = 0.5$, and a constant dynamic modulus $E = 1 \text{ MPa}$ (a pessimistic estimate that includes stiffening from viscous effects and hyperelastic compression), then the torus dissipates $1/2\epsilon^2 EV = 83 \text{ J}$, approximately equal to the kinetic energy of Skippy falling from a height of 3 m.

In reality, the sponge is not a single solid body and can for instance consist of sheet segments. The filling density can then be reduced by increasing the gaps between the segments. Thus, note that, even if the above calculations provide only a rough estimate, the sponge's energy capacity, stiffness and weight can be easily reduced by reducing the bumper's filling density if experiments conclude that the bumper is too stiff. The first of these experiments can be done with the catapult, which is introduced in section 6.3.3.

6.3.2. REDUCING INTERNAL PEAK LOADS

Internal peak loads are reduced by (1) permitting motion where it is permissible and by (2) choosing proper fastening techniques, both of which are further explained below.

PERMITTING MOTION

For balancing purposes the robot is desirably as stiff as possible, as discussed in chapter 5. High stiffness is however not desirable for physical robustness. One possible consideration is designing for Skippy to have a low stiffness only in certain directions, such as for out-of-plane motions. However, since Skippy is a 3D robot with a yet undiscovered full 3D behaviour, and since torsional vibrations have proven to be problematic in Tippy, it is decided not to experiment in this direction.

To accommodate both desires for improved performance and physical robustness, some parts can be designed so that they are stiff during normal operation but permit motion by becoming flexible for excessive loads. A classic industrial example of this principle is a slip clutch. A slip clutch has also been implemented in the design of the swivel actuator, as further explained in section 6.8.

³<https://uk.misumi-ec.com/vona2/detail/110300279330/>

Buckling crossbar beams The slip clutch protects the strain wave gear from torsional impacts, but it does not protect the joint against axial and bending impacts. To protect it from bending impacts, each crossbar beam is designed (1) to be sufficiently stiff during normal operation, (2) to become flexible for excessive (impact) loads and (3) to be able to autonomously restore to its original shape after substantial deformation. The latter implies that no plastic deformation is allowed to occur and that there is no local minimum in potential energy storage of the spring other than at its rest position. [Mintchev, de Rivaz, and Floreano \(2017\)](#) presents a frame design of a quadcopter with similar properties. The frame remains stiff during flight and becomes flexible at impact.

An affordable and practical design solution is a self-restoring compliant buckling mechanism consisting of two curved spring steel blades mounted back-to-back, rigidly glued to 3D printed plastic end pieces. The curved spring steel blades are those found in tape measure. A prototype is depicted in [fig. 6.3](#). Two blades are used per beam to provide an equal buckling load in both directions. The blades are sufficiently stiff during normal operation, but buckle at excessive loads. The blades snap back to their original position after buckling, which has been verified experimentally. For safety reasons, the blades could be shielded by a fabric sleeve (not depicted in any of the figures), such as that used for electric cabling or chain locks.

The crossbar's dead weights are black rubber knobs that are connected to the blades via its 3D printed end pieces, as depicted in [fig. 6.1](#)). The knobs are the objects to first make contact with the environment during a crash. A notable property of the blades is that they are not very torsionally stiff. This suggests that the dead weights are ideally axially aligned with the blades to reduce vibrations during use. Nevertheless, this alignment also introduces the possibility for a crash that does not cause any bending moment on the blades (namely that corresponding to an impact that is axially aligned with the blades too), causing them not to buckle instantly (as the Euler buckling load of the blades is substantially higher), resulting in a high radial shock on the swivel shaft. To guarantee a bending moment on the blades for any type of crash, the knobs are placed slightly off-axis. Investigating the influence on the performance of vibrations as caused by the off-axis alignment of the knobs requires substantial modelling and analyses. Instead, by following the design philosophy presented in [chapter 2](#), it is suggested to investigate this through experimental trial-and-error: if the vibrations appear to be of significant trouble, then the knobs can easily be moved closer in line with the blades' axes by reprinting the end pieces that connect the knobs to the blades.



Figure 6.3: A prototype beam of the crossbar, consisting of two tape measure blades and 3D printed end pieces. The design is a self-restoring compliant buckling mechanism that protects Skippy from breaking when falling on the crossbar and enables it to get back up unaided. The blades are not depicted with protection sleeves.

PROPER FASTENING & FEWER PARTS

If a robot physically breaks, it is either due to one or multiple parts breaking in half or due to parts breaking loose from other parts. The latter implies failure of fastening. Many fastening mechanisms are governed by pre-tension and high internal loads (e.g. structural screws), which require a heavy design. In addition, failure of fastening is typically less predictable due to complications of properly modelling contact constraints and different fastening techniques, which results in over-engineered designs. Proper fastening design improves the physical robustness by further reducing peak loads in the robot. The challenge here is to find lightweight solutions.

A simple strategy is lowering the number of fastening mechanisms by using fewer parts. However, this typically also increases individual part's complexity, which introduces a trade-off for manufacturability and assembly (and thus cost). Most of Skippy's structural frames are made of single parts, and special attention is paid to manufacturability by employing a specific design strategy, as further explained in section 6.4.

In addition, most of the robot's parts are fastened with an increased surface area by using bonds and by reducing the number of structural screws. Bonds include both adhesive bonds and pressure bonds (i.e. interference fits). Screws have limited surface area and require pre-loading (their tightening), which increases peak stress and thus requires a heavy design. They are also less reliable than bonds because of possible loosening due to structural vibrations, like experienced in Tippy, although this can largely be solved by using thread-locking adhesives. Examples of bonds in Skippy's design include:

- glueing of carbon fibre reinforced plastic (CFRP) tubes to structural frames;
- clamping of bumper inner (3D printed) frames to structural frames using an interference or snap-fit (visible for the head in fig. 6.19).
- clamping of the battery frame to the head frame (held in place by the addition of only two screws);
- fitting the swivel actuator to the head through a plastic sleeve (held in place by the addition of screws).

Skippy still makes use of threaded fastening mechanisms, but in many cases these consist of a single (big diameter) fine thread rather than a flange with multiple screws, such as found in tactical torches (see fig. 6.4). This solution is lighter than using a flange, and the increased surface area of the big thread improves force distribution. The disadvantage is limited applicability if radial alignment is important and if torques are transferred between parts. In Skippy, single fine thread fastenings can be found in the following parts.

- Most of the connections in the main actuator make use of a single fine thread rather than a flange with screws, including those of the main spring, the nut and its guides (fig. 6.17). Note that some torque is transferred between the nut and the coupler. However, the highest torque is right-handed (as a result from pulling the nut during the thrust phase), which would lead to further tightening of right-handed thread. In addition, thread locking adhesives can be used to prevent loosening.

- The frame of the swivel actuator is connected using a single fine thread (fig. 6.18).
- The rocker axles are fixed with a single fine thread (fig. 6.17a).
- All of the hip axles and ankle axle are fixed with a single screw (fig. 6.13).

Screws in Skippy of a size smaller than M4 are only used for fixation (e.g. of electronics), not for structural fastening.



Figure 6.4: A disassembled tactical torch with two fine thread fastenings at display, designed to be reliable and robust. Many fastenings in Skippy are based on this technique.

6.3.3. INCREASING PERMITTED LOADINGS

Other than reducing peak loads and accelerations we can focus on increasing permitted loads by choosing components with better shock resistance. This includes selection of proper bearings, such as that of the needle bearings in the hip (see fig. 6.13), but also of electronics. As a response to the encountered issues with electronics in Tippy, the catapult project has been established, which serves to test various electronic components for their impact resistance. In addition, many of the small and feeble cable connectors have been replaced for either more rugged connectors or soldered connections, depending on the required need for frequent removal/replacement. The catapult is further explained below.

CATAPULT

The catapult is a mechanical device meant to do shock-testing of various components, such as encoders and batteries, see fig. 6.5. Unlike ordinary catapults, this catapult does not release a ballistic payload. Instead, the payload remains attached to the catapult's arm and smashes it into a bumper. The bumper can be either a sponge or solid rubber, of the same materials used in Skippy's bumpers and end stops. The catapult has a secondary purpose of identifying the energy capacity of the bumpers.

The purpose of the catapult is to mimic crash scenarios of Skippy in a controllable way. A crash is governed not only by peak acceleration and force, but also by an acceleration profile (i.e. a duration), which is influenced by the impact velocity. The required energy to accelerate the catapult's arm to a desired velocity is realized by pre-loading it with passive springs. The catapult is designed so that it can accelerate the arm to an impact velocity of at least 9 m s^{-1} , which corresponds to mimicking a fall from a height of at least 4 m.

Principle The potential energy of the springs is transferred to kinetic energy as the springs pull on a pulley that is rigidly connected to the catapult arm, on which the payload is mounted. The catapult is pre-loaded by a human by rotating the arm with an angle $0 < \alpha_0 < 270^\circ$ from the bumper. The bumper can be placed either close to the joint




Figure 6.5: The Catapult: a mechanical device designed and built to mimic impact conditions of Skippy in a controlled way, for the purpose of testing various of its components for their sensitivity to shock.

(higher force, lower velocity) or far from the joint, at the payload (lower force, higher velocity), as would be desired for using either a rubber end stop or sponge bumper respectively. The impact velocity can be controlled by changing the pre-load angle, or by changing the number of springs. The catapult arm, bumper and springs are all mounted on the same (rigid) frame, which is a workbench.

Experiments⁴ The catapult was at first not complete with the pulley encoder depicted in fig. 6.5. Nevertheless, initial experiments have been carried out to find a suitable encoder for Skippy, based on reasonably accurate knowledge of system inertia and spring properties. The tested encoders include CUI AMT20 (absolute) and AMT31 (relative) encoders, Netzer DS-25-16 absolute encoders and RLS (Renishaw) AksIM 2nd generation $\varnothing 29$ and $\varnothing 39$ absolute encoders. The encoders have been mounted on a custom encoder test rig, depicted also in fig. 6.5, fixed to the catapult arm. The test rig can host up to two encoders, other than the already present ENX Easy 16 relative encoder on the Maxon motor. The motor is used to enable controlled rotation of the encoders during crashes. Of the tested encoders, only the AksIM encoders did not suffer from decalibration or failure. The $\varnothing 29$ encoders were adopted for Skippy's design, and the $\varnothing 39$ was later adopted for the catapult pulley encoder.

6.4. STRUCTURAL DESIGN

The structural design regards the design of the frames that make up for its individual rigid bodies, in particular the leg and torso, since they are the largest. It also involves the design of the smaller bodies' frames, including the foot, lever, follower, rocker and coupler. The leg and torso are designed using CFRP tubes, which are rigidly glued to

⁴*Acknowledgements* — The design of the catapult and encoder test rig and its dynamic analyses have been done by the author; the experiments have been carried out by B. Roodra P. Singh.

aluminium end pieces. These end pieces have a design style that is shared also by the foot and lever.

6.4.1. DESIGN STYLE

Decisions behind the design style of structural components are the result of a trade-off between performance and affordability. High performance is governed by a high specific stiffness and strength.

Highest performance designs can be achieved using topology optimisation and a hybrid of materials. An example of such a design is presented in [Tuerk, Zueger, Klahn, and Meboldt \(2015\)](#), which takes the quadrupedal HyQ as a case study ([Semini, 2010](#)) and presents a custom lightweight shank design. It comprises custom wrapping of CFRP about titanium alloy printed end pieces. The CFRP provides high specific stiffness and strength and the end pieces allow for complex shapes and precise machining, such as required for bearing and encoder assembly. The development and production of such optimised designs is very costly. In addition, the optimised design requires knowledge of all load cases. As such, this approach is more suitable for an evolved project, but not for one that is still in its infancy.

In contrast, a design that features multiple yet simple low-quality COTS and custom aluminium parts screwed together—such as found in Tippy and the original HyQ—allows for an affordable design, but has significantly reduced performance (i.e. a low stiffness and strength to weight ratio). In particular, it was found that the structural vibrations in Tippy were problematic as they limited the balancing controller's performance.

A solution with a performance and affordability in between that of the two above extremes is the use of COTS fibre-reinforced plastic (FRP) tubes that are rigidly glued to millable aluminium end pieces. The use of COTS FRP can also be found in the designs for instance Cassie ([Agility Robotics, 2018](#)) and Stoch ([Dholakiya et al., 2019](#)), both of which make use of CFRP tubes in their structural design. Tubes are closed-section beams, which provide more stiffness than open-section beams, as further explained in section 6.4.2. The design of custom structural aluminium parts—including end pieces to which tubes are connected—are inspired from an II-beam section, as a hybrid between an open- and closed section design, which provides a compromise between performance and manufacturability, as further explained in section 6.4.3.




Figure 6.6: Examples of open and closed beam section designs based on which structural components in robots can be designed.

6.4.2. TUBE-BASED DESIGN

The leg and torso are designed using COTS CFRP tubes, which have a high stiffness and strength and a low density of ca. $\rho = 1.3 \text{ kgdm}^{-3}$, for which they offer high structural performance.

CLOSED VERSUS OPEN SECTION

Since the balancing controller has shown susceptibility to low structural vibrations, it has been decided to maximize the stiffness by basing the design on closed-section beams. Tubes are closed-section beams, which provide much higher torsional stiffness than open-section beams. See fig. 6.6 for various examples of open and closed beam sections. For comparison, torsional stiffness is calculated as $J = 2\pi r^3 t = Pr^2 t = Ar^2$ for thin O-beams, whereas it is calculated as $J = 2/3\pi r t^3 = 1/3 P t^3 = 1/3 A t^2$ for a thin circular C-beam of otherwise identical shape as an O-beam (i.e. an O-beam with a slit). $r = D/2$ is the tube's radius, P the perimeter and A the sectional surface area. Results for differently shaped beam sections (i.e. non-circular ones) are similar. We find that, especially as $t \ll r$, the stiffness is substantially lower for open than for closed beam sections. At a ratio of $r/t = 10$ (e.g. a tube of size $D = 20 \text{ mm}$ and $t = 1 \text{ mm}$) the formulae would imply a factor of 300 difference in stiffness for otherwise identically dimensioned beams. This ratio is furthermore designed to be high in Skippy by choosing a minimum wall thickness (1 mm) to further reduce weight and increase stiffness.

These calculations suggests that closed-section designs are generally preferred over open-section designs. Note nevertheless that open-section designs are not uncommon in robotic design. We find examples of open-section beam designs in for example limbs of the well-known quadrupeds SpotMini and Centauro, featuring I-beam and C-beam sections respectively.⁵ The C-beam section design allows for placement of bigger electronics, but it is loaded less efficiently.

An advantage of open-section design is its assemblability and manufacturability for methods of subtractive manufacturing (e.g. milling), which is the most common applied method for manufacturing custom aluminium frames. However, the COTS tubes in Skippy need not be machined, and do not need to house many components (such as electronics). Nevertheless, the custom aluminium end pieces to which the tubes are mounted do. This subject is further discussed in section 6.4.3.

CFRP TUBING CHOICE

All CFRP tubes in the design are COTS and purchased from Mateduc⁶, which provides several types based on the manufacturing process and material usage.

Structural CFRP tubes in the leg and torso are produced using the process of filament winding. This process allows for any combination of fibre angle lay-up, the use of high performance uni-directional fabrics and efficient use of additives. The default lay-up, consisting of various layers of 90° and $\pm 20^\circ$, provides a high combined performance in flexion, compression and torsion. It is more resistant to impact and has a higher specific strength than tubes produced through pullwinding or wrapping. Compared

⁵The C-beam section design of Centauro is not instantly noticeable, since the limbs appear closed by its plastic covers.

⁶<https://www.mateduc-composites.com/index.php/en/>

to the process of wrapping, filament winding produces tubes with reduced tolerances on the inner and outer diameter ($\pm 0.2\text{mm}$), but this is not problematic for the use of adhesives. Guiding tubes of the main actuator are produced using the process of wrapping instead, as further explained in section 6.7.

END FITTING DESIGN

A disadvantage of CFRP and GFRP is that they are not very machinable: adding holes or features in the tubes leads to destruction of fibres and introduces substantial weakness. The CFRP tubes are glued in fittings that are part of custom aluminium end pieces. The design of the fittings is inspired from those found in industry and other robot designs.

The fittings are external to the tube, as for example observed also in the design of ATRIAS (Hubicki et al., 2016), so that the inside of the tube can still be accessed for placement and fixation of electronics, and so as to reduce the local bending stress by having an increased diameter. As per rule of thumb, the length L of the fitting is designed so that $1.5D < L < 2D$, where D is the tube's external diameter. In addition, similar to the designs described in for instance Fanni and Jweda (2016), and referring to fig. 6.7, the design furthermore includes the following features:

1. a transition fit of $\varnothing D - 0.1\text{ mm}$ to eliminate translational misalignment and reduce orientation misalignment;
2. a conically shaped hull so as to avoid unnecessary material near the opening (which for reasons of manufacturability has a minimum wall thickness of 0.5 mm);
3. a clearance fit of $\varnothing D + 0.3\text{ mm}$ for adhesive;
4. an escape chamfer to reduce strain on the (imperfect) rim of adhesive.

The COTS CFRP tubes have a diameter of $D \pm 0.2\text{ mm}$, and are to be roughened and sanded to a diameter with a tolerance of maximum $D_{-0.2}^{+0}$, so that they fit into the transition fit and to ensure a desired gap thickness for the adhesive of $0.2 \pm 0.05\text{ mm}$.




Figure 6.7: Section view of a schematic design of the external fitting that connects the CFRP tube (dark grey) to the aluminium end piece (blue) with glue (yellow), featuring (1) a transition fit (2) a conical hull shape (3) a clearance for glue (4) an escape chamfer. Not to scale.

6.4.3. DESIGN OF CUSTOM ALUMINIUM STRUCTURAL COMPONENTS

To warrant affordability, all metal components are to be designed so that they are manufacturable using conventional CNC machines, rather than for instance by relying on additive manufacturing technologies. A milled piece of AL7075-T6 aluminium furthermore has superior properties than that of a printed piece, both in ultimate strength and

isotropy. In particular, the parts should be manufacturable by IIT's internal workshop, which has a 3+2-axis CNC-milling machine, comprising 3 automated axes and 2 manually adjustable axes. This requires the part's geometry to be well-defined, and the number of feature planes to be kept to a minimum.

CLOSED- VERSUS OPEN-SECTION DESIGNS

Designs of structural beams can be used as a base to design the structural end pieces to which the tubes are mounted. Ideally, these end pieces also have a closed-section (i.e. hollow) design, to ensure torsional stiffness and to allow for a simple transition to the (likewise closed-section) fitting for the tube. A hollow end piece design would also offer a practical solution for routing cables and shielding other electronics from being susceptible to any type of mechanical interference, whereas an open-section design does not.

Another argument against open-section designs is the high D/t ratio of the parts. The structural components have a width on the order of 4(1) cm, as required for housing the tubes, encoders, and other electronics. The wall thickness is generally to be kept to a minimum of 1 mm, as constrained by manufacturability (with the exception of a few small ribs and flanges of 0.5 mm). This would correspond to a high D/t ratio of approximately 40; which is favourable for improving the specific bending strength and stiffness. This ratio also implies that the torsional stiffness between that of a closed-beam section and open-beam section would be on the order of 10^3 , which further reduces attractiveness of open-section designs. Moreover, it deteriorates bending performance of simple I- or C-beam section designs containing flat material surfaces, such as those depicted in fig. 6.6: the flat upper and lower ribs will barely take any stress as they will bend away, as illustrated in fig. 6.8. Note that high (torsional) stiffness at encoded joints is increasingly important to ensure functionality of the encoders, which are easily affected by a change of the *ride height*, which is the distance between the encoder's reading chip and the magnetic disk it reads.




Figure 6.8: Bending illustrated of a vertically loaded thin I-beam ($D/t = 40$), which shows the top and bottom open end plates turning away and taking less stress: effectively losing their effectiveness.

However, closed-section designs have the chief disadvantage that they are not (easily) machinable from the inside, and they furthermore complicate assembly. The parts do require precise features for bearing and encoder placement. One possible solution is to split the part in two (or more) and screw it together during assembly. To guaran-

tee high performance (as if it part was not split), the matching parts should be properly locked to each other by either using many positioning rims/pins or many screws (such as observed on the thighs of ANYmal), of which the former (rims/pins) requires high tolerance (and thus costs) and the latter makes the robot heavier. Note furthermore that the tube fitting must be present on only one of the halves, so that the parts are asymmetric. The exact behaviour and loading of these (asymmetric) halves is more unpredictable, which is undesired especially (1) for lightweight design (the design would have to be over-engineered) and (2) since parts are going to be rigidly glued to each other (if one end piece fails, the whole limb has to be disposed of).

With the advantages and disadvantages in mind, an attempt was made to adopt a closed-section hollow design style. Figure 6.9 shows a conceptual design of Skippy's four-bar linkage that comprises hollow end pieces that each consist of two halves screwed together. The hollow design features placement of the lever on the inside of the end pieces, which would not be possible if the end piece was a single part. The advantage of this configuration is that the lever itself can be designed to be simpler, but the disadvantage is that the end piece of the torso now has to be longer, as the lever cannot physically move through the tube. The complexity (regarding manufacturability) and mass of these parts were disappointing, and so it was decided to overhaul the design style.

6



Figure 6.9: The left and centre figure depict Skippy's four-bar linkage in the new and old design style respectively. The right image shows a prototype 3D print of the old design, which was found to be unsatisfactory regarding both manufacturability and performance.

HYBRID DESIGN

A design style was adopted based on the II-beam section depicted in fig. 6.6, which is a hybrid between closed-section and open-section design. The inside allows only for rough material removal of simple shapes, where the open sides can be machined to improved tolerance and detail, such as required for bearing holes, encoder mounts and cavities/ribs for strengthening the parts where necessary. The hollow inside significantly improves stiffness and strength of the part, and makes the side cavities more manufacturable and its corresponding ribs more effective. It furthermore enables a smooth transition to the tube fitting for end pieces and provides space to guide cables. The design procedure of the lower end piece of the torso (the piece connecting to the four-bar linkage) is presented below. It is presented as an example, because the design procedure of

other frame parts is very similar, including that of the foot, lever and leg end pieces.

- First, the contours of the part are defined, which are typically constrained by the range of motion and size of other parts (and is as such also the result of iterative design between other parts). For this part, they are constrained mainly by the locations of the joints, the size of the encoder ($\varnothing 38\text{mm}$), and the size and motion range of the lever and follower, which are both mounted to the exterior. See fig. 6.10a.
- Second, the necessary side features are defined, such as those required to mount the encoder PCB. The features also include tentative dimensioning of pockets and ribs for weight-saving. See fig. 6.10b.
- Third, the part is hollowed with simple features, leaving bridges where stiffness and strength are required (e.g. at encoders and bearings). The number of undercuts is kept to a minimum, by minimizing the number of machine setups. For this part in particular, inner cuts can be made from only two different machining angles, as illustrated in fig. 6.10c.
- Last, the part is finalized through a few iterations of FEAs, based on which small features are applied or changed. These include change of radii and wall thickness and the addition of holes in the side. The main purpose of the FEAs is validating that the parts are sufficiently strong; the secondary purpose is weight reduction (i.e. validating that the parts have not been substantially over-engineered). Essentially, following the design philosophy presented in chapter 2, the parts are designed to be good enough, by making an educated trade-off between validating its functionality and weight reduction; no time-consuming topology-optimisations have been performed towards finalizing designs. The finalized part is depicted in fig. 6.10d and the resulting weight reduction is typically 10-20% per part. As a side effect, the resulting design is aesthetically pleasing.

FEA

Finite element analyses are done in Siemens NX 11, using its built-in linear static solver, and are used to finalize a part after the completion of relevant optimisation studies that are further described in chapter 8. The FEAs are done for several force equilibriums that correspond to extreme scenarios of Skippy, which include that of (1) Skippy preparing its highest jump while being close to the nadir during the stance phase (corresponding to the second and third frame in fig. 3.2), and (2) Skippy crashing. Joint forces in the sagittal plane are typically dominated by the former, whereas out-of-plane forces and moments by the latter.

As an example, fig. 6.11 depicts the approximate force diagram as implemented in an FEA study for the torso, which corresponds to the scenario of Skippy preparing its highest jump. In this analysis, the lever joint was decided to be fixed and other forces were defined in such a way that they are approximately in torsional equilibrium about the lever joint.

$F_f = 4\text{ kN}$ and $F_s = 2\text{ kN}$ are the maximum force estimates as exerted by the follower and rocker (main actuator) respectively. F_f is estimated based on analysing the moment




Figure 6.10: Frame design

arms of the lever a_l and follower a_f . The moment arm of the lever is the distance between the line defined by the actuator and the lever joint on the torso. The moment arm of the follower is the distance between the line defined by the follower joints and the lever joint on the leg, see fig. 6.14. This follows from the fact that both the actuator and follower act as mechanical rods, since they are hinged freely at both ends. When the main actuator exerts its maximum force $F_s = 2\text{ kN}$, Skippy is nearly flexed, for which a_f is designed to be 4.5 cm to 5 cm⁷, and the lever is designed to be approximately proportional to the actuator to maximize the mechanical advantage, so that $l_l \approx a_l \approx 10\text{ cm}$, which is approximately double a_f . Since most of the moment generated by the actuator and lever is transferred through the hip (neglecting inertial forces), it follows that $F_f a_f \approx F_s a_l$, from which it follows that the load on the follower joints reaches up to approximately double that on the coupler.

$F_g = 0.2\text{ kN}$ is the approximated inertial force as caused by Skippy's upward acceleration of ca. 10g. Inertial moments are not considered in this equilibrium, as they are negligible for the relatively slow rotating bodies. Note furthermore that it is not important these forces be very accurate, as they only serve the purpose of obtaining a rough indication of whether material is over- or under-stressed. The exact forces are moreover not known and differ per experimental trial in practice. It is anticipated the forces are

⁷Parameters of the four-bar linkage have been carefully chosen to take the follower's desired moment arm when fully flexed $a_{f,\text{flx}}$ and extended $a_{f,\text{ext}}$ into account, as further explained in section 8.3.3.

correct within a margin of $\pm 25\%$ in magnitude and ± 5 to 30° in direction.

Material is locally added (e.g. by increasing round radii and wall thickness) or removed (e.g. by reducing round radii wall thickness or adding holes) if the part is locally over- or under-stressed respectively, after which finite element simulations are repeated. Approximately two iterations are required to obtain satisfactory results. FEA results of the finalized torso lower frame are depicted in fig. 6.12a for the force diagram in fig. 6.11, which shows elemental Von Mises stress in colour, reaching values up to 0.3 GPa (ca. 60% of the yield stress of Al7075-T6). A prototype of the finalized part has been machined by the internal workshop to demonstrate its manufacturability, see fig. 6.12b.

Similar studies are also performed for the leg, foot and lever.



Figure 6.11: Free body diagram of the torso depicting forces as implemented in an FEA. The forces correspond approximately to those of Skippy close to its nadir while preparing its highest hop, which is one of the most demanding scenarios. The red dot is the fixed joint in FEA; its reaction forces are not drawn.




Figure 6.12: Images depicting the lower end piece of the torso's frame, with left an image showing stress distributions of the FEA corresponding to the load case depicted in fig. 6.11 and right the machined part.

6.5. HIP DESIGN

Figure 6.13 depicts the design of the hip, featuring the lever, follower and coupler. Its morphology is unchanged with respect to that depicted in fig. 3.1a from earlier design studies. The follower, torso lower frame, leg upper frame and lower half of the lever make up the four-bar linkage. All of the bodies of the four-bar linkage are hollow, so that cables

can easily be guided from the leg to the torso through the sagittal plane. A new addition to the hip is its end stop, which acts as a parallel spring between the lever and torso.

DESIGN STYLE AND LOADING

The hip is designed so that all torsional loads and bending loads out of the sagittal plane on the torso and leg are transferred through the lever, instead of through the follower. The structural design of the lever is like that described in section 6.4, based on a closed-section II-profile, and offers both bending and torsional stiffness. The follower instead only transfers a radial load in line with itself, for which it acts as a rod. The load is mainly a compressive load, and has been designed so that it would not fail by buckling. It consists of two identical parts with C-beam profiles, held together by two pivot pin axles and two screws. The Euler buckling load is calculated as 21 kN, under the pessimistic assumption that the follower is a C-beam over its full length (ca. 120 mm). As explained in section 6.4.3, the expected maximum force on the follower is approximately 4 kN, which is well under the calculated buckling load.

BEARINGS

Because the lever axles also transfer non-radial loads, the torso and leg joints to which the lever is hinged are equipped with combined needle roller/axial ball bearings (NX7-TV-XL). These have a combined static radial load rating of 5.3 kN, and an axial load rating of 4.3 kN. The advantage of using these type of bearings compared to a set of angular contact ball bearings, is that

- they require simpler mechanical features with less precise geometrical tolerances;
- they have a smaller diameter, which enables the use of simpler (non-hollow) COTS axles, and offers advantages for placement of the selected encoders (which too have a small inner diameter);
- they are generally more resistant to impact (roller bearings make a line contact rather than a point contact);
- they offer a higher load-to-weight ratio, and thus offer a good solution for lightweight design.

The disadvantage of these type of bearings is that (1) they offer slightly less precision (increased play), (2) they take up more axial space and (3) they require hardened axles, since the bearings do not have an inner ring. However, these effects and requirements are not problematic for Skippy's four-bar linkage. First, the joints are not actuated, and only one joint is encoded, so precision is not an issue for most joints. Accurate axial alignment, as important for the encoder in the lever-torso joint, can be realized by using shim rings (red in fig. 6.13). The needle bearings are furthermore of the type 'massive', in which play is smaller than in that of ordinary 'drawn cup' needle bearings. Play can furthermore be reduced by choosing an axle tolerance class that is one or two sizes bigger than nominal, and can be eliminated by axially pre-loading the bearing using a screw. Second, the extra required axial space is no issue, because there is no space constraint in the axial direction. Instead, space is radially constrained in one of the joints because of the the encoder. If the axle diameter were too big, the encoder would have to be placed on the exterior of the body, where it is more vulnerable. Third, the requirement

for hardened axles is no issue, because the axles are simple $\varnothing 7$ mm COTS ‘precision pivot pins’, available from MISUMI (GFCLSHDF) with chrome hardening.

Both the follower and coupler axles transfer only radial loads. The bearings in these joints are NK7-10-TV-XL ‘massive’ needle roller bearings, identical to the needle bearings in the combined bearings NX7-TV-XL, with a radial load rating of 2.65 kN per bearing. The coupler joint has one bearing, whereas the follower joints have two bearings each. The radial load on the follower joints is approximately double that on the coupler joint, which is related to the size of the follower arm a_f and lever length, as explained in section 6.4.

The axles of the follower and coupler joints are also COTS precision pivot pins, but of different length. In contrast to the lever joints, they are not axially preloaded. The screws depicted in the left side of fig. 6.13 are merely used to fix the axles, and nylon spacers (orange) are used to centralize the bodies. Note that in total, the entire hip features only five structural screws (excluding screws to fix the encoder), equal to the number of axles.

HIP END STOP

The hip end stop is indicated in fig. 6.13. It is mounted to the lever and compressed by the torso, which is more clearly visible in fig. 6.14. The end stop has been introduced in response to the anticipated increase of Skippy’s mass to boost its jumping performance. It has been designed to do so without compromising on Skippy’s in-flight and balancing performance, by (1) only being effective for a near-flexed configuration and (2) by not exerting substantially more than double the torque that would be required for gravity compensation while statically balancing for a fully flexed configuration. Compression occurs for only the first 10° of lever rotation from Skippy’s flexed configuration. The configuration of the hip depicted in fig. 6.13 corresponds to the touching condition of the end stop. Since the maximum force of the end stop is designed to not be substantially larger than that required for countering the gravitational force, it will not store more than several joules. Corresponding calculations are further discussed in section 8.3.3.

COLLISIONS AND PROTECTION

The hip design is carefully tailored to prevent internal collisions, which is challenged by its 180° joint range and bumper placement. The hip’s motion range is shown in fig. 6.14.

The first decision involves centralization and order of bodies. As described in section 6.4, it was decided to centralize the frames of the torso and leg, so that they consist of only a single part and so that they are symmetric about the sagittal plane, which is beneficial for fitting them with end fits for the CFRP tubes. Both the lever and follower are then placed external to these frames.

The decision for the follower to be placed external to the frame is such that it can be fitted with a bumper that rotates with the follower, as shown in fig. 6.14. Bumper placement on the follower (instead of on the torso) is beneficial as it rotates towards the direction it needs to protect from impact, and also introduces fewer design problems regarding internal collisions.

The lever is placed external to the frames so as to avoid collisions with the torso tube when Skippy is extended (as would be the case if it was internal, see fig. 6.14a). External placement of the lever is furthermore beneficial for its required out-of-plane stiffness,




Figure 6.13: The hip design, including the lever, follower and the coupler and their axles and bearings.

6

as this allows it to be wider. The follower is furthermore designed to move through the lever, and has a width approximately equal to that of the torso and leg frames.

The lever is equipped with two bumpers, which also protect the hip from sideways impacts (out of the sagittal plane), in contrast to the bumper on the follower, which points only to the back in the sagittal plane. The bumper of the follower is aligned with the upper follower joint, because its frame is not designed to take bending loads.

Most of the contours of the lever, torso and leg frames are defined based on collision prevention, of which some considerations are described below.

- The lever outer contours are defined so that it and the main spring do not collide with the torso and follower-torso axle and bumper when fully extended (see fig. 6.14a) and the follower-leg axle when fully flexed (see fig. 6.14c). This has resulted in the lever's S-shape.
- The lever upper cavity and torso lower frame contours are defined so that they do not collide with each other when fully flexed and extended.
- The lever lower cavity contours are defined so that it does not collide with the follower and leg upper frame when fully flexed and extended.
- The torso lower frame features cavities on its sides so that it does not collide with the follower when fully flexed; idem for the leg upper frame to prevent collisions when fully extended.




Figure 6.14: Skippy's hip range. The distance between the red dashed lines is the moment arm of the follower a_f . When Skippy is fully flexed (more torque) $a_f = a_{f, \text{flx}} = 5 \text{ cm}$ and when it is fully extended (more speed) $a_f = a_{f, \text{ext}} = 2 \text{ cm}$.

6.6. ANKLE AND FOOT DESIGN

The design of the ankle and foot is displayed in fig. 6.15.⁸ It features the ankle spring, an end stop and the same Renishaw/RLS encoder that is used in the four-bar linkage and swivel actuator.




Figure 6.15: The design of the foot and ankle. This figure shows the ankle and foot design as finalized in chapter 8, which is slightly different from that depicted in fig. 6.1.

⁸Note that this figure depicts the final design following the optimisations in chapter 8, which is slightly different from that shown in fig. 6.1.

6.6.1. ANKLE SPRING DESIGN

The ankle torque-rotation profile ideally has an S-shape, as described in [Driessen \(2015\)](#), where stiffness around zero and maximum rotation is high, but low in between. This implies that the initial part of the stiffness profile is regressive and the second part progressive. The desired regressive behaviour is realized by the use of a curved compression leaf spring made from GFRP, and the progressive behaviour is realized by the use of a rubber end stop.

CURVED COMPRESSION LEAF SPRING

The GFRP spring is hinged between the foot and lower leg frame. This construction forms a triangle that has one edge of variable length (i.e. the spring's edge), as shown in [fig. 6.16a](#). The compression spring has a regressive force-displacement profile of its own, and the regressiveness can be altered by changing the spring's curvature. [Chapter 7](#) explains in detail how the stiffness profile of curved compression springs can be accurately described using a pseudo-rigid body model (PRBM), and how this model can be used to predict failure. The resulting torque-rotation profile of the joint is typically also regressive and can be further altered by changing parameters that make up the triangle kinematics.




Figure 6.16: The ankle spring design. This figure shows the ankle and foot design as finalized in [chapter 8](#), which is slightly different from that depicted in [fig. 6.1](#).

GFRP GFRP is used because it has a high gravimetric energy, especially if choosing proper fibre alignment. The spring cannot move or rotate out of the sagittal plane due to constraints imposed by its hinges, which allows for unidirectional (UD) fibre alignment in the longitudinal direction of the spring. In the fibre direction, realistic assumptions for material properties are a Young's Modulus of $E = 35 \text{ GPa}$ and a maximum tensile and compressive strain of $\epsilon_{xt} = 0.03$ and $\epsilon_{xc} = 0.021$ respectively ([Kussmaul, 2016](#)). Note that

these values depend also on the quality of the glass fibre, and could be even higher for high quality S-glass fibres. With a density of 1.9 kg dm^{-3} , the gravimetric energy density of UD GFRP can then be estimated as $1/2\epsilon^2 E/\rho = 4 \text{ kJ/kg}$. This value assumes failure in compression; it also assumes failure by the fibres instead of the bond. In comparison, the gravimetric energy density of a high quality spring steel with a yield strength of $\sigma_y = 1.2 \text{ GPa}$, $E = 200 \text{ GPa}$ and $\rho = 8 \text{ kg dm}^{-3}$, can be estimated as $1/2\sigma_y^2/E/\rho = 0.45 \text{ kJ/kg}$, approximately a factor of 10 lower.

Mounting The GFRP spring is rigidly glued to aluminium end pieces with a $\varnothing 3 \text{ mm}$ hole, through which it is hinged to the foot and leg using an off-the-shelf (Misumi) axle and two miniature needle bearings per side. The two end pieces, two axles and four bearings weigh 15 g and can take up to a load of 1.6 kN (as limited by the bearings). The hollow foot design and a prismatic cavity in the leg allow it to move freely without colliding for the ankle's full range of motion.

The interiors have a width of 40 mm, which allow for a spring of up to 38 mm wide to be mounted. The width is based on the width of the wide ankle bearings (NX7-TV-XL), rather than the smaller leg tube ($\varnothing 25 \text{ mm}$). This wide foot design allows for flexibility in design for strong springs.

ANKLE END STOP

The desired progressive behaviour is realized by the use of an end stop made from rubber, as shown in fig. 6.15a. The end stop is only compressed when the ankle joint approaches its maximum permissible rotation. During normal operation, the end stop should only be touched rather than be compressed substantially, but it should provide sufficient resistance to prevent the GFRP from breaking and the foot from hitting the leg, as could be the result after an unfortunate impact. The rubber end stops are off-the-shelf components from Misumi, available in various hardness and size, to allow for easy replacement and adjusting of end stop characteristics.

6.6.2. ANKLE TORQUE PROFILE CALCULATIONS

The ankle's torque-rotation profile is the sum of that contributed by the compression spring and that of the end stop. Only the former needs to be calculated accurately, as the end stop is not to be compressed substantially during normal operation.

CURVED COMPRESSION LEAF SPRING

The contribution of the compression spring to the ankle joint's torque-rotation profile can be calculated using the triangle kinematics function γ_{\uparrow} defined in appendix B as

$$\{l_a, \theta_b, \theta_c, \mathbf{G}_a\} = \gamma_{\uparrow}(\theta_a, l_b, l_c).$$

$\theta_a = \theta_{a,0} - q_5 + q_{5,0}$ is the triangle angle opposing the spring, with $\theta_{a,0}$ the rest angle, q_5 the ankle joint angle and $q_{5,0}$ the ankle joint angle for which the ankle spring is at rest (i.e. for which $\theta_a = \theta_{a,0}$). l_b and l_c are the fixed edge lengths of the triangles, as shown in fig. 6.16a. $l_a = l_{a,0} - x_a$ is the spring length, with $l_{a,0}$ the spring rest length and x_a the spring's displacement in compression. \mathbf{G}_a is the Jacobian, for which we find

$$G_{a,1} = \frac{\partial l_a}{\partial \theta_a} = \frac{-\partial x_a}{-\partial q_5} = \frac{\partial x_a}{\partial q_5}.$$

The compression spring force F_a can be calculated as function of the spring compression x_a , as described in chapter 7. It follows that the ankle torque τ_5 can then be calculated as follows:

$$\tau_5 = G_{a,1}F_a.$$

6.6.3. PARAMETER DESIGN CONSIDERATIONS

Sizing and shaping the foot is an involved process due to conflicts of interests and various adverse effects on performance and collision constraints. For example, for optimal spring behaviour, it is desired for the ankle spring rest angle to be substantially larger than the joint range (further explained below). This implies that the attachment point of the spring is desirably “buried” deep in the leg, as shown in fig. 6.16, which is facilitated if the lower leg frame is designed to not have a big “kink” like that shown in fig. 3.1a (effectively placing the ankle joint above the line of the leg tube). Yet, the end stop is desirably placed as far as possible from the ankle joint so as to reduce the forces on both the ankle joint and bearings; a kink facilitates its placement far from the joint (effectively bringing the leg and foot closer together for maximum flexion) without requiring either of the foot or lower leg frames to be extended (introducing extra mass).

As a compromise, in the design depicted in fig. 6.1 the ankle joint is placed only 12.5 mm above the leg tube. However, in the final design (following the optimisations in chapter 8), it was decided to remove the kink completely by placing the ankle joint in line with the leg tube, which is the design depicted in figs. 6.15 and 6.16.

6

END STOP PLACEMENT AND FOOT CURVATURE

The end stop is placed on the foot at an arm of approximately 5.5 cm from the joint, and makes contact with the leg at a joint angle of $q_5 = 45^\circ$.

The arm length is chosen based on being sufficiently far from the joint to withstand impact forces during an unfortunate crash, whilst preventing the necessity of extending the foot frame far beyond the required size of the prismatic cavity in which the spring moves. The short arm length is made possible partly due to the choice of the same powerful combined needle roller/axial ball bearings as found in the four bar linkage joints (NX7-TV-XL), which offer substantially more peak load than required during normal operation. An impact can lead to a force up to approximately 5 kN on the end stops and bearings, equal to the bearings’ peak load.

The mounting angle was decided based on early simulation studies, and can be tweaked in practice by installing end stops of different thickness. These studies indicated a desired joint range of $\Delta q_5 \approx 60^\circ$, and furthermore showed that the foot would start rolling on its back if it would bend too far towards the leg. This could be solved by removing the foot’s curvature (i.e. creating a straight foot), but would in turn cause the spring to collide with the foot’s interior. Referring to fig. 6.16a, ϕ_{tx} is used to define the foot’s curvature ($\phi_{tx} = 90^\circ$ would correspond to a straight foot) and also defines the minimum toe angle $q_{4,\min}$ (the foot would start rolling as $q_4 \lesssim \phi_{tx}$). ϕ_{tx} is set to $\phi_{tx} = 75^\circ$ based on observing that $q_4 > -75^\circ$ for various simulations of the stance phase, whilst leaving sufficient space for designing the spring and its operation space, which is further explained below.

INCREASING $\theta_{a,0}$ AND l_c/l_b

Through early simulation studies it was found that the rest angle of the spring $\theta_{a,0}$ is ideally substantially larger than the range of the ankle joint, likely so as to avoid a bi-stable profile of the ankle torque profile and to decrease maximum compression of the ankle spring, each of which are further elucidated below.

A bi-stable profile implies that the ankle torque-rotation profile is not monotonic. Whereas $F_a(y_a)$ is monotonically increasing (see chapter 7), $\tau_5(q_5)$ is not necessarily monotonically increasing, which is due to the triangle kinematics. The spring's moment arm is maximal for $\theta_a = \theta_{a,opt} = \arccos(l_b/l_c)$ (assuming $l_b < l_c$), and will thus decrease if the ankle is flexed beyond this angle ($\theta_a < \theta_{a,opt}$), subsequently reducing the mechanical advantage (i.e. $\partial^2 x_a / \partial q_5^2 < 0$). This will not instantly cause the spring profile to become bi-stable, because $F_a(y_a)$ is still monotonically increasing, and bi-stable behaviour at the very highest loads are not necessarily problematic. Nevertheless, angles substantially lower than $\theta_{a,opt}$ should be avoided.

$\theta_{a,opt}$ is typically an angle between 60° ($l_c = 2l_b$) and 75° ($l_c \approx 4l_b$), which suggest that $\theta_{a,min} \gtrsim 45^\circ$. Since the range of the ankle joint is approximately $\Delta q_5 \approx 60^\circ$, $\theta_{a,0} \gtrsim 105^\circ$.

In addition, reduced maximum compression of the ankle spring is desired to reduce the spring's required operation space and to reduce its chance of failure, which can be also achieved by increasing $\theta_{a,0}$ beyond the range of q_5 , and furthermore by increasing the ratio l_c/l_b .

The triangle's short side l_b is placed on the leg so as to reduce the required size of the aluminium end-piece for mounting the spring. Its minimum length and $\theta_{a,0}$ are constrained by collisions with the ankle encoder and its hull. The maximum length of l_c is constrained by the foot length l_p : l_c is defined to be 24 mm shorter than the foot: $l_c = l_p - 0.024$. These subsequently constrain l_c/l_b ; depending on $\theta_{a,0}$ and l_p , we find $l_c/l_b \approx 3$. Below follows a more elaborate explanation of how collisions constrain various parameters and the design.

SPRING ATTACHMENT AND COLLISIONS

The attachment points of the spring to the leg and foot can be defined by triangle sides l_b and l_c and mounting angles ϕ_b and ϕ_c (see fig. 6.16a), so that the spring's rest angle is related to $q_{5,0}$ as $\theta_{a,0} = \pi/2 + \phi_b + \phi_c - q_{5,0}$. The attachment points must be chosen so as to avoid collisions between the spring, foot and leg.

Non-linear FEA studies of the flexible spring have been combined with shapes of the foot and leg as defined in CAD software to create figures such as that depicted in fig. 6.16b, based on which is checked for collisions, and subsequently allow for the design to be finalized.


- ϕ_c is chosen firstly, and set to $\phi_c = -5^\circ$ to prevent collisions of the spring with the inside of the foot when fully flexed. This choice also assumes a spring of mild curvature, with a radius of $R = 180$ mm.⁹
- The shape of the prismatic cavity in which the spring is mounted to the leg is designed to prevent collisions with the spring.

⁹ *Acknowledgements* — Early optimisations conducted by Dr. Featherstone showed that the spring must not be too regressive.

- ϕ_b and l_b are chosen as a compromise between maximizing $\theta_{a,0}$ and minimizing l_b whilst preventing collisions with the encoder and its hull (see red circle in fig. 6.16b), and preventing the need of excessive frame that would protrude next to the heel, as required for housing the spring's upper axle. In the design depicted in fig. 6.1, $l_b = 5$ cm and $\phi_b = 5^\circ$. ϕ_b was later changed to $\phi_b = 15^\circ$, as permitted by results of the optimisations in chapter 8, which have indicated a desire for a reduced initial angle $-q_{5,0}$ and have resulted in the updated ankle design depicted in figs. 6.15 and 6.16. Collisions are prevented as long as $l_p \gtrsim 15$ cm and $q_{5,0} \gtrsim -15^\circ$.

6.7. MAIN ACTUATOR DESIGN

The main actuator consists of the Maxon DCX32L brushed motor, screw drive and the main spring, all connected in series. The screw drive is by default a ball screw, but can be replaced for a roller screw for comparison, and is guided and protected by a lightweight telescopic guiding mechanism. The main motor, screw spindle and their housing make up the rocker body, which is hinged to the head. The actuator is furthermore designed according to the desired weight distribution of the torso, as described in chapter 3, where mass is concentrated at either of its ends, rather than in the centre. For this reason, the main spring has been placed as close as possible to the hip joint. The rocker, guiding mechanism and main spring are further explained below.



(a) Main actuator using a ball screw, depicted for Skippy fully extended (above) and flexed (below). The actuator length is limited by the guiding arrangement, a light-weight solution to resist buckling and provide ingress protection.

(b) Main actuator using a roller screw, depicted for Skippy fully extended. The roller screw spindle has a smaller diameter and is thus lighter, yet its nut is larger and thus requires a larger diameter guiding system.

Figure 6.17: The main actuator with either a ball or roller screw.

6.7.1. ROCKER

The rocker joint is designed to be as close as possible to the motor so as to reduce the relative motion of the motor end with respect to the torso as caused by the rocking motion of the rocker. This allows for a smaller head and reduced motion of power and encoder (ENX 16 EASY) cables, which are connected to the back end of the motor. The motor is screwed directly to the rocker via its customized flange through 4 out of its 8 screw holes with M4 flat head screws.

The rocker is hinged in the head with two narrow 61803 thin section bearings, as depicted in fig. 6.17a, which have a combined static radial load rating of 3 kN. The axes of the actuator (motor/spindle) and rocker joint are designed to intersect to prevent bending loads on the spindle. The rocker joint axle is therefore designed to consist of two pieces, to be screwed in the rocker from the side, so that it does not intersect with the actuator axis, and to offer space for mounting a coupling.

The screw spindle is held by special lightweight axial (thrust) angular-contact bearings (ZKLN0619-2Z-XL), at a mass of 0.02 kg, featuring a contact angle of 60° and a static axial load rating of 6 kN, preclamped at 2 kN to 3 kN by a ZM06 (M6 × 0.5) precision lock nut to remove play.

The interior of the rocker offers space for the coupling, which is preferably rigid and lightweight. The spindle's end is $\varnothing 5$ mm and the motor axle is $\varnothing 6$ mm. The two can be connected using (for example) the custom hub depicted in fig. 6.17, which is glued directly to the motor shaft, and connects to the spindle using 4 M4 grub screws (2 × 2 under an angle of 90°). A possible lightweight alternative is the use of a COTS 6 mm N-coupling¹⁰ in combination with an open-section adapter to increase the spindle's mounting diameter from 5 mm to 6 mm.

6.7.2. GUIDING SYSTEM

The actuator is freely hinged at both ends to prevent bending loads on the spindle. However, this also makes the arrangement more prone to buckling during compression loads. The system is designed to be used mainly in tension, but minor compressive loads are inevitable during balancing and in-flight control, resulting in compression forces up to several 100 N. The arrangement is weak to buckling because of the flexible (series) main spring and the nut, which are both weak in bending, meaning that even a load as low as 100 N might be threatening.

To make the actuator more robust against buckling, it has been equipped with a guiding system, and the main spring has been made hollow. The presented guiding solution was found to be a lighter and more compact solution than the alternative of using guiding rod(s) to either side of the actuator, which moreover does not offer the additional advantages such as ingress protection. The guiding system weighs 61 g, which includes the tubes, end pieces (except the coupler frame) and bushings.

TELESCOPIC ARRANGEMENT

The guiding system is a dual-stage telescopic arrangement, and consists of three tubes that are coaxial with the actuator to guide both the nut and the spring. The part connecting the nut to the spring, referred to as *nut-spring tube* and depicted green in fig. 6.17,

¹⁰Misumi CPN6 <https://us.misumi-ec.com/vona2/detail/110300127250>

double-functions as the inner guide for the *nut guide tube* and outer guide for the *spring guide tube*. Both the spindle and spring guide tube are depicted in black in fig. 6.17. The tubes are best distinguished in the second image of fig. 6.17a. The nut guide double-functions as an ingress protector for the screw drive, and has a third function as the outer clamp of the thrust bearings in the rocker, so that the bearing housing need not be designed with a press fit.

Note furthermore that the system relies on the ability of the spindle to move inside the (smallest) spring guide tube, as visible in the upper picture of fig. 6.17a and fig. 6.17b. The spring is made hollow (tube-shaped) in order for the spring guide tube to be placed on the inner side of the spring. A hollow spring furthermore improves the spring's bending and torsional stiffness (and thus resistance to buckling, were there no guides). A gap between the spring and its guide ensures that the two do not come in contact with each other when the spring is stretched.

CONNECTIVITY

The tubes are rigidly glued to aluminium end pieces, using the end fitting design discussed in section 6.4, which are in turn connected to their respective parent parts using a fine metric thread, so that the arrangement requires only four threaded connections, which—compared to using ordinary connector flanges with metric screws—saves significant weight and space.

Telescopic sliding is realized by the use of IGUS G-Sleeve bushings, which are fixed on the inside of the outer tubes. The inner tubes slide directly in the bushings.

TUBING CHOICE

To minimize weight, the tubes are made from CFRP and have a wall thickness of 0.5 mm, which is the minimum thickness available. Of the three tubes, only the nut-spring tube is loaded axially. With an outer diameter of $\varnothing 20$ mm, and an assumed failure stress of 200 MPa in tension (pessimistic), it should be able to withstand at least 12 kN, which is a factor of 6 higher than the expected maximum force exerted by the nut.

The tubes desirably have a high tolerance outer diameter and a smooth outer surface finish, since they are directly fitted with bushings and used as sliders. For this reason, the CFRP guides are produced using the process of wrapping, rather than filament winding as used for frame tubes, as explained in section 6.4. This technique provides higher tolerance on the inner and outer diameter (± 0.05 mm) and a smoother surface finish. Unlike the frame tubes, these wrapped tubes are not as isotropic since they provide strength mainly in flexion (rather than also in torsion and compression), which fits the guides' loading profile. An additional advantage is that the COTS wrapped tubes are generally available in more sizes with wall thickness as low as 0.5 mm.

6.7.3. MAIN SPRING

The main spring is a series spring that connects the nut to the coupler. Like the ankle spring (for initial deflection), its force-displacement profile is desirably regressive, which suggests that the previously introduced curved compression leaf springs are also good candidates for the main spring. However, unlike the ankle spring, the main spring works mainly in tension, and it stores substantially more energy than the ankle spring, for

which it was decided to use a rubber tension spring instead. The rubber spring takes the shape of a tube, so that the guiding mechanism can be centralized with the spring, and to make it stiffer in torsion and bending.

ENERGY DENSITY

The main spring should be able to store up to approximately $E_s = 80\text{J}$, which is several factors higher than the energy stored by the ankle spring, and as such implies that the need for a spring with a high gravimetric energy density is more critical. Rubbers exist with an energy density exceeding $\rho_{\text{Eg}} = 10\text{kJkg}^{-1}$, higher than the ca. $\rho_{\text{Eg}} = 4\text{kJkg}^{-1}$ of GFRP. In addition, the high ultimate tensile strain of rubber allows it to be used in pure tension, which leads to efficient material deformation of nearly 100%. In coil and leaf springs, material deforms less uniformly, effectively reducing the gravimetric energy density of the spring as a whole by approximately a factor of $1/\eta_E = 10$, where η_E can be described as a spring's geometrical energy storing efficiency. A spring's effective gravimetric energy density could then be calculated as $\bar{\rho}_{\text{Eg}} = \eta_E \rho_{\text{Eg}}$. A GFRP leaf spring would then be required to have a mass of $E_s/\rho_{\text{Eg}} \approx 0.2\text{kg}$, and a spring made of high quality spring steel would have a mass of ca. 2 kg. Note that these figures still exclude the mass required for their assembly.

Unlike the ankle spring, the main actuator mainly functions in tension rather than in compression. This implies that the use of curved compression leaf springs would require an extension and fortified guiding mechanism to turn the tension spring into a compression spring, which easily adds another 0.2 kg to the already heavy GFRP springs of 0.2 kg.

REGRESSIVENESS

One disadvantage of a pure tension spring compared to a spring that makes use of geometry (such as the curved compression leaf springs), is that it offers less flexibility in designing for a specific stress-strain profile and thus for regressiveness. Rubbers naturally have a regressive stress-strain profile, but it is not as regressive compared to what could be achieved with the curved compression leaf springs. According to the Neo-Hookean model for incompressible hyper-elastic materials, the elastic modulus at $\epsilon = 0$ is (only) three times as high as that for $\epsilon \rightarrow \infty$.¹¹ A more regressive profile could be obtained by combining the rubber tension spring with a rubber compression end stop. Such an end stop could be incorporated in the design by slightly altering the guiding mechanism.

This end stop is not incorporated in the current design, for the reason that it adds more weight and complexity whereas its necessity is disputed and undetermined. For jumping, the regressive behaviour is desired to maximize energy storage for a given maximum force and displacement, but it is likely not as critical as for balancing, which could benefit significantly from improved stiffness. However, the exact influence of stiffness on balancing performance has not been quantified and is not going to be further investigated in this thesis. If it is concluded from future studies that the spring should be substantially more regressive, then the end stop can be added in a subsequent design cycle by slightly altering the guiding mechanism.

¹¹Note that even though the Neo-Hookean model is only accurate for strains up to ca. $\epsilon = 0.5$ (Treloar, 1973), these figures provide an indication of the regressiveness that can be expected from rubber tension springs.

RUBBER SOURCING

Many types of rubber exist with different material properties. The main spring's rubber is required to have a high gravimetric energy density and a high elastic resilience (ability to store instead of dissipate energy). It is furthermore desired that the rubber has a regressive stress-strain profile.

A major drawback of most rubbers is that they typically have undesired material properties, like viscosity and creep, and that they are difficult to characterize, because most of the properties are governed by non-linear relations and time-dependent. In addition, rubbers can permanently change (deteriorate) over time, due to vulnerability of being exposed to oxygen, oil, humidity or sunlight.

Most industrial applications of rubber focus on its use as shock absorbers, which requires a low elastic resilience instead, or rely on its longevity in exposed environments, which requires the addition of chemicals that typically reduce a rubber's mechanical properties. As such, it is difficult to find a COTS rubber with the right properties. Natural rubbers typically have a high elastic resilience, but were not available in useful form from our suppliers. A synthetic rubber with a high elastic resilience is that used in the Super Ball bouncing ball (Wikipedia, 2019), which is a synthetic polymer polybutadiene with other ingredients. An attempt has been made to mould a custom high quality rubber spring using a (similar) liquid polybutadiene rubber, but the result was not satisfactory due to its weakness (having a low gravimetric energy density).

The main spring design currently consists of a COTS tube-shaped synthetic rubber component from Misumi, rigidly glued to end pieces that allow for a threaded connection to the coupler and nut. The rubber is either a synthetic standard urethane or nitrile rubber (NBR), which are not advertised to have a high elastic resilience, but do have a high gravimetric energy density. The rubber type with the highest elastic resilience should be selected for the final spring design, which is best identified through experimentation with a universal testing machine. The rubbers are furthermore available in different hardnesses, which allow for different spring shape designs. For example, the urethane rubber of Shore hardness A70, A90 and A95 have a 25% secant modulus $E_c = E_{25}$ of ca. 1 MPa, 2.2 MPa 5.3 MPa respectively, as extracted from the datasheet¹². Tentative spring dimensions can be estimated assuming a linear spring profile with a stiffness based on the secant modulus.

TENTATIVE SPRING DIMENSIONS

Assuming a linear stress-strain profile using the secant modulus, we can estimate the spring force and energy as follows:

$$F_s \approx \frac{E_c A x_s}{x_{s,0}} \quad \text{and} \quad E_s \approx \frac{E_c A x_s^2}{2x_{s,0}}$$

where A is the spring's cross-sectional surface area, x_s main spring elongation, $x_{s,0}$ the main spring's rest length and F_s the spring force. A spring that stores $E_s = 80$ J of energy at a spring force of $F_s = 2$ kN has a spring elongation of maximum $x_s = 8$ cm.¹³ If the spring is strained to 100% ($\epsilon = 1$) (which is well within bounds of the specified maximum

¹²Available here: https://uk.misumi-ec.com/pdf/fa/2014/P2_0389-0390_F38_EN.pdf

¹³This figure assumes a linear spring profile and is lower for regressive spring profiles.

strain, which exceeds 300% for all rubbers), then the spring has a rest length of $x_{s,0} = 8$ cm too. By solving for A using the shore A90 rubber, we find $A = 924 \text{ mm}^2$, which can be realized by a tube with an inner diameter of ca. 2 cm and an outer diameter of ca. 4 cm ($A = (20^2 - 10^2)\pi = 943 \text{ mm}^2$). At a density of $\rho = 1.2 \text{ kg dm}^{-3}$, the rubber weighs $Ax_{s,0}\rho = 90$ g, and each of the end-pieces weigh approximately 6 g each, resulting in a spring design that is substantially lighter than one based on GFRP springs.

The exact shape of the spring can be finalized based on optimisation studies, and can furthermore be changed by choosing a rubber of different hardness. A softer rubber results in a fatter but shorter spring, and a harder rubber in a longer but thinner spring. The presented design is considered to be good a trade-off between the two extremities. The spring length limits the minimum actuator length, as further discussed in the next section. A wider spring requires a bigger bumper on the lever to protect it from impacts and a longer spring further increases the minimum torso length, as further discussed in the next section.

Note furthermore that a thin spring has an increased maximum stress also on the glue that connects the rubber to the end pieces, because of the decreased surface area. For $F_s = 2 \text{ kN}$ and $A = 943 \text{ mm}^2$, the normal stress is approximately $\sigma = F_s/A = 2 \text{ MPa}$, which is still well below the glue's listed maximum bonding stress of at least 8 MPa for rubber¹⁴. However, caution should be taken for (chemical) deterioration of the rubber due to the application of adhesive, and for increased peak stress caused by shear and local imperfections at the bond, which could lead to propagating tear and subsequent early failure. Experiments on a small spring with a surface area of 25 mm^2 have shown failure at 73 N, corresponding to $\sigma = 2.9 \text{ MPa}$. The failure occurred in the rubber close to the adhesive rather than in the adhesive itself. Whereas this is within range for normal use, simple design measures could be taken if the bonding is not sufficiently strong, such as changing the shape of the bonding surface to being conical, increasing its surface area.

6.7.4. ACTUATOR LENGTH

The presented design of the main actuator imposes a constraint of minimum length on itself and thus on the torso. The actuator length (as measured from the coupler joint to the rocker joint, when Skippy is fully extended) is constrained by the sum of the length of the coupler frame, the main spring at rest, the guiding arrangement and the rocker:

- The length of the coupler is limited by the presence of the bumper, as can be seen in fig. 6.14, and adds 4.5 cm to the length of the actuator.
- The length of the main spring is defined by the sum of length of its rubber and end pieces (1 cm). The length of the rubber is assumed to be 8 cm, based on a spring elongation of approximately 4 to 5 cm ($\epsilon \approx 0.5$). It might be shortened depending on the final choice of rubber.
- The minimum required length of the guiding arrangement equals the sum of the lengths of the stroke, the minimum tube overlap (equal to the length of the nut guide bushings: 4 cm), and collision safety margins as required for (unintentional) spring elongation/compression (ca. 3 cm). The stroke is limited by the lever length

¹⁴Weicon Contact VA 5000 THIX: https://www.weicon.de/media/pdf/08/79/f2/TDS_12551030_EN_Contact_VA_5000_THIX_Cyanoacrylate_Adhesive.pdf

and its angular range. For a lever with a (maximum) length and range of 10 cm and 90° respectively, the stroke is maximum 14 cm.

- The rocker adds ca. 4 cm to the length of the actuator due to the presence of the coupling, bearings and lock nut.

These numbers add to an actuator length of ca. 38.5 cm, which—in combination with a lever of length 10 cm in Skippy's extended configuration—constrain the rocker joint to be at a distance of at least ca. 47.5 cm from the lever-torso joint.

SPINDLE LENGTH AND RESONANCE

The nut covers a stroke length of 134 mm for Skippy's full work space of $\Delta\phi_{\text{hip}} = 180^\circ$, assuming a lever length of 10 cm and the main spring to be at rest. The threaded length of the ball screw is 184 mm, and the nut has a length of 30 mm. This leaves a margin of $184 - 134 - 30 = 20$ mm. Tube lengths have been chosen so that there is a margin of threaded screw rod of 10 mm if Skippy is fully flexed or extended, assuming the spring to be at rest. This should be sufficient, given that the main spring should not be stretched or compressed when Skippy is fully flexed or extended.

The spindle is furthermore checked for resonance. The resonance frequency can be approximated by assuming the spindle to be a cylindrical beam, for which the resonance frequency of the first eigenmode is calculated as follows:

$$\omega_0 = K \sqrt{\frac{Er_0^2}{4\rho l_0^2}},$$

where K is a constant depending on beam support, E the Young's modulus, ρ the mass density, r_0 the inner radius and l_0 the unsupported length of the spindle. For our simply supported (i.e. not supported on one end) spindle, $K = 3.52$ and for steel $E = 200$ GPa and $\rho = 8\text{ kg dm}^{-3}$. The unsupported length is approximated by the distance between the nut's centre and its unsupported end, which is $l_0 = 16$ cm if Skippy is flexed, and the inner radius is $r_0 = 3.5$ mm. We find $\omega_0 = 1369\text{ rad s}^{-1} = 13\text{ krpm}$, which is well above the 6000 rpm speed limit.

6.8. SWIVEL ACTUATOR DESIGN

The swivel actuator is depicted in fig. 6.18. It comprises the encoded Maxon DCX22S motor, a Harmonic Drive HFUC-2A-11 strain wave gear, an RLS Aksim absolute encoder at the output shaft and a slip clutch.

As described in section 6.2, the strain wave gear has been upgraded one size compared to Tippy, to comply with higher torque at the output shaft and to provide a more rugged grub screw connection at the input shaft. The input hub now has an inner diameter of $\varnothing 5$ mm instead of $\varnothing 3$ mm, which requires the use of a shaft adapter.

In addition, the housing is custom-made to reduce weight, similar to that of Tippy's main actuator discussed in section 5.4.3. The custom housing allows for integration of the slip clutch, output encoder and strong output bearing. It is furthermore designed to be mounted to the inside of the frame of Skippy's head, which has inner dimensions of 48×48 mm.

SLIP CLUTCH

The slip clutch permits the *output shaft* (dark blue) to slip with respect to the *intermediate shaft* (output from the strain wave gear) (fuchsia), and consists of two friction rings (neon yellow) that are pre-loaded with 16×8.2 disc springs (Belleville washers) and a lock nut. The diameter of the slip clutch is ideally as large as possible, to minimize the required pre-load, but is limited by the inner diameter of $\varnothing 14$ mm of the encoder read head. The pre-load can be estimated using ballpark figures, as explained below, but is to be finalized through trial-and-error.

The pre-load force is calculated as $F_{pl} = \tau_{slip} / \mu / r_{slip}$, where τ_{slip} is the slipping torque, μ the static friction coefficient between the friction rings and the shaft (for PEEK-steel/aluminium $\mu \approx 0.5$) and $r_{slip} = 5$ mm the effective radius of the friction rings. For a supply voltage of 30 V and a transmission ratio and efficiency of 1:30 and $\eta = 0.7$ respectively, the actuator can deliver ca. 3.2 Nm stall torque, and approximately double this torque if rotating at no-load velocity in opposite direction. The strain wave gear can handle up to 8.5 Nm momentary peak torque. The slipping torque should be in between these values, e.g. $\tau_{slip} = 7.5$ Nm, for which we find $F_{pl} = 3$ kN. One $16 \times 8.2 \times 0.9$ -T disc springs (DIN2093) at 75% deflection provides ca. 1 kN load. Realizing $F_{pl} = 3$ kN would require the use of 3 of these springs preloaded at 75% deflection in parallel (note that fig. 6.18 depicts only two of them, mounted in series).

ABSOLUTE ENCODER

The RLS encoder is placed after the slip clutch, so that it measures the absolute position of the crossbar, but before the output bearing, so that the encoder is protected. The output bearing fixes the output shaft axially and radially to the housing, so as to fix the reading height of the encoder irrespective of wear of the slip rings. The reading height can be controlled by using shim rings on the actuator side of the output bearings (red). Holes in the frame permit encoder cables to be directly wired to the centre of the head.

BEARINGS AND ALIGNMENT

The output bearing is a 3805-2RS precision thin section double row angular contact ball bearing (25 g), which can handle the force if Skippy were to crash frontally on its head, and the moment or radial load if it were to crash on the tip of its crossbar. The bearing is fixed using a lightweight nylon lock nut (light brown), as typically used in piping industry. The bearing is axially stiff, but if nonetheless a crash would cause the inner ring to displace axially, it would be stopped by the locknut, which has an outer diameter slightly larger than that of the bearing. It is nominally spaced by the outer ring of the bearing by using a shim ring on the crossbar-side of the bearing (red). This arrangement protects both the bearing and encoder.

In addition, a 61703/6703-2RS precision thin section ball bearing (4 g) has been installed to radially align the intermediate shaft with the strain wave gear, and to prevent dust ingress from the slip clutch into the strain wave gear. This bearing is free to move axially, because axial alignment is already determined by the slip clutch and output bearing.




Figure 6.18: The swivel actuator, featuring Maxon's DCX22S brushed DC motor, a strain wave gear, encoder, slip clutch and a strong double row output bearing.

6.9. HEAD

Figure 6.19 shows Skippy's head, which is the heaviest part of the robot, containing most of the robot's features. It has the following main functions:

- Providing housing space for the rocker and main motor, the swivel actuator and most of the electronics, including the brain and batteries.
- Providing an environment for regulated heat flow and cooling of the main motor, drivers and swivel motor.
- Providing a method and easy access to change the main electronics and batteries.
- Providing protection for all of its contents, and an end stop for the leg.

The protection includes the frontal bumper on the swivel axle and the toroidal bumper, which have already been discussed in section 6.3. Note furthermore that the end stop is not visible in fig. 6.19, but it is in fig. 6.1 (next to the brain). The end stop uses the same sponge that is used in the bumpers, and also has a thickness of 3 cm.

HOUSING SPACE

The size and shape of the head is a compromise between accessibility and weight. Accessibility is generally achieved by spacious design. Whereas there is no direct volume budget on Skippy's design, a more spacious design was found to be heavier too.

The interior of the head frame features a lengthy *chamber* with an inner cross-section of 48×48 mm and a length of ca. 120 mm, which fits most of the electronics and swivel actuator. These dimensions have been determined based on (1) comparing it to the size of Tippy's torso interior, which at 38×38 mm was too small, (2) fitting the new swivel actuator and electronics (3) an iteration of sample prints and assembly (i.e. trial-and-error).¹⁵

¹⁵The iteration involved an earlier design featuring a slightly shorter length (ca. 100 mm) and smaller cross-




Figure 6.19: Skippy's head, with various points of interest indicated. Heat waves as caused by the presence of the motors, drivers and fan are drawn red. The approximate location of the torso's CoPc (⊙) that corresponds to landing from a vertical hop is also shown. The two possible mounting points for the IMU are indicated as well, complete with IMUs (in practice, only one mounting point has to be used).

The upper side of the head frame features a clevis, which houses the rocker, as discussed in section 6.7. Ideally, the rocker is located as close as possible to the torso tube to reduce the bending moment it exerts on the tube. In the final design, its hinge is placed 47.5 mm above the line defined by the torso tube. The rocker rocks in a range of ca. 7.5° for a lever length of $l_l = 10\text{ cm}$, so that the back side of the motor displaces over a span of ca. 12 mm. This space is made available by aligning the upper side of the chamber with the upper side of the torso tube, rather than designing for them to be

section ($46 \times 46\text{ mm}$), which was printed and successfully fitted with the electronics, but of which was concluded that it was slightly too cramped for easy assembly and disassembly.

coaxial.

The design style of the head frame is also closed-section design. Its interior can be milled from the front side, where the swivel actuator is to be mounted, and partly from above. The interior features inner rounds of 10 mm or higher to facilitate its machining and reduce machining time.

COOLING

Cooling of the main motor is realized by clamping an actively cooled heat sink on the motor's housing. The active cooling is done by a fan that is attached to the heat sink.

The swivel motor and drivers might also require some cooling, depending on their use. For this reason, both the swivel motor and drivers are placed in the same cooling environment, below the fan. To guide the air flow, the head frame has two air inlets/outlets at the bottom, next to the drivers. The maximum winding temperature of the main motor is rated at 155 °C, whereas that of the swivel motor is rated at 125 °C. Nevertheless, their maximum ambient temperature is rated at 100 °C. The maximum driver temperature is limited by its ceramic capacitors at 85 °C, which is slightly lower than the permitted ambient temperature of the motors. This is likely not an issue, since the ambient temperature should not be permitted to increase beyond 85 °C anyway. However, even if it would, then either of two measures can be taken.

- The fan can be configured so that it blows air from the bottom to the top, so that the temperature of the air below the fan is lower than that above the main motor.
- The temperature limit of the drivers can be increased to 105 °C by replacing the capacitors for more thermal resistant capacitors, as suggested by the Pololu company.

ELECTRONICS

Skippy's head houses most of its electronics, which are mostly the same as those of Tippy. Several revisions have been done by reflecting on Tippy section 5.7, which mainly concern issues with cabling.

Other than the two motors and their encoders (2x ENX EASY 16 relative encoders and an RLS Aksim II absolute encoder), the head houses the following electronics:

- Two Pololu G2 24v21 High-Power motor drivers, one for each motor (identical to Tippy), placed on both sides next to the swivel motor.
- A Pololu 5V step-down voltage regulator, mounted on one side. The other side offers an extra space for a miscellaneous PCB of similar size, such as for battery wiring or monitoring.
- Two mounting options for an IMU (indicated with I and II in fig. 6.19). Placement of the IMU is further explained in section 6.10.
- A *junction PCB*, to which all electronics are wired through solder connections. The PCB has a 60-pin plug sticking out of the lower side of the head frame, to which Skippy's brain is connected.

- The brain¹⁶ is plugged into the junction PCB, and screwed on the exterior of the bottom of the head frame. The brain consists of three layers of PCB's. The two PCBs closest to the body are custom. The first PCB takes care of wiring and basic electronic circuitry, similar to that found in TP-PCB. The second PCB is essentially a downscaled version of the TI Launchpad F28377S, featuring an extra SPI port and no more breakout pins. The third PCB is the Raspberry PI Zero.
- The batteries are placed on the exterior on both sides of the head frame, and reside in a 3D printed housing to which they are strapped with hook-and-loop fastener. The batteries have a capacity of 1000 mAh, in contrast to the 1300 mAh used on Tippy, with an identical 30C/60C peak rating, implying they can deliver up to a current of 60 A, which is still enough. Smaller batteries have been chosen to save weight.

Note that the brain and the batteries are the only electronics mounted to the exterior of the head frame, because they have to be easily replaceable. The electronics differ from Tippy in the following ways:

- Skippy no longer has a separate 3.3 V step-down voltage regulator (required to power the MCU), as it is integrated in the first PCB of the brain.
- The IMU has been designed with a custom PCB and housing so that its SPI is accessible and to decrease its size. The housing is 3D printed to reduce costs.
- The third SPI peripheral of the MCU is now accessible, as required for using the IMU over its SPI.
- Multiplexers have been added for reasons of expansibility, enabling connectivity of additional peripherals that make use of SPI. The multiplexers are integrated in the first PCB of the brain.
- The size of the brain has been reduced, as realized by removal of unused interfaces and (bulky) breakout pins.
- All cable connections that do not require (frequent) removal are replaced for rugged solder connections. Connections that require more frequent removal are equipped with more rugged connectors (no more break-out pins).

6.10. IMU PLACEMENT

The IMU is preferably placed in the robot where peak accelerations are minimal, so as to avoid saturation of the IMU readings. For this purpose, it is common in robot design to place the IMU on a heavy body close to a robot's CoM. Generally, shock propagation is lower in heavy bodies that are central to the robot (i.e. far from end effectors). In addition, accelerations as caused by rotational motions are typically lower close to the CoM. However, if a robot's typical behaviour is governed by a particular and extreme impact, then more tactical placement of the IMU is at a body's *centre of percussion* (CoPc). Unlike the CoM, which is only a function of a body's inertial properties, the CoPc

¹⁶*Acknowledgement* — The development of the brain including the junction PCB has been a joint effort between the author, Antonios E. Gkikakis, B. Roodra P. Singh and Phil E. Hudson, the latter of who has implemented its electronic design in electronic CAD software.

is specific to the impact conditions. For a legged robot, the most extreme impact results from landing during its running or hopping gait, and is typically well described.

6.10.1. CENTRE OF PERCUSSION DEFINITION

The CoPc is most commonly calculated for long rigid bodies (i.e. beams) that have a *single* well-defined pivot point (i.e. revolute joint), such as a bat, racket or sword (in which the pivot points are their respective handles). The CoPc is then defined as follows:

The CoPc is the point where an impact perpendicular to the body's extension produces no reactive impulse at the pivot point.

This implies that no instantaneous change of translational acceleration and thus velocity would occur at the pivot point, even if the pivot point is unconstrained. For a body of mass m and moment of inertia I about the CoM, it can be derived that

$$c_x(p_x - c_x) = \frac{I}{m} = r^2 \quad (6.1)$$

where c_x and p_x are the (1D) coordinates of the body's CoM and CoPc respectively with respect to the pivot point, and r is its radius of gyration. Referring to the slender beam depicted in fig. 6.20 of length l with a uniform mass distribution so that $I = 1/12ml^2$ and pivoted at one of its ends so that $c_x = 1/2l$, we find $p_x = 2/3l$.

The fact that the inertias I and m are always positive implies that also $c_x(p_x - c_x)$ is positive, which implies that the pivot point and the CoPc are always on opposite sides of the CoM (i.e. if $c_x > 0$ then $p_x > c_x$). Note furthermore that the appearance of the product $c_x(p_x - c_x)$ in eq. (6.1) implies reversibility: if instead an impulse is applied at the (unconstrained) pivot point, then no reactive shock is produced at the CoPc. This concept follows from conservation of momentum, which enables us to extend the definition of the CoPc for a rigid body of any shape with any number of joints in the following way:¹⁷

For a given impulse on (e.g. the joints of) a rigid body, the CoPc is the point at which no reactive instantaneous change of translational velocity occurs.

Note that in this definition the CoPc is a function of the applied impulse, whereas in the former definition it was not because the impulsive conditions were included in the definition. In 2D this point is always defined except if the applied impulse results in a pure translation (which would correspond to an impact being applied at the CoM). In 3D this point is generally not defined, as a rigid body does not include a point or set of points with zero displacement in the case of the most general rigid body displacement. This follows directly from Chasles' theorem, which says that the most general rigid body displacement is described by a translation along a line – the screw axis – in combination with a rotation about an axis parallel to that line. The displacement is minimal along the screw axis, and zero only along the axis for the special case of a pure rotation.

¹⁷Despite efforts of the author in searching for such a generalized definition of the CoPc in literature, none have been found.




Figure 6.20: A slender beam suspended from a pivot on a frictionless slider, with the CoM (●) at $c_x = 1/2l$ and the CoPc (⊗) at $p_x = 2/3l$ from the pivot point. The initial response is shown in grey: if struck by an impact at the CoPc (left), then no initial displacement occurs at the pivot, and the other way round (right).

6.10.2. CENTRE OF PERCUSSION DERIVATION

The CoPc of a rigid body is a function of the applied impulse. The impulse results in a specific instantaneous change of spatial velocity $\Delta \hat{\mathbf{v}}$ of the rigid body, based on which the CoPc can be calculated. Based on conventions used by Featherstone (2008), we can express $\hat{\mathbf{v}}$ as a Plücker coordinate vector (and thus also $\Delta \hat{\mathbf{v}}$) so that it describes the velocity in coordinates of the frame O in the rigid body with respect to the global reference frame (i.e. the global velocity in body coordinates):

$$\hat{\mathbf{v}}_O = \begin{bmatrix} \boldsymbol{\omega} \\ \mathbf{v}_O \end{bmatrix} = \begin{bmatrix} \omega_x \\ \omega_y \\ \omega_z \\ v_{Ox} \\ v_{Oy} \\ v_{Oz} \end{bmatrix} \Rightarrow \Delta \hat{\mathbf{v}}_O = \begin{bmatrix} \Delta \boldsymbol{\omega} \\ \Delta \mathbf{v}_O \end{bmatrix} = \begin{bmatrix} \Delta \omega_x \\ \Delta \omega_y \\ \Delta \omega_z \\ \Delta v_{Ox} \\ \Delta v_{Oy} \\ \Delta v_{Oz} \end{bmatrix},$$

where $\Delta \boldsymbol{\omega} = [\Delta \omega_x \ \Delta \omega_y \ \Delta \omega_z]^\top$ describes the change of angular velocity of the rigid body about the x , y and z axes of coordinate frame O , and $\Delta \mathbf{v}_O = [\Delta v_{Ox} \ \Delta v_{Oy} \ \Delta v_{Oz}]^\top$ the change of translational velocities along those axes, both with respect to the global reference frame. O can be any frame in the rigid body, but it will be used to describe the frame that coincides with the joint that connects the rigid body to its parental body, following the conventions for describing rigid body kinematics in Featherstone (2008).

The change of translational velocity $\Delta \mathbf{v}_P$ of the point P in the rigid body can be calculated as

$$\Delta \mathbf{v}_P = \Delta \mathbf{v}_O - \overline{OP} \times \Delta \boldsymbol{\omega}, \quad (6.2)$$

where $\overline{OP} = [\overline{OP}_x \ \overline{OP}_y \ \overline{OP}_z]^\top$ is the vector describing the translation from O to P in coordinates of O . If $\Delta \mathbf{v}_P = \mathbf{0}$ then P defines the CoPc: $\mathbf{p} = \overline{OP}$. However, solving eq. (6.2) for \overline{OP} does not generally have a solution in 3D, as explained previously (Chasles' theorem). In 2D we have $\Delta v_{Pz} = \Delta v_{Oz} = \Delta \omega_x = \Delta \omega_y = 0$ (assumes the xy -plane) so that

eq. (6.2) simplifies to

$$\begin{bmatrix} \Delta v_{p_x} \\ \Delta v_{p_y} \end{bmatrix} = \begin{bmatrix} \Delta v_{O_x} \\ \Delta v_{O_y} \end{bmatrix} - \begin{bmatrix} \overline{OP}_y \Delta \omega_z \\ -\overline{OP}_x \Delta \omega_z \end{bmatrix}.$$

For $\Delta \mathbf{v}_p = \mathbf{0}$ we find

$$\begin{bmatrix} p_x \\ p_y \end{bmatrix} = \begin{bmatrix} \overline{OP}_x \\ \overline{OP}_y \end{bmatrix} = \frac{1}{\Delta \omega_z} \begin{bmatrix} -\Delta v_{O_y} \\ \Delta v_{O_x} \end{bmatrix} = \frac{1}{d\omega_z} \begin{bmatrix} -dv_{O_y} \\ dv_{O_x} \end{bmatrix}, \quad (6.3)$$

which is defined as long as $\Delta \omega_z \neq 0$ (non-pure translational displacement). Note from eq. (6.3) that only the ratio of $\Delta \mathbf{v}_O$ matters, not the magnitude.

A method to calculate $\Delta \hat{\mathbf{v}}_O$ of a body in a multi-rigid-body system as the result of an impulse is explained below. We first calculate the change of joint velocity variables $\Delta \hat{\mathbf{q}}$ as the result of an impulse and then calculate $\Delta \hat{\mathbf{v}}_O$ from $\Delta \hat{\mathbf{q}}$.

IMPULSIVE EQUATION OF MOTION

The EoM of a multi-rigid-body system can be written in the canonical form $\mathbf{H}\ddot{\mathbf{q}} + \mathbf{C} = \boldsymbol{\tau}$, where \mathbf{q} are the joint position variables and $\boldsymbol{\tau}$ the corresponding joint forces, \mathbf{H} the joint-space inertia matrix and \mathbf{C} the joint-space vector of force terms. For an open-loop system, the impulsive equation of motion amended for inelastic impulsive contact can be expressed as

$$\begin{bmatrix} \mathbf{H} & \mathbf{J}(\mathbf{D}_c, \mathbf{q}) \\ \mathbf{J}(\mathbf{D}_c, \mathbf{q})^\top & \mathbf{0} \end{bmatrix} \begin{bmatrix} \dot{\mathbf{q}}^+ \\ \boldsymbol{\rho}_c \end{bmatrix} = \begin{bmatrix} \mathbf{H}\dot{\mathbf{q}}^- + \boldsymbol{\iota} \\ \mathbf{0} \end{bmatrix}. \quad (6.4)$$

$\dot{\mathbf{q}}^-$ and $\dot{\mathbf{q}}^+$ are the joint velocities before and after impact respectively, so that $\Delta \dot{\mathbf{q}} = \dot{\mathbf{q}}^+ - \dot{\mathbf{q}}^-$. \mathbf{D}_c is the impulsive contact condition, as function of \mathbf{q} , so that $\mathbf{D}_c(\mathbf{q}) = \mathbf{0}$ at contact. $\mathbf{J}(\mathbf{D}_c, \mathbf{q})$ is the Jacobian of \mathbf{D}_c to \mathbf{q} . $\boldsymbol{\iota}$ are the applied impulses at the joints (zero if the impulse is only caused by impulsive contact). The length of \mathbf{D}_c equals the number of contact points. $\boldsymbol{\rho}_c$ are Lagrange multipliers so that $\boldsymbol{\iota}_c = \mathbf{J}(\mathbf{D}_c, \mathbf{q})\boldsymbol{\rho}_c$ are the contact impulses.

$\Delta \dot{\mathbf{q}}$ can be found by solving for $[\dot{\mathbf{q}}^+ \ \boldsymbol{\rho}_c]^\top$ in eq. (6.4) for specified impulsive contact conditions. Note that if the impulsive contact occurs directly at one or multiple of the joints, then eq. (6.4) simplifies to

$$\begin{bmatrix} \mathbf{H}_{11} & \mathbf{H}_{12} & \mathbf{1} \\ \mathbf{H}_{21} & \mathbf{H}_{22} & \mathbf{0} \\ \mathbf{1} & \mathbf{0} & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{0} \\ \dot{\mathbf{q}}_2^+ \\ \boldsymbol{\iota}_c \end{bmatrix} = \begin{bmatrix} \mathbf{H}_{12} \\ \mathbf{H}_{22} \end{bmatrix} \dot{\mathbf{q}}_2^+ + \begin{bmatrix} \boldsymbol{\iota}_c \\ \mathbf{0} \end{bmatrix} = \begin{bmatrix} \mathbf{H}_{11} & \mathbf{H}_{12} \\ \mathbf{H}_{21} & \mathbf{H}_{22} \end{bmatrix} \begin{bmatrix} \dot{\mathbf{q}}_1^- \\ \dot{\mathbf{q}}_2^- \end{bmatrix} + \begin{bmatrix} \boldsymbol{\iota}_1 \\ \boldsymbol{\iota}_2 \end{bmatrix} = \mathbf{H}\dot{\mathbf{q}}^- + \boldsymbol{\iota}, \quad (6.5)$$

in which the subscripts $_1$ and $_2$ are defined to distinguish between joints at which impulsive contact occurs and at which not, respectively, so that $\dot{\mathbf{q}}_1^+ = \mathbf{0}$. In this case it follows that $\Delta \dot{\mathbf{q}} = \dot{\mathbf{q}}^+ - \dot{\mathbf{q}}^-$ can be directly calculated as

$$\Delta \dot{\mathbf{q}} = \begin{bmatrix} \Delta \dot{\mathbf{q}}_1 \\ \Delta \dot{\mathbf{q}}_2 \end{bmatrix} = \begin{bmatrix} -\dot{\mathbf{q}}_1^- \\ \mathbf{H}_{22}^{-1}(\mathbf{H}_{21}\dot{\mathbf{q}}_1^-) \end{bmatrix}. \quad (6.6)$$

Closed-loop systems Remember that eqs. (6.4–6.6) only apply to open-loop systems (i.e. for systems where $\dot{\mathbf{q}}$ are independent joint velocity variables). For closed-loop systems, we have to take reaction impulses into account that result from joint constraints.

This can be done by replacing the joint velocity variables $\dot{\mathbf{q}}$ for *independent* joint velocity variables $\dot{\mathbf{y}}$ and \mathbf{H} for $\dot{\mathbf{H}}$ in eq. (6.4), so that $\dot{\mathbf{q}} = \mathbf{G}\dot{\mathbf{y}}$ and $\dot{\mathbf{H}} = \mathbf{G}^T \mathbf{H} \mathbf{G}$, where \mathbf{G} is the velocity mapping or Jacobian $\mathbf{G} = \mathbf{J}(\mathbf{q}, \mathbf{y})$. Substitutions in eqs. (6.5) and (6.6) are similar, i.e. we substitute $\dot{\mathbf{q}} = [\dot{q}_1 \ \dot{q}_2]^T$ for $\dot{\mathbf{y}} = [\dot{y}_1 \ \dot{y}_2]^T$ and make corresponding substitutions for \mathbf{H} . Equation (6.6) becomes

$$\Delta \dot{\mathbf{y}} = \begin{bmatrix} \Delta \dot{y}_1 \\ \Delta \dot{y}_2 \end{bmatrix} = \begin{bmatrix} -\dot{y}_1^- \\ \dot{\mathbf{H}}_{22}^{-1} (\dot{\mathbf{H}}_{21} \dot{y}_1^-) \end{bmatrix}. \quad (6.7)$$

$\Delta \dot{\mathbf{q}}$ can then be calculated using \mathbf{G} :

$$\Delta \dot{\mathbf{q}} = \mathbf{G} \Delta \dot{\mathbf{y}}. \quad (6.8)$$

VELOCITY IN BODY COORDINATES

The change of global velocity in body coordinates of a body $\Delta \hat{\mathbf{v}}$ can be calculated from $\Delta \dot{\mathbf{q}}$ and coordinate transforms that define the kinematics of the multi-rigid-body system. Following the conventions of Featherstone (2008), body velocities can be calculated recursively according to

$$\hat{\mathbf{v}}_i = {}^i \mathbf{X}_{\lambda(i)} \hat{\mathbf{v}}_{\lambda(i)} + \mathbf{S}_i \dot{q}_i, \quad (6.9)$$

where $\hat{\mathbf{v}}_i$ is the global velocity of body i in body coordinates of the frame that coincides with the 1-DoF joint to which body i is connected to its parent. The corresponding joint velocity of the 1-DoF joint of body i is \dot{q}_i , which is part of the tuple $\dot{\mathbf{q}} = [\dot{q}_1 \cdots \dot{q}_i \cdots \dot{q}_n]^T$. λ is the parent array of the kinematic tree, so that $\lambda(i)$ is the parent of body i ; $\lambda(i) = i - 1$ if the kinematic tree is unbranched. $\hat{\mathbf{v}}_{\lambda(i)}$ is the parent body velocity, and $\hat{\mathbf{v}}_0 = \mathbf{0}$. \mathbf{S}_i is the joint axis vector, which is a function of joint type (e.g. revolute, prismatic). ${}^i \mathbf{X}_{\lambda(i)}$ is the coordinate transform that describes the position and orientation of body coordinates of body i relative to those of $\lambda(i)$. It is calculated as the product of the joint transform $\mathbf{X}_{J_i}(q_i)$ and the tree transform \mathbf{X}_{T_i} , i.e. ${}^i \mathbf{X}_{\lambda(i)} = \mathbf{X}_{J_i} \mathbf{X}_{T_i}$. \mathbf{X}_{T_i} is a function of the displacement and rotation between coordinate frames of bodies $\lambda(i)$ and i for $q_i = 0$, whereas \mathbf{X}_{J_i} is a function of joint type and position q_i . For descriptions and templates (for various joint types) of \mathbf{X}_J , \mathbf{X}_T and \mathbf{S} refer to Featherstone (2008).

Since $\hat{\mathbf{v}}_i$ is a vector, and thus eq. (6.9) a vector equation, we can write

$$\Delta \hat{\mathbf{v}}_i = {}^i \mathbf{X}_{\lambda(i)} \Delta \hat{\mathbf{v}}_{\lambda(i)} + \mathbf{S}_i \Delta \dot{q}_i. \quad (6.10)$$

Having obtained $\Delta \hat{\mathbf{v}}_i$, the CoPc for body i can be directly calculated using eq. (6.3) for $\hat{\mathbf{v}}_O = \hat{\mathbf{v}}_i$.

IMU PLACEMENT

The IMU is placed on the torso. By using the dynamic model of Skippy described in section 8.2, we calculate the CoPc of its torso for a landing from a vertical jump, by assuming initial conditions similar to that of the landing of the stance phase in preparation for a high jump. Recall furthermore that Skippy is a closed-loop system.

The initial position variables are set so that Skippy's initial hip angle $\phi_{\text{hip}} = q_6(0) + q_7(0) \approx 38^\circ$ and the foot angle $q_4(0)$ so that Skippy's CoM is approximately above its support point. The the independent joint velocity before impact is set to

$\dot{\mathbf{y}}^- = [0 \ -1 \ 0 \ 0 \ 0]^T$, to emulate a landing from a vertical jump, corresponding to only a non-zero vertical velocity and zero joint velocities. Note that the magnitude of $\dot{\mathbf{q}}(0)$ is irrelevant for the calculation of the CoPc, implying that the CoPc is identical for falls from different heights. Impulsive contact occurs directly at joint 1 and 2, so that we can use eq. (6.7) for $\dot{\mathbf{y}}_1^- = [\dot{y}_1^- \ \dot{y}_2^-]^T$ and $\dot{\mathbf{y}}_2^- = [\dot{y}_3^- \ \dot{y}_4^- \ \dot{y}_5^-]^T$ to calculate $\Delta\dot{\mathbf{y}}$. We then proceed to use eq. (6.8), eq. (6.10) and eq. (6.3) successively to calculate the CoPc. We find $\mathbf{p} = [p_x \ p_y]^T = [0.50 \ 0.14]^T$ m for the torso ($i = 7$), which is approximately 8 cm above the heat sink of the main motor (see fig. 6.19). We furthermore find that changing initial conditions for a vertical landing within limits imposed by optimisations do not significantly change \mathbf{p} , especially not p_x . For example, for a vertical landing from a somersault, for which we set (per estimate) $\dot{\mathbf{y}}(0) = [0 \ -8 \ 0 \ 4 \ 0]^T$, we find $\mathbf{p} = [0.52 \ 0.25]^T$ m.

The IMU is to be placed close to the CoPc to reduce the shock it receives at landing. The performed calculations show that the CoPc is close to Skippy's head, which suggests that the IMU be also placed in this region. More precisely, the CoPc is calculated to be close to and above the main motor and its heat sink. This is not a practical location for the IMU because of (1) its potentially hot environment (the IMU is sensitive to change of temperature), (2) the magnetic field of the motor (which could disturb the IMU's magnetic sensor) and (3) the lack of structural frame and bumpers (to mount and protect the IMU from crashes). The closest mountable position is that below the rocker (position I), as shown in fig. 6.19, which is out of the heat stream of the cooling unit, and still only ca. 1 to 2 dm away from the CoPc.

The position second closest to the CoPc (a few centimetres further away) is that inside the torso tube (position II), as also shown in fig. 6.19. This position is further away from the heat stream, and furthermore allows the IMU to be mounted under an angle of 45°. The latter increases the saturation limit of vertical acceleration by up to a factor of $\sqrt{3/2}$ to $\sqrt{2}$, which is a consequence of the way saturation limits in the IMU are built-in. The translational saturation limit of 16g applies to each of its axes (x , y and z) separately, which implies that the acceleration is limited by a cube rather than a sphere in geometrical terms, so that its orientation matters. Note that this effectively makes placement of the IMU at position II more robust to impact than placement at position I, because it is located less than a factor of $\sqrt{3/2}$ away from the CoPc. Nevertheless, position II is more difficult to reach and assemble, partly due to the presence of cables from the foot and hip crossing the torso. For this reasons, both mounting positions of the IMU have been kept available in the final design.

6.11. CONCLUSION

This chapter has presented Skippy's mechanical design in detail, including its actuators, springs, hip (four-bar linkage), and head. The design comprises various novelties and unconventional design choices, which have been the result of reflecting on previous design studies and a desire for lightweight, physically robust and affordable design. They include (1) a structural design comprising CFRP tubes and closed-section frame parts, (2) helmet-inspired bumpers and a tape-measure buckling system for physical robustness, and (3) light-weight spring design of a rubber tensile spring and a GFRP curved compression leaf spring. To warrant or support the design's feasibility, various of the part's functionality, survivability and manufacturability have been tested in isolation.

Design decisions of which feasibility is less certain have been accompanied by suggestions for easy alterations, which is in conformity with an educated-risk strategy that supplements the trial-and-error design philosophy as presented in chapter 2. For example, this includes the possibility of easily changing the axial alignment of the dead weights on the crossbar blades, of replacing the main and ankle spring to obtain a different spring stiffness and of altering the main actuator design to include an end-stop for obtaining a more regressive main spring, if so desired. The design has also resulted in new theoretical contributions, such as modelling of curved compression leaf springs and revising the centre of percussion for placement of components that are sensitive to shock, the former of which is further treated in chapter 7. Having obtained accurate inertial parameters and physical constraints with the new CAD model, few remaining parameters are finalized in optimisations which are further treated in chapter 8.

7

MODELLING OF CURVED COMPRESSION LEAF SPRINGS

This chapter presents a plug-and-play formula for curved compression leaf springs, based on a simple 1-DoF pseudo-rigid body model. The spring profile covers a great range of regressiveness. Comparisons with simulation studies show good correspondence with the spring profile for large deflections. In addition, the theory can be used to predict failure of the spring. The model is used to describe Skippy's ankle spring for simulation and optimisation.

7.1. INTRODUCTION

The work in this chapter was inspired from the search for a simple analytic representation of Skippy's ankle spring, which can be described as a curved compression leaf spring, to be usable in simulation and optimisation. The spring has a regressive spring profile, which is a popular and promising choice for various robotic applications, and seem to be advantageous for tasks such as jumping and hopping (Driessen, 2015; Seyfarth et al., 2006). In comparison to ordinary linear springs, they allow for up to twice as much energy storage per unit of elongation for identical peak force, and they offer more bandwidth for delicate control tasks that require little force. For the particular application in Skippy, the regressive spring enables an increase of the stroke shortly before take-off that can be used to steer the robot (i.e. to control its take-off momentum).

To obtain a regressive spring profile, one either has to exploit material properties or geometry. This research focuses on the latter, as it describes the regressive behaviour that results from the geometry of the curved compression leaf spring, of which the ordinary 'buckling column' is a special case.

BUCKLING COLUMN

An extreme but yet simple example of a regressive spring is the buckling column, which remains stiff until it buckles. Euler's critical load theory can be used to calculate the




Figure 7.1: Compression leaf spring at rest ($x = 0$, $F = 0$), with left the continuous representation and right the corresponding 3-link pseudo-rigid body model (PRBM).

load at which buckling occurs, but it cannot be used to obtain the spring profile after buckling.

7

Numerical non-linear finite element analyses offer accurate solutions, but are typically employed for case-by-case studies that lead to look-up tables, which are usually not generalizable to different shapes. A more practical solution is that of algebraic approximations, which can be obtained using pseudo-rigid body models (PRBM). These models divide the flexible body in multiple rigid bodies that are interconnected with linear springs. The springs are rotational springs, axial springs or both. Some studies have been done on generalizing these models (Venkiteswaran, 2017), whereas others are dedicated to specific structures, such as the previously mentioned buckling columns (Boyle, Howell, Magleby, & Evans, 2003).

However, the buckling column is not always practical to be used as spring, because (1) its behaviour around the critical load is very unpredictable due to high sensitivity to the direction of the applied load and imperfections of the shape, and (2) the spring profile is only C^0 -continuous at the point of buckling.

CURVED COMPRESSION LEAF SPRING

The buckling column can be generalized to a compression spring with curvature, which offers a great variety of regressiveness. Figure 7.1 (left) shows such a compression spring with rest length L and rest arc width a . A spring of this type has seen previous usage in robotics in for instance a bowleg hopper (Zeglin & Brown, 2002) and a climbing robot (Degani et al., 2011), and it is to be used in Skippy's ankle as described in section 6.6.1. If $a = 0$, the spring represents a buckling column. As a/L is higher the profile is less




Figure 7.2: Half a compression leaf spring.

regressive and less sensitive to imperfections in its shape and load. The PRBM shown in fig. 7.1 (right) is used to derive the spring profile as function of its shape.

Many studies on PRBMs focus on deriving dynamics (frequency response) (Akano & Fakinlede, 2013; Boyle et al., 2003; Y.-Q. Yu, Li, & Xu, 2018) or generalize the use of PRBMs (Venkiteswaran, 2017). This study focuses on only statics, but is novel in that it focuses specifically on accurately describing the spring profile (1) of a generic curved compression leaf spring (2) with a single algebraic formula (3) that is accurate for large deflections reaching 50% of the spring length ($L/2$). In addition, (4) a formula is presented to predict failure of the spring. In particular, the algebraic formulae for the spring profile and its failure prediction are useful for optimisation studies, such as those of Skippy as presented in chapter 8.

7

7.2. DERIVATION

The PRBM consists of three rigid bodies and two linear torsion springs that interconnect the three links. The torsion springs are defined to have identical stiffness, and the top and bottom link have the same length. The mechanism is symmetric, which allows us to further simplify the system to a 1R PRBM having only two links and one spring, of which the second link is always aligned vertically, as shown in fig. 7.2.

7.2.1. KINEMATICS

To preserve the shape of the original spring, the link joints are constrained to coincide with the arc that describes the shape of the original spring.¹ The arc is defined by lengths

¹Note that this implies that the sum of link lengths ($2l_1 + l_2$) does not equal the arc perimeter. Alternatively, one could define the sum of link lengths to equal the arc perimeter, but this does a poorer job on preserving the shape for low radii and, correspondingly, was found to deteriorate results.

L and a , but can also be defined by its radius R and spanning angle Θ :

$$R = \frac{L^2 + 4a^2}{8a} \approx \frac{L^2}{8a} \quad \text{if } L \gg a \quad (7.1)$$

$$\Theta = 2 \arcsin \frac{L}{2R} \quad (7.2)$$

The arc perimeter P is then calculated as

$$P = R\Theta \quad (7.3)$$

The link lengths l_1 and l_2 are furthermore determined by the relation:

$$\lambda = \frac{p_1}{P/2} = \frac{P - p_2}{P} = \frac{\theta_1}{\Theta/2} = \frac{\Theta - \theta_2}{\Theta} \quad (7.4)$$

where λ is a characteristic radius and $p_1 = R\theta_1$ and $p_2 = R\theta_2$ are the perimeters of the parts of the arc that cover the first and second link respectively, see fig. 7.2. The characteristic radius is typically defined as a ratio between lengths, not angles or perimeters (Boyle et al., 2003). The definition used here was found to give more accurate solutions for this particular shape of spring for a wide variety of spring curvatures (i.e. different angles Θ). Note that for $\Theta \rightarrow 0$, λ is also the ratio of beam lengths (i.e. $\lambda \approx l_1/(l_1 + l_2/2)$).

We can calculate l_2 as follows:

$$\frac{l_2}{2} = R \sin \frac{\theta_2}{2} = R \sin \left(\frac{p_2}{2P} \Theta \right) = R \sin \left(\frac{(1-\lambda)}{2} \Theta \right) \quad (7.5)$$

Subsequently, for l_1 we find

$$l_1 = \sqrt{l_{1x}^2 + l_{1y}^2}, \quad \text{with} \quad (7.6)$$

$$l_{1x} = \sqrt{R^2 - \left(\frac{l_2}{2} \right)^2} - (R - a) \quad \text{and}$$

$$l_{1y} = \frac{L - l_2}{2}.$$

7.2.2. STATICS

The bending moment M exerted by the torsion spring at the joint that connects both rigid bodies equals

$$M = k(\phi - \phi_0) \quad (7.7)$$

where k is the spring constant, ϕ the joint rotation with respect to the vertical and ϕ_0 the joint rest angle. ϕ_0 follows from previously derived kinematic relations:

$$\phi_0 = \arctan \frac{l_{1x}}{l_{1y}} \quad (7.8)$$

Displacement y as function of ϕ equals

$$\frac{y}{2} = l_{1y} - l_1 \cos \phi \quad (7.9)$$

From the principle of virtual work ($F dy/2 = M d\phi$) it follows that the force F equals

$$F(\phi) = \frac{2 d\phi}{dy} M = \frac{k(\phi - \phi_0)}{l_1 \sin \phi} \quad (7.10)$$

If we then express F directly as function of y we find:

$$F(y) = \frac{k \left(\arccos \left(\frac{l_1 y - y/2}{l_1} \right) - \phi_0 \right)}{\sqrt{l_1^2 - (l_1 y - y/2)^2}} \quad (7.11)$$

The stiffness k can be approximated using the Euler-Bernoulli beam theory:

$$k = \lambda k_c \frac{EI}{L_c} \quad (7.12)$$

with E the material's Young's modulus, I the beam's cross-sectional second moment of area, L_c the characteristic length and k_c the stiffness coefficient.

For an Euler buckling column $k_c = \pi^2/4 \approx 2.47$ and for a cantilever beam $k_c = 3$. Since our system very closely represents that of the Euler buckling column, we use the former $k_c = \pi^2/4$, which should be accurate especially for initial deflection.

The characteristic length L_c represents the length over which bending occurs, which equals $L_c = P/2$ for the original flexible body, whereas it would equal $L_c = l_1 + l_2/2$ for the PRBM. Since we are aiming to represent the flexible body, the former definition leads to slightly more accurate results, even when applied to the PRBM.

The characteristic radius factor λ is required to effectively place the spring in the joint rather than in the centre of the body, such that the moment in the joint felt by F is only caused by the arm of the first body. A possibly even more sensible definition of λ in eq. (7.12) is $\lambda = l_1/(l_1 + l_2/2)$ or $\lambda = 2l_1/P$, but for now we will stick with the definition in eq. (7.4).

7.3. VALIDATION

Validations of the stiffness profile defined by eq. (7.11) have been performed by comparing it with that generated by non-linear FEAs. The FEAs are done using Siemens' NX Nastran 11. The model used for analysis is the symmetrical model displayed in fig. 7.2, with 24 unidirectional elements of each 3 mm long. A load of $F = 300$ N is applied and the static displacement is calculated over at least 50 iterations of re-calculating the stiffness matrix.

The shape of the spring is defined by the length $L = 144$ mm and various arc widths a ranging from $a = 0.1$ (practically straight) to $a = 16$ mm. The cross-section is rectangular with a width $w = 24$ mm and thickness $t = 1.5$ mm, which has a second moment of area of $I = \frac{t^3 w}{12} = 6.75 \text{ mm}^4$. For the Young's modulus we choose $E = 35$ GPa, to represent a unidirectional glass-fibre reinforced plastic (GFRP), as used in Skippy's ankle spring.

For the algebraic model, this leaves only λ undefined. λ mainly affects the stiffness profile after initial deformation and is ideally (like k_c) a function of displacement. We chose $\lambda = 0.6$ empirically, to fit well with FEA results for the range of 0% to 50%




Figure 7.3: Spring profile: force F versus displacement y according to FEAs (solid) and algebra (dashed) for various arc widths a (see legend). Parameters: $k_c = \pi^2/4$, $\lambda = 0.6$, $L = 144$ mm, $I = 6.75$ mm⁴, $E = 35$ GPa. Black x-marks correspond to a maximum compressive strain of $\epsilon_{xc} = 0.21$ according to FEAs, which is assumed to be the compressive strain at which material failure is imminent (see also fig. 7.6).

7


deformation (i.e. $0 < y < 72$ mm). For better fits for displacements exceeding $y \approx 72$ mm, λ should be slightly increased.

See fig. 7.3 for the results of the FEAs and the algebraic model. Results correspond well, both for low and high deformation. See also results for different combinations of λ and k_c in fig. 7.4. These include combinations of choices in (Boyle et al., 2003), where $k_c = 2.65$ and $\lambda = 0.85$, and $\lambda = \frac{2}{3}$, which corresponds to all rigid bodies having equal length (i.e. $l_1 = l_2 = l_3$). Of these alternatives, the choice $\lambda = \frac{2}{3}$ and $k_c = \pi^2/4$ (fig. 7.4b) leads to the best results, yet not outperforming those in fig. 7.3 for the range $0 < y < 72$ mm.

NORMAL STRESS

Note that our model does not include deformation due to normal stresses. The stiffness due to deformation caused by normal stresses can be included by also modelling axial springs in the bodies of the PRBM, like in (Venkiteswaran, 2017). However, this introduces a significantly more complicated system and the resulting additional deformation is negligible if the slenderness ratio² (P/t) of the beam is sufficiently high, which is the case for our system. As a ballpark figure for our system, the deformation due to normal stress can be approximated by $FL/(AE) = FL/(wtE)$, which equals a mere 0.01 mm for $F = 100$ N.

²The slenderness ratio is proportional to the inverse of the beam geometry factor $f_{bg} = I/(AP^2)$ (Venkiteswaran, 2017); for the specific case of the rectangular beam profile $f_{bg} = t/(12P)$.

(a) $\lambda = 0.85$, $k_c = 2.65$ (b) $\lambda = 2/3$, $k_c = \pi^2/4$ Figure 7.4: Results for different values of λ and k_c .

7.4. FAILURE

Failure of the material occurs if the maximum tensile or compressive strain in the spring exceeds that of the material's failure (ultimate) tensile strain ϵ_{xt} or failure compressive strain ϵ_{xc} respectively. For springs made of GFRP, these values are assumed to be $\epsilon_{xt} = 0.03$ and $\epsilon_{xc} = 0.021$ respectively, as copied from [Kussmaul \(2016\)](#). However, note that these values largely depend on the type of fibres used, and they furthermore assume that failure occurs in the fibres rather than the bond. For $E = 35$ GPa, the corresponding ultimate tensile and compressive stresses are $\sigma_{xt} = E\epsilon_{xt} = 1050$ MPa and $\sigma_{xc} = E\epsilon_{xc} = 735$ MPa respectively.

Maximum strain in compression and tension are approximately equal (compressive strain is only slightly higher due to normal stress), so failure is likely to occur due to compression. This strain is maximum on the inner side in the middle of the spring.

The maximum compressive strain ϵ_{cmax} is calculated as

$$\epsilon_{cmax} = \frac{M_b t}{2IE} + \frac{F_n}{AE} \quad (7.13)$$

to which both the bending moment M_b and normal force F_n contribute. The term introduced by F_n is typically negligible (if the slenderness ratio is high). We can set $F_n = F$ and as a first approximation of M_b we can try $M_b = M$ (eq. (7.7)), which is according to the PRBM.

See fig. 7.5 for a comparison between maximum compressive strain according to FEAs and the strain calculated by eq. (7.13) with $F_n = F$ and $M_b = M$. The strains from FEAs have been recorded as the worst principle compressive strain at the centre node that is located on the mirror plane of the spring. It is evident that the theory significantly underestimates strains. For the system with $a = 0.1$ mm, failure occurs at $y = 46$ mm according to simulations and at $y = 67$ mm according to theory.




Figure 7.5: Maximum compressive strains corresponding to data in fig. 7.3, with strains according to eq. (7.13) with $M_b = M$ and $F_n = F$ dashed. The fit is poor. The failure compressive strain $\epsilon_{c,c} = 0.021$ is plotted dash-dotted in black.

SHAPE FACTOR

The discrepancy is best explained by the poor prediction of the spring's shape by the PRBM. In the PRBM, the second body remains always vertical, which implies that the moment experienced at the centre of that body is equal to that experienced at its joints ($M_b = M$). In reality however, the body always deflects further from the line of F at the centre than at its joints, which indicates that M always underestimates M_b at the centre of the body.

To rid this discrepancy, we introduce a shape factor f_s such that $M_b = f_s M$. For our choice of $\lambda = 0.6$ and the range $0 < y < 72$ mm, it was found empirically that a constant shape factor of $f_s = 1.25$ leads to good results, see fig. 7.6.

The shape factor is largely dependent on the relative size of the second body, i.e. of λ . A more general expression for f_s , as function of λ , was empirically found to be $f_s = 1 + (1 - \lambda)^{1.5}$.

THINNER SPRINGS

To increase the maximum displacement of the spring before failure, one could increase the slenderness ratio of the spring by reducing the thickness t . See figs. 7.7a and 7.7b for results of springs with various t , $a = 0.1$ mm and all other properties unchanged ($L = 144$ mm, $E = 35$ GPa, $w = 24$ mm, $k_c = \pi^2/4$, $\lambda = 0.6$).






Figure 7.6: Maximum compressive strains, similar to fig. 7.5, but with theoretical values compensated by a shape factor such that $M_b = 1.25M$.



(a) forces



(b) maximum strains

Figure 7.7: Results for different values of t and $a = 0.1$ mm.

7.5. FUTURE CONSIDERATIONS

The shape of the analysed curved compression spring is simple and well-defined, comprising a constant thickness, width and radius. This has the advantage that the spring is simpler to analyse and produce. The disadvantage of this simple shape is that the spring does not deform equally over its full length, which makes the spring heavier than necessary. Since the bending moment is proportional to the moment arm (and since normal stress is negligible for springs with a high slenderness ratio), the spring is deformed much less close to its hinges than at its centre. This suggests that spring material close to its hinges is not utilized efficiently. An improved spring shape can be obtained by making the spring thicker or wider at the centre, so that the maximum bending stress is approximately constant over the full length of the spring. Similar shapes are observed in bows and conventional leaf springs.

Nevertheless, the simple shape welcomes the use of GFRPs for affordable custom orders, which (as a material) has an energy density that is at least several times higher than the spring steel that is used in conventional helical springs. The GFRP ankle spring with dimensions equal to those resulting from the optimisations in chapter 8 weighs only $m = 28\text{ g}$, calculated as $m = Pwt\rho$ (with $w = 38\text{ mm}$, $t = 2.4\text{ mm}$ and $P = 159\text{ mm}$ and an assumed mass density of $\rho = 1.9\text{ kg/dm}^3$), and it does not fail under the assumption that $\epsilon_{xc} = 0.21$. At a mass of only 28 g, improving the shape of the spring will not lead to substantial weight reductions. In addition, difficulties were already encountered in finding a company that could produce even just the simple spring shape (as presented in this chapter) from a unidirectional high quality (S-glass) GFRP. As such, the spring is considered *good enough*, which is in line with the design philosophy as outlined in chapter 2; further improving the spring's shape is instead recommended for future work. Such an improved shape could be of particular interest for a future design iteration of Skippy if it would require a more powerful spring or if $\epsilon_{xc} = 0.21$ turns out to be overestimated. For this purpose, it is useful to understand to which extend the spring shape could be improved. This is further explained below.

7

7.5.1. SHAPE EFFICIENCY

To understand the extend to which the spring can be improved by altering its shape, we can compare its potential energy storage \mathcal{E} with that of a bending beam \mathcal{E}_{bb} of identical shape that is loaded with a pure bending moment, which is considered to be the ideal case.

It can be calculated that $\mathcal{E}_{bb} = 1/6 E \epsilon_{xc}^2 Pwt$, where Pwt is the spring's volume.³ The energy stored in the spring \mathcal{E} can be found by integrating the force F over the displacement y until failure (indicated by the x-mark in fig. 7.3), i.e. $\mathcal{E} = \int_0^{y_x} F dy$, where y_x corresponds to elongation at which maximum compressive strain $\epsilon_{cmax} = \epsilon_{xc} = 0.21$. This integration can be approximated numerically using FEA data or analytically using eqs. (7.11) and (7.13).

We can define the spring's shape efficiency $\eta_{\mathcal{E}} = \mathcal{E} / \mathcal{E}_{bb}$, for which we find $\eta_{\mathcal{E}} \approx 43\%$. The value is calculated using FEA data and differs slightly for different spring parameters,

³Essentially, the volumetric energy density of a beam loaded in pure bending equals one-third the volumetric energy density $1/2 E \epsilon_{xc}^2$ of its material loaded in pure tension.

but not more than 1 %, as shown in table 7.1.

$\eta_{\mathcal{E}} \approx 43\%$ suggests that in the most ideal case the spring's volume (and thus mass) can be reduced by up to 57 %. Such an improvement is however not achievable in practice, for several reasons:

- *Hingeability* — Since the bending moment is exactly zero at the hinges, it is suggested the spring's cross-sectional surface area has to be infinitesimally small at its hinges. Since the spring needs to be physically held in place at its hinges, this is not possible. The spring needs to have sufficient width and thickness at its hinges to be held in place and to account for material stresses other than those caused by the bending moment.
- *Other stresses* — In addition, stresses other than that caused by the bending moment are present in the material, and these stresses become more prominent if the spring has a variable cross-section. This is especially noticeable for a spring of varying width: if the spring is not hinged over its full width, then it will bend more out of the plane. When using a fibre-based material (non-isotropic), it might be required to add more off-axial (e.g. diagonal) fibre layers, which weaken it in its axial direction.⁴
- *Machinability* — A spring of varying width would have to be trimmed with curved rather than straight lines. This is substantially more difficult for FRPs than for steel. Local imperfections could lead to premature cracking and failure of the spring. A spring of various thickness made of a FRP faces similar issues, as the spring thickness will have to be discretized according to the layer thickness of the FRP.

With $\eta_{\mathcal{E}} \approx 43\%$ and the above listed culprits, it is estimated that an optimised spring cannot reduce the weight by more than a factor of two. For the case of Skippy's spring as presented in chapter 8, which weighs 28 kg, no more than 14 g could be saved by optimising its shape.

a [mm]	0.1	1	2	4	8	12	16
Pwt [cm ³]	5.18	5.18	5.19	5.19	5.23	5.28	5.35
\mathcal{E} [J]	5.70	5.71	5.71	5.71	5.72	5.75	5.79
\mathcal{E}_{bb} [J]	13.3	13.3	13.3	13.4	13.4	13.6	13.8
$\eta_{\mathcal{E}}$ [%]	42.7	42.8	42.8	42.7	42.6	42.3	42.0

Table 7.1: Energy storage of springs in fig. 7.3 (i.e. springs of various arc widths a , and $L = 144$ mm, $w = 24$ mm, $t = 1.5$ mm, $\epsilon_{xc} = 0.021$ and $E = 35$ GPa) and for pure bending beams of identical dimensions. The spring's shape efficiency approximately equals $\eta_{\mathcal{E}} \approx 43\%$, independent of shape. Calculations have been repeated for a third of the failure strain ($\epsilon_{xc} = 0.007$, not listed), for which also $\eta_{\mathcal{E}} \approx 43\%$ was found.

7.6. CONCLUSION

This chapter has presented an analytic representation of a curved compression leaf spring, based on a 1-DoF PRBM. Despite its simplicity, the model provides an accurate fit for the spring's force-displacement profile for at least 50% compression, and for

⁴Axial here refers to the line of y and F , which is the line passing through both hinges.

a large range of spring curvatures. This is partly due to the choice of the stiffness coefficient k_c , the definition of the characteristic length L_c and the definition and choice of the characteristic radius factor λ . The stiffness coefficient $k_c = \pi^2/4$ is based on that of the buckling column rather than that of a cantilever beam, because the system more closely resembles the buckling column. The buckling column is in fact a special case of the system. Both the characteristic length and characteristic radius factor are based on perimeters rather than beam lengths of the PRBM, which further improve the model accuracy. A fixed characteristic radius factor at $\lambda = 0.6$ was empirically found to provide accurate results for a wide range of curvatures. The model is amended by failure prediction based on calculating maximum bending strain and a shape factor to compensate for the difference in shape between the PRBM and that of the real spring. The simplicity and coverage of the models make them suitable for optimisation purposes.

8

DESIGN OPTIMISATIONS

To conclude a process of iterative design, single-objective optimisation studies have been performed to finalise Skippy. The optimisation studies make use of inertial data from the detailed CAD model presented in chapter 6. They have shown that Skippy is good at making somersaults, but not at making ordinary straight-up hops. As a result, few changes have been applied to Skippy's design, comprising a longer leg, which have resulted in an increase of Skippy's jump height, especially for making straight-up hops. Skippy's design is then finalized based on this design. The optimisation problem is furthermore an example of machine and behaviour co-design and of designing at the performance envelope, and demonstrates its complexity despite (1) being a simplified (single-objective) optimisation problem (2) applied to a simplified system (2D) of (3) a robot with only few actuators (Skippy has 1 actuator in the sagittal plane). Nevertheless, it has resulted in a high performance design capable of outperforming other robots and animals.

8.1. INTRODUCTION

The design presented in chapter 6 has been the result of iterations and is nearly finished. This chapter finalizes a first version of that design through optimisation studies, in which few remaining physical parameters are chosen for achieving controlled high hops.

The optimisation studies are single-objective optimisations, with a focus on preparing Skippy's highest jump. As such, the resulting design will be good at a particular type of hopping (high hopping), but not necessarily at all types of hopping, or at any behaviour other than hopping. The optimisations are furthermore limited to 2D (the sagittal plane), and have been done in various stages, which has mainly been done to simplify the optimisation problem and to accelerate decision-making.¹ This chapter in particular presents the optimisations towards the robot's finalisation, which considers only

¹ *Acknowledgements* — Dr. Featherstone has performed earlier stages of optimisations and has implemented most of the corresponding Skippy model, as described in section 8.2. Several specific ideas and implementations concerning the modelling and optimisation are also that of Dr. Featherstone, for which due credits are given throughout the remainder of this chapter in separate footnotes.

few physical parameters for optimisation. Earlier iterations between mechanical design, simulations and optimisations have already resulted in selection of most of the robot's physical parameters. These include parameters such as those of the ankle and the torso length, as discussed in section 6.6 and section 6.7 respectively.

The iterative approach does not guarantee that the results are a global optimum, but this is neither the objective. Instead, as for Skippy, the objectives of the optimisations are to

- find a system that is 'good enough' at making high jumps,
- obtain an understanding of which and how parameters influence the jumping performance (i.e. as a guide for heuristic design), and
- select the few remaining parameters in an optimal way (i.e. to provide a justification for their selection).

The optimisations also serve a more general purpose, namely to

- be exemplary for machine and behaviour co-design,
- be exemplary for designing at the performance envelope, and
- illustrate the inherent complexity of such an optimisation problem, even for such a simplified system.

The optimisations are an example of machine and behaviour co-design, as only the final outcome is specified in advance, namely that of reaching a specific jump height. It is the task of the optimiser to find both an optimal command signal and optimal physical parameters. It is also an example of designing at the performance envelope, as the model and command signal are not tied to a specific controller, but have instead been implemented to be operable in saturation. This has been done by imposing physical constraints on the velocity, voltage and current. The idea is similar to that of a bang-bang controller, but it is worth noting that

- we consider the limits of multiple parameters (i.e. velocity, voltage and current) instead of one;
- the physical limits are directly related to the capabilities of the hardware (they are not software limits like maximum joint torque);
- the command signal is slightly more complicated than a single step up and step down, as further explained in section 8.3.2.

Skippy's model that is used for optimisations is described in section 8.2. The optimisation set-up and results are described in sections 8.3 and 8.4 respectively, based on which Skippy's design is finalised. A discussion follows in section 8.5.

8.2. SKIPPY MODEL

This section describes the model of Skippy, as used for optimisations and for calculations performed in section 6.10. The presented model is more realistic than that of previous simulation studies on Skippy, such as those conducted in [Driessen \(2015\)](#). The model includes a four-bar linkage instead of the 1:1 geared hip joint depicted in fig. 3.1b, and

implements a foot with rolling contact rather than a point foot. In addition, the models of the motor, main spring, ankle spring and end stops are more realistic compared to those of previous studies.




Figure 8.1: Full kinematic model of Skippy in its *open-loop zero* configuration $\mathbf{q} = \mathbf{0}$, not to scale. Loop closure constraints would be satisfied if the points connected by the two dotted lines coincided, forming revolute joints. Joint coordinate frames and corresponding joint numbers are drawn in red; body names and corresponding numbers are written in black. Their transformations and loop closure kinematics are fully defined by the 14 indicated parameters.

8.2.1. DYNAMIC MODEL DESCRIPTION

Skippy's planar closed-loop system is modelled as an open-loop planar mechanism with loop closure constraints. The open-loop kinematic model is depicted in fig. 8.1, and the closed-loop kinematic model in fig. 8.2. It consists of ten generalized joints and ten bodies, both of which are numbered 1 to 10. Joint i connects body $\lambda(i)$ to body i , where λ is the parent array describing the connectivity of the kinematic tree, as listed in table 8.1. Joints are either revolute or prismatic. Most of the numbered joints and bodies are named, with body $\lambda(1) = 0$ being the fixed ground. Joint types and names are also listed in table 8.1.

The ten joint positions are elements of the tuple \mathbf{q} , and corresponding joint forces




Figure 8.2: Full kinematic model of Skippy in its *closed-loop* configuration. Parts that are not modelled as rigid bodies (ankle spring, main spring and nut) are drawn in gray. The 14 kinematic parameters indicated in fig. 8.1 are to scale, according to their values described in section 8.3.3 and to Skippy shown in fig. 6.1. The left figure depicts Skippy's closed-loop *zero* configuration, for which $\mathbf{y} = \mathbf{0}$ and $\phi_{\text{hip}} = 0$ (fully flexed). The right figure depicts a non-zero hip angle ϕ_{hip} , corresponding to the touching condition of the hip end stop (like in fig. 6.13).

are elements of the tuple $\boldsymbol{\tau}$:

$$\mathbf{q} = \begin{bmatrix} q_1 \\ q_2 \\ \vdots \\ q_{10} \end{bmatrix} \quad \text{and} \quad \boldsymbol{\tau} = \begin{bmatrix} \tau_1 \\ \tau_2 \\ \vdots \\ \tau_{10} \end{bmatrix}.$$

Corresponding joint variables are either angles (in rad) or distances (in m), and joint forces are either torques (in Nm) or forces (in N), for revolute and prismatic joints respectively. Some of the numbered joints and bodies are fictitious, as they do not physically appear in Skippy, whereas other physical joints in Skippy do not appear in \mathbf{q} , because they are the result of kinematic loop closure constraints. This is further explained below.

BODY AND JOINT DESCRIPTIONS

Bodies 1 to 3 are fictitious massless bodies, and their coordinate frames all coincide. Joints 1 and 2 are prismatic joints in the horizontal (x) and vertical direction (y) that determine the position of Skippy's toe, which effectively make Skippy a floating robot. q_1 and q_2 are held at zero by fictitious actuator forces τ_1 and τ_2 , so that Skippy always makes contact with the ground. Consequently, τ_1 and τ_2 are the GRFs. These joints have been introduced to (1) monitor the GRFs and (2) solve impulsive equations of motions for Skippy's landing from a previous jump, as described in section 6.10. τ_1 and τ_2 are calculated using hybrid dynamics, as further explained in section 8.2.2.

Like joint 1, joint 3 is a prismatic joint in the horizontal direction. It is introduced to implement—together with joint 4—the wheel-like 1-DoF rolling of Skippy's curved toe, which has a finite radius of r_{tx} . Joint 4 describes the foot angle and the robot's overall orientation, and is referred to as base joint. The 1-DoF rolling is implemented

i	$\lambda(i)$	joint type	joint name	body name
1	0	px		
2	1	py		
3	2	px		
4	3	r	base	foot
5	4	r	ankle	leg
6	5	r	lever-leg	lever
7	6	r	lever-torso	torso
8	7	r	follower-torso	follower
9	7	r	rocker	rocker
10	9	px	main	coupler

Table 8.1: Skippy body and joint names and properties. A joint is either revolute (r) or prismatic in its local x axis (px) or y axis (py). Joint numbers of independent joint variables are bold.

with a kinematics loop closure function that takes q_4 as independent position variable, as further explained in section 8.2.2.

Bodies 4 to 10 are the foot, leg, lever, torso, follower, rocker and coupler respectively. The inertia of the swivel actuator and crossbar have been merged with the torso. Furthermore, to simplify model dynamics, some of Skippy's smaller and flexible bodies have been merged with other rigid bodies. These include the ankle and main springs, nut and actuator shaft. The mass and inertia of the ankle spring at rest have been added to the foot, as the effect of its deformation on dynamics is negligible.

The main spring introduces an additional DoF and thus additional state variables. However, the main spring and nut are not modelled as separate rigid bodies; the mass and inertia of the main spring (at rest) and nut have been added to the coupler. The additional DoF is modelled separately as a massless 1D dynamic system acting on the coupler. The rotational inertia of the actuator shaft (consisting of the screw spindle, motor rotor and coupling) has been added to this dynamic model as reflective terms. This is further explained in section 8.2.3. Its mass has been added to the rocker.

Skippy has two more joints that do not appear in \mathbf{q} and table 8.1, because they are the result of kinematic loop closure constraints of the hip. These are the revolute joints that connect the lever to the coupler (lever-coupler joint) and the follower to the leg (follower-leg joint). As fig. 8.1 depicts the open-loop kinematics, the joints appear dismantled. The points connected by dotted lines are to be coincided to satisfy the loop-closure constraints.

The loop-closure constraints reduce the number of DoFs of the linkage mechanism from 5 (q_6 to q_{10}) to 1. The independent joint variable that is used to describe the position of the hip is q_{10} . The number of independent joint variables of the model has thus reduced from 10 to 5, and the fictitious actuators of Skippy's foot (ground fixation $q_1 = q_2 = 0$) further reduce the number of DoFs to 3. Corresponding calculations of the dynamic model (including loop closure) are explained below in section 8.2.2.

8.2.2. HYBRID DYNAMIC AND LOOP CLOSURE CALCULATIONS

Hybrid dynamics and loop closure of the robot is done as described in Appendix B of [Driessen \(2015\)](#), except that they are adjusted to incorporate the new joint variables and loop closure constraints. A summary of their implementation follows below.

LOOP CLOSURE

The EoM of an open-loop multi-rigid-body system can be written as

$$\boldsymbol{\tau} = \mathbf{H}\ddot{\mathbf{q}} + \mathbf{C}, \quad (8.1)$$

where \mathbf{H} is the joint-space inertia matrix and \mathbf{C} a joint-space tuple of force terms. Solving eq. (8.1) for $\ddot{\mathbf{q}}$ (forward dynamics calculations) requires joint position variables \mathbf{q} to be independent from each other, which is not the case for closed-loop mechanisms. Instead, we can write

$$\dot{\boldsymbol{\tau}} = \dot{\mathbf{H}}\dot{\mathbf{y}} + \dot{\mathbf{C}}, \quad (8.2)$$

where \mathbf{y} are independent joint position variables, $\dot{\boldsymbol{\tau}}$ the corresponding joint torques, and

$$\dot{\mathbf{H}} = \mathbf{G}^T \mathbf{H} \mathbf{G} \quad (8.3)$$

$$\dot{\mathbf{C}} = \mathbf{G}^T (\mathbf{C} + \mathbf{H} \mathbf{g}), \quad (8.4)$$

with the velocity mapping or Jacobian $\mathbf{G} = \mathbf{J}(\mathbf{q}, \mathbf{y}) = \partial \mathbf{q} / \partial \mathbf{y}$ and the tuple of convective accelerations $\mathbf{g} = \dot{\mathbf{G}}\dot{\mathbf{y}} = \mathbf{J}(\mathbf{G}\dot{\mathbf{y}}, \mathbf{y})\dot{\mathbf{y}}$, so that $\dot{\mathbf{q}} = \mathbf{G}\dot{\mathbf{y}}$, $\dot{\boldsymbol{\tau}} = \mathbf{G}^T \boldsymbol{\tau}$ and $\ddot{\mathbf{q}} = \mathbf{G}\ddot{\mathbf{y}} + \mathbf{g}$. The joint position variables \mathbf{q} are calculated from \mathbf{y} with the loop closure function $\mathbf{q} = \boldsymbol{\gamma}(\mathbf{y})$. For Skippy we choose

$$\mathbf{y} = \begin{bmatrix} y_1 \\ y_2 \\ y_3 \\ y_4 \\ y_5 \end{bmatrix} = \begin{bmatrix} q_1 \\ q_2 \\ q_4 \\ q_5 \\ q_{10} \end{bmatrix}, \quad \text{and correspondingly} \quad \dot{\boldsymbol{\tau}} = \begin{bmatrix} \tau_1 \\ \tau_2 \\ \tau_4 \\ \tau_5 \\ \tau_{10} \end{bmatrix}.$$

The corresponding loop closure function $\boldsymbol{\gamma}(\mathbf{y})$, \mathbf{G} and \mathbf{g} can be calculated as

$$\mathbf{q} = \boldsymbol{\gamma}(\mathbf{y}) = \begin{bmatrix} q_1 \\ q_2 \\ -r_{\text{tx}} q_4 \\ q_4 \\ q_5 \\ \lambda_{\text{hip}}(q_{10}) \end{bmatrix}, \quad \mathbf{G} = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & -r_{\text{tx}} & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{G}_{\text{hip}} \end{bmatrix}, \quad \mathbf{g} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ \mathbf{g}_{\text{hip}} \end{bmatrix},$$

where γ_{hip} , \mathbf{G}_{hip} and \mathbf{g}_{hip} are the loop closure function, velocity mapping and convective accelerations of the hip mechanism respectively, as calculated in appendix C.² $q_{10} = 0$ is furthermore defined to correspond to a hip angle $\phi_{\text{hip}} = q_6 + q_7 = 0$, which is defined as Skippy's fully flexed configuration, see fig. 8.2. The hip angle ϕ_{hip} also describes the

²In software, the definition of \mathbf{g} includes also a stabilization factor, which is further explained in [Featherstone \(2008\)](#).

angle between the physical leg and torso tubes, which are parallel to the extended body segments of the diagram depicted in fig. 8.1 (although they are not coaxial). Figures 6.1 and 6.14c also depicts Skippy in its fully flexed configuration, which shows the torso and leg tubes parallel to each other. Skippy is defined to be fully extended for $\phi_{\text{hip}} = 180^\circ$, for which the leg and torso tubes are also parallel, as depicted in fig. 6.14a.

HYBRID DYNAMICS

Obtaining joint forces of the locked joints 1 and 2 calls for inverse dynamics calculations, whereas obtaining joint acceleration of all other joints calls for forward dynamics calculations. Hybrid dynamics calculations are implemented to combine the two. This is done by separating the state variables so that $\mathbf{y} = [\mathbf{y}_1 \ \mathbf{y}_2]^\top$, in which the subscript $_1$ corresponds to the variables of which joint accelerations are known and $_2$ to variables of which joint forces are known. We have $\mathbf{y}_1 = [y_1 \ y_2]^\top = [q_1 \ q_2]^\top$ and $\mathbf{y}_2 = [y_3 \ y_4 \ y_5]^\top$, and correspondingly divide $\dot{\mathbf{H}}$, $\dot{\mathbf{\tau}}$ and $\dot{\mathbf{C}}$ in quadrants so that

$$\begin{bmatrix} \dot{\mathbf{\tau}}_1 \\ \dot{\mathbf{\tau}}_2 \end{bmatrix} = \begin{bmatrix} \dot{\mathbf{H}}_{11} & \dot{\mathbf{H}}_{12} \\ \dot{\mathbf{H}}_{21} & \dot{\mathbf{H}}_{22} \end{bmatrix} \begin{bmatrix} \dot{\mathbf{y}}_1 \\ \dot{\mathbf{y}}_2 \end{bmatrix} + \begin{bmatrix} \dot{\mathbf{C}}_1 \\ \dot{\mathbf{C}}_2 \end{bmatrix}.$$

Since the foot is constrained to make contact with the ground, we have $\mathbf{y}_1 = \mathbf{0}$ and thus $\dot{\mathbf{y}}_1 = \ddot{\mathbf{y}}_1 = \mathbf{0}$. It follows that

$$\begin{aligned} \ddot{\mathbf{y}}_2 &= \dot{\mathbf{H}}_{22}^{-1} (\dot{\mathbf{\tau}}_2 - \mathbf{C}_2) \\ \dot{\mathbf{\tau}}_1 &= \dot{\mathbf{H}}_{12} \ddot{\mathbf{y}}_2 + \mathbf{C}_1, \end{aligned}$$

which complete the hybrid dynamics calculations. Note that \mathbf{y}_1 also corresponds to joints at which impact occurs from landing. As such, the terms $\dot{\mathbf{H}}_{21}$ and $\dot{\mathbf{H}}_{22}$ correspond exactly to those in eq. (6.7).

8.2.3. MAIN ACTUATOR MODEL

Skippy's series elastic main actuator is modelled as a 1D dynamic system, to interact with the coupler (body 10) of the main dynamic model described in section 8.2.1. The model incorporates rotational inertia of the actuator shaft. The implementation is similar to that described in [Driessen \(2015\)](#), and is based on the following equations:

$$\tau_{10} = F_s(x_s, \dot{x}_s) \quad (8.5)$$

$$q_{10} = x_n - x_s \quad (8.6)$$

$$\tau_m = I_m \ddot{x}_n n - \frac{1}{n} (F_s + F_{\text{fric}}) \quad (8.7)$$

F_s , x_s and \dot{x}_s are the main spring force, displacement and velocity respectively. x_n is the nut position. $n = 2\pi/p$ is the transmission ratio in rad m^{-1} with a pitch of $p = 4 \text{ mm}$. τ_m is the main motor rotor (output) torque and I_m the rotational inertia of the rotor, spindle and coupling. Their corresponding rotational position is $\phi_m = x_n/n$. F_{fric} models the friction of the ballscrew as sliding friction, and is calculated as $F_{\text{fric}} = \text{sign}(\dot{x}_n) \frac{1-\eta}{\eta} |F_s|$ for $\dot{x}_n \neq 0$, with an efficiency of $\eta = 0.9$.

Compared to the implementation in [Driessen \(2015\)](#), the motor model (i.e. the model relating motor torque to voltage and current) and spring model ($F_s(x_s, \dot{x}_s)$) have been implemented to be more realistic, each of which are further elucidated below.

MOTOR MODEL

A DC motor model is implemented for the purpose of being able to (1) command the motor to operate in saturation based on a maximum voltage and current limit, and (2) calculate heat generation. The motor model that is used is identical to that defined in chapter 4 and [Driessen \(2015\)](#), except that it has been amended with temperature calculations and a temperature-dependent winding resistance.³ Motor inductance is not modelled for the reason that it is computationally costly (as it requires many tiny integration steps due to the fast dynamics of the current) whereas its effect is negligible for the simple voltage profile that is commanded to Skippy.

The temperature-dependent winding resistance is calculated as $R = R_{25}(1 + \alpha_{Cu}(T - 25^\circ\text{C}))$, where R_{25} is the winding resistance for a temperature of $T = 25^\circ\text{C}$ (as provided by the catalogue) and $\alpha_{Cu} = 0.0039$ is the thermal resistance coefficient for copper ([Maxon Motor, 2017](#)). As a result, an increased temperature effectively reduces the motor performance. However, it was found that during a single stance phase the coil temperature does not increase more than 6°C , so the effect on performance reduction is hardly noticed.

MAIN SPRING MODEL

The main spring has been implemented with a (continuous) regressive polynomial as force-displacement profile, rather than a regressive piecewise linear function like in [Driessen \(2015\)](#). As explained in section 6.7, a regressive spring profile is generally preferred for improved energy storage and control bandwidth, but the required regressive-ness for balance control has not yet been investigated. As discussed in section 6.7 it is first attempted to make Skippy operable with only a linear rubber element, without the addition of an end stop to make the profile more regressive, so as to save weight and reduce complexity. The linear rubber spring is modelled as an incompressible neo-Hookean solid and a linear damper. According to the neo-Hookean model, the spring force under pure elastic deformation is

$$F_{s,\lambda} = \frac{4}{7}F_{s,2} \left(\lambda - \frac{1}{\lambda^2} \right) \quad (8.8)$$

where $\lambda = \epsilon + 1 = x_s/x_{s,0} + 1$ is the extension ratio of the main spring, $x_{s,0}$ the length of the unloaded spring and $F_{s,2}$ is the spring force under pure elastic deformation at an extension of $\lambda = 2$ ($x_s = x_{s,0}$), proportional to its cross-sectional surface area, to be determined through optimisation (tentative sizing of the spring has been provided in section 6.7). This model is only mildly regressive, as it implies that for $x_s \rightarrow \infty$ the stiffness is only three times lower than for $x_s = 0$. Note however that the model is only accurate up to strains of ca. $\epsilon = 0.5$ ([Treloar, 1973](#)); rubbers typically become slightly more regressive for higher strain, as more accurately described by for instance the Mooney-Rivlin solid model. Since regressive-ness is desired, the Neo-Hookean model thus makes a *safe* (pessimistic) assumption about the rubber's actual regressive-ness.

The spring's energy losses are taken into account by including a linear damper in

³*Acknowledgements* — The temperature calculations and respective implementations have been done by Dr. Featherstone.

parallel to the non-linear neo-Hookean element:

$$F_s = F_{s,\lambda} + d_s \dot{x}_s \quad (8.9)$$

where d_s is a damping coefficient.

8.2.4. ANKLE SPRING

The GFRP ankle spring is modelled using the compression leaf spring formulae derived in chapter 7 and the triangle kinematics defined in section 6.6.1. The end stop that is used to implement the progressive part of the ankle spring profile (effectively limiting ankle joint rotation), is implemented using the end stop model described next.

8.2.5. END STOP MODELS

Skippy's model includes two visco-elastic end stops, one in the ankle and one in the hip. Both end stops are implemented with the same model, which uses a quadratic spring profile and a damping that is proportional to compression x , so that the spring force calculation takes the shape of

$$F = kx^2 + dx\dot{x} \quad \text{for } x > 0, \quad (8.10)$$

where F and x are either a force and displacement or torque and rotation respectively, k a stiffness coefficient and d a damping coefficient. Non-linearity in this function stems from contact geometry rather than from hyperelastic deformation like that described by the Neo-Hookean model (Azad & Featherstone, 2014b).

The ankle end stop is attached on the foot and makes contact with the torso. It is implemented to act directly on the ankle joint, so that $x = q_5 - \phi_{ea,0}$, where $\phi_{ea,0}$ is the angle at which the ankle end stop starts being compressed.

The hip end stop is attached on the lever and makes contact with the torso. It is implemented as a torsional spring to act directly on joint 7, so that $x = -q_7 + \phi_{eh,0}$, where $\phi_{eh,0}$ is the lever angle at which the hip end stop starts being compressed.

8.3. OPTIMISATION ENVIRONMENT

Simulations are done in Matlab and Simulink, using Dr. Featherstone's Spatial v2 library Featherstone (2014). Optimisations are performed with Matlab's built-in `fmincon` function, using the sequential quadratic programming (SQP) algorithm. The algorithm was found to give better results than the default but faster interior point algorithm, possibly because some of the optimisation parameters were found to converge to either their lower or upper bound. The SQP algorithm allows for the parameters to be exactly equal to their bounds and for solutions to slide along constraints.

The optimiser attempts to minimize the cost function $f(\mathbf{x})$, where \mathbf{x} are optimisation parameters, subject to bounds of x and constraints. With a focus on Skippy's most demanding task of reaching its maximum jump height, the cost function has been defined to maximize the vertical take-off velocity v_{end} , implemented as $f(\mathbf{x}) = -v_{\text{end}}$. The velocity is calculated as the result of a simulation of Skippy's stance phase, using the model described in section 8.2, which is affected by the choice of optimisation parameters and subject to constraints.

The constraint function is further explain in section 8.3.1. The optimisation parameters \mathbf{x} contain both parameters that define the command signal and physical parameters. The command signal is defined in section 8.3.2 and the selection of optimisation parameters is explained in section 8.3.3.

8.3.1. CONSTRAINT FUNCTION

Constraints can be either or a combination of

- Linear inequality constraints, so that $Ax \leq B$ is satisfied, where A and B are numerical matrices.
- Linear equality constraints, so that $A_{\text{eq}}x = B_{\text{eq}}$ is satisfied for numerical matrices A_{eq} and B_{eq} .
- Non-linear inequality constraints, so that $C(x) \leq 0$ is satisfied for the function $C(x)$.
- Non-linear equality constraints, so that $C_{\text{eq}}(x) = 0$ is satisfied for the function $C_{\text{eq}}(x)$.

All the constraints are based on simulation results, which require the use of non-linear constraints. The most important ones are those related to steering, as described by the robot's momentum at take-off. The momentum of the robot's mass in the horizontal direction p_x is constrained to be zero at take-off (i.e. $p_{x,\text{end}} = 0$, which is an equality constraint), so that the robot hops on the spot (rather than it moving either forward or backward). The angular momentum of the robot about it's CoM p_r must be constrained at take-off so as to avoid uncontrolled rotation of Skippy during its flight phase. Either it must be constrained to zero, so that it does not rotate at all, or it must be within a feasible range in order for Skippy to make one somersault, (rather than for it to land on its head), such as demonstrated by the biped of [Hodgins and Raibert \(1990\)](#).

TAKE-OFF MOMENTUM FOR MAKING A SOMERSAULT

The angular momentum that Skippy should have to make one full somersault can be estimated as follows, based on Skippy's flight time t_f , mass m and radius of gyration r in the sagittal plane (or rotational inertia):

$$p_r = \frac{2\pi m r^2}{t_f}$$

Skippy's radius of gyration is a function of Skippy's configuration, and is significantly higher when fully extended than when fully flexed: $r = 0.242\text{m}$ when fully flexed and $r = 0.393\text{m}$ when fully extended. The flight time can be approximately related to the jump height as $t_f \approx 2\sqrt{(2h/g)^4}$, from which we find $t_f = 1.6\text{s}$ for a jump height of $h = 3.2\text{m}$. For a mass of $m = 3\text{kg}$, we then find $p_r = 0.69\text{Nms}$ or $p_r = 1.82\text{Nms}$ to make a full somersault if Skippy would be either fully flexed or extended respectively. Note that, in reality, Skippy cannot change its orientation instantly, due to having a specific orientation and velocity shortly before landing and after take-off. A more conservative range of feasible angular momentum to make a somersault is approximately $1 < p_r <$

⁴This relation is exact if Skippy would land in the same configuration it had at take-off, and if aerial friction is neglected.

1.5Nms. The exact landing conditions can then be tuned by the in-flight controller, which controls Skippy's configuration during the flight phase.

SOMERSAULT VERSUS STRAIGHT JUMPS

The angular momentum at take-off is constrained at either $p_{r,\text{end}} = 0$ (equality) or $1 < p_{r,\text{end}} < 1.5\text{Nms}$ (inequality), enabling Skippy to jump straight up or to make a backward somersault respectively. Earlier studies have shown that Skippy is more inclined to generate positive angular momentum instead of zero (having simulated a higher jumping height if it is allowed to have positive angular momentum at take-off), which is related to its asymmetric design and heavier upper body (Driessen, 2015). As such, Skippy is expected to be able to jump higher when it aims for a somersault, rather than for it to jump straight up. It is perfectly acceptable if Skippy makes its highest hop while somersaulting, but it should still be able to jump straight up sufficiently high so that it can meet the initial conditions in preparation for the stance phase of the highest (somersault) jump. In the presented simulation studies, these initial conditions assume that Skippy lands from a previous jump of 1.5 m high, as further explained in section 8.3.3. To build in room for error and to be conservative on Skippy's actual capabilities, it was decided that Skippy should be able to jump straight up without somersaulting to a height of at least 2 m.

OTHER CONSTRAINTS

Other than momentum constraints, the non-linear constraint function includes various other inequality constraints:

- To prevent Skippy from flexing further than is physically possible (because it would collide with the head), it is constrained to $q_{10} \geq 0$ for the full duration of the simulation (correspondingly $\phi_{\text{hip}} = q_6 + q_7 \geq 0$), which corresponds to the configuration shown in fig. 6.1.
- To reduce the effect of possible slipping (which is not modelled because the robot is constrained to the ground), the tangential GRF τ_1 at take-off is constrained to be less than 15 N, i.e. $|\tau_{1,\text{end}}| < 15$.⁵
- For purposes of in-flight control, it is not desired that the main motor is still in saturation and at maximum speed at take-off. The speed limit of the motor at take-off is $|\dot{\phi}_{m,\text{end}}| < 300\text{rad s}^{-1} \approx 3\text{krpm}$, which is approximately half the speed of the 6000 rpm speed limit of the ball screw. To prevent the motor from operating in saturation at take-off, the total duration of the stance phase is constrained to be $t_{\text{end}} > t_{\text{rise}} + t_{\text{sat}} + 0.02\text{s}$.

⁵Slip occurs if the ratio of the tangential to the normal GRF $|\tau_1/\tau_2|$ exceeds a certain value, as determined by the coefficient of static friction between the ground and (rubber) toe. This suggests that it is better to check for this ratio rather than for an absolute value of the tangential GRF. However, shortly after landing and before take-off, τ_2 is nearly zero, for which the ratio most certainly exceeds the permissible value and thus causes the modelled robot to slip. However, the induced slipping would only last for several milliseconds, and if the tangential GRF is minimal, the effect is negligible in reality (or even non-existent, due to the contact material's unmodelled finite stiffness, which influence contact dynamics). Nevertheless, after simulations, the ratio $|\tau_1/\tau_2|$ is also checked for the remainder of the stance phase, to assure that the robot does not slip while thrusting. Results are shown in fig. 8.5a.

8.3.2. VOLTAGE COMMAND SIGNAL

The optimised command signal is motor voltage. Since voltage is the actual command sent to the motors on the real robot, it directly allows for the motor to operate in saturation by applying maximum voltage. As such, the use of a voltage command instead of for instance a commanded position, velocity or torque of the motor or an end effector (e.g. the position of the CoM or a joint variable) provides more transparency for discovering the robot's behaviour at the performance envelope. Especially commanding end effector position (requiring inverse dynamics for motor torque calculations) has shown to be particularly inefficient for Skippy due to the low stiffness of the main spring for high loadings (Driessen, 2015).

The voltage command is generally feed-forward, but it is clipped if either the ball screw velocity or motor current would exceed their respective maxima. This effectively enables Skippy to be operated in saturation of multiple variables (voltage, current and speed), by sliding along either of their limits. The implementation is according to Driessen (2015). These studies have also shown that maximum voltage is commanded during most of the stance phase in order to reach a maximum jump height, except for the final part of the stance phase, during which the voltage may be dropped for the purpose of steering. By steering is meant controlling the robot's horizontal and rotational momentum at lift-off. The motor should furthermore not operate in saturation at landing, to prevent slipping of the foot. Nevertheless, the command signal should be defined in such a way that it is maximal during most of the stance phase, and the optimiser should be allowed to tweak only the beginning and the end of the commanded profile. This is different from more conventional approaches in optimal control that define the command signal to be a smooth spline.

The simplest command signal would consist of only a step up to maximum voltage $U = U_{\max}$ at $t = 0$ for a predetermined time t_{sat} , followed by a step down to zero voltage.⁶ However, step changes of voltage (discontinuity) are not realistic due to unmodelled motor inductance and increase the likelihood of the robot to slip at impact. They furthermore provide the optimiser with only 1 parameter to optimise (the duration t_{sat}), leaving little freedom for finding the optimal command signal. The steps have been replaced for quadratic functions, providing the optimiser with 5 tunable parameters that fully define the command signal.⁷ Each of the quadratic functions are defined by a duration (t_{rise} and t_{fall}) and a curvature parameter (c_{rise} and c_{fall}), so that the entire command signal U_c is determined by the following equation:

$$U_c = \begin{cases} U_{\max} \left(t \frac{1+c_{\text{rise}}}{t_{\text{rise}}} - t^2 \frac{c_{\text{rise}}}{t_{\text{rise}}^2} \right) & \text{if } 0 \leq t < t_{\text{rise}} \\ U_{\max} & \text{if } 0 \leq t - t_{\text{rise}} < t_{\text{sat}} \\ U_{\max} \left(1 - (t - t_{\text{rise}} - t_{\text{sat}}) \frac{1+c_{\text{fall}}}{t_{\text{fall}}} + (t - t_{\text{rise}} - t_{\text{sat}})^2 \frac{c_{\text{fall}}}{t_{\text{fall}}^2} \right) & \text{if } 0 \leq t - t_{\text{rise}} - t_{\text{sat}} < t_{\text{fall}} \\ 0 & \text{otherwise} \end{cases} \quad (8.11)$$

⁶Negative voltage is avoided to prevent the motor from doing negative work, although in principle it is possible and it possibly leads to improved steering performance.

⁷*Acknowledgements* — The replacement of steps for quadratic functions has been introduced and implemented by Dr. Featherstone.




Figure 8.3: The commanded voltage profile that is used for stance phase optimisations is fully defined by the following 6 parameters: the maximum supply voltage U_{\max} , the durations t_{rise} , t_{sat} and t_{fall} and the curvature parameters c_{rise} and c_{fall} . The different colours show the effect of different values for curvature parameters.

The curvature parameters take values between -1 and 1 . The resulting profile is shown in fig. 8.3 for various values of the curvature parameters.

8.3.3. OPTIMISATION PARAMETERS AND BOUNDS

Parameters that are to be optimised include those defining (1) the command signal, (2) the robot's initial conditions and (3) the physical system, each of which are explained subsequently. Of all the parameters that define the system, only few are selected for optimisation. Below follow explanations regarding the selection of optimisation parameters. An overview of them is provided at the end of this section and in table 8.3.

COMMAND PROFILE

The commanded voltage profile is defined by the six parameters listed in section 8.3.2, and is furthermore affected by the current limit and speed limit of the ball screw. The effective maximum supply voltage U_{\max} is increased from 30 V to 31 V, in correspondence with the experienced average voltage of 8 nominally charged LiPo cells during the Tippy experiments reported in section 5.6. The maximum current is $I_{\max} = 50$ A, as limited by the drivers, and the speed limit of the ball screw is 6000 rpm.⁸ This leaves the five parameters t_{rise} , t_{fall} , t_{sat} , c_{rise} , c_{fall} for optimisation, which have been given the following bounds:

$$0.02 \leq t_{\text{rise}} \leq 0.12 \text{ s} \quad (8.12)$$

$$0.05 \leq t_{\text{sat}} \leq 0.2 \text{ s} \quad (8.13)$$

$$0.02 \leq t_{\text{fall}} \leq 0.12 \text{ s} \quad (8.14)$$

$$-1 \leq c_{\text{rise}} \leq 1 \quad (8.15)$$

$$-1 \leq c_{\text{fall}} \leq 1 \quad (8.16)$$

INITIAL CONDITIONS

Skippy has five independent generalized joint positions, two of which are constrained, leaving three DoFs, and thus six effective state (position and velocity) variables. Initial

⁸This speed limit was previously based on recommendations of colleagues in the robotics community, as mentioned in section 3.1.2. However, it has recently been experimentally verified that the ball screw is still operational at this speed (under load) using the custom experimental rig that has been developed and used to test the ring screw by Heijmink (2018) (see section 6.2.2).

velocities are chosen to emulate landing with a vertical impact velocity of 5.5 ms^{-1} , corresponding to landing from a previous (straight up) hop of ca. 1.5 m. This height corresponds to the decreased minimum target jump height of 3 m (see section 6.2) and the requirement that the target jump height should be achievable after landing from a previous hop that is not higher than half this target height (see chapter 1 and section 3.1.2).

The initial joint velocities are calculated assuming inelastic impact using eq. (6.7) for $\dot{\mathbf{q}}_1^- = [\dot{q}_1^- \ \dot{q}_2^-]^T = [0 \ -5.5]^T$. This leaves only the choice of initial position states $q_{4,0}$ (robot orientation), $q_{5,0}$ (ankle) and $q_{10,0}$ (main), which have been given the following bounds:

$$-0.47 \text{ rad} \leq q_{4,0} \leq -0.25 \text{ rad} \quad (8.17)$$

$$-15^\circ \leq q_{5,0} \leq -10^\circ \quad (8.18)$$

$$0.2 \leq q_{10,0} \leq 0.75 \text{ m} \quad (8.19)$$

The bounds of $q_{4,0}$ and $q_{10,0}$ are based on a range of their values that resulted from previous optimisations. The range of $q_{10,0}$ corresponds to a hip range of $20^\circ < \phi_{\text{hip}} < 80^\circ$. Bounds for $q_{5,0}$ are based on iterative design, and have been determined based on feasible ankle spring designs (preventing collisions) described in section 6.6.3.

PHYSICAL PARAMETERS

Deciding on which of Skippy's physical parameters to optimise and their bounds is challenging, as (1) there are many, (2) many of them have dependencies that are not necessarily easily found (or their relevance to the optimisation problem is small), (3) many combinations of parameter choices lead to infeasible designs, although this is not always clear prior to implementation in CAD (4) many combinations of parameters choices form a feasible but poor design, as caused by the optimiser having converged to a poor local minimum. As such, the strategy has been to solve the optimisation problem in various (iterative) stages, and to intermittently lock parameters based on the notion that they were 'good enough' and their implementation in the CAD model.

The physical parameters that define Skippy can be subdivided in

- 14 kinematic parameters shown in fig. 8.1, defining joint coordinate transforms and loop closure,
- 28 inertial parameters (a mass, CoM coordinates and rotational inertia for each of the 7 non-massless rigid bodies depicted in fig. 8.1),
- miscellaneous parameters, relating to the main actuator, ankle and end stops.

KINEMATIC PARAMETERS

The 14 kinematic parameters depicted in fig. 8.1 fully define the rigid body kinematics. Whereas these 14 parameters are independent if all seven bodies are not massless, some are barely effective since they do not or only marginally alter the transmission ratio from joint q_{10} to the end effector (toe). For instance, the triangle defined by the sides l_l , l_r and l_m (depicted in red in fig. C.1) as a whole can be rotated without affecting its transmission ratio $\partial q_{10} / \partial q_7$ by keeping $\phi_l - \phi_r - \phi_u$ constant. Likewise, the four-bar linkage (defined by link lengths l_u , l_f , l_k and l_v) as a whole can be rotated without affecting its transmission

ratio $\partial q_7 / \partial q_6$ by keeping $\phi_u + \phi_v$ constant. The transmission ratio can also be kept constant by equally scaling all the lengths l_v , l_u , l_k and l_f .

Simply allowing for the optimiser to change all these kinematic parameters thus leads to an ill-conditioned optimisation problem. In addition, a change of kinematic parameters also requires a change of inertial parameters, as a body's mass and inertia scale up with its size. Loading scenarios on a body also change with size: typically, a larger body demands a stronger structural design, making the robot heavier still. Inadequate consideration of these effects could lead to the optimiser shooting off in an infeasible direction. For instance, a larger robot will always result in higher jumps if the mass of the robot is not changed. Lastly, a change of parameters could lead to physical collisions, which in some cases only become clear when updating the CAD model, and are thus not easily implemented as constraints.

Most of the kinematic parameters are fixed by choices in (1) earlier stages of simulation and optimisation and (2) the mechanical design presented in chapter 6, which is further elucidated below.

Four-bar linkage The four-bar linkage kinematics are defined by the link lengths l_v , l_u , l_k and l_f and the angle $\phi_u + \phi_v$, which have all been assigned values preliminary based on required and desired properties that are 'good enough'. As such, none of the parameters are being optimised, although this does not exclude them from being optimised in a future more elaborate optimisation study. The following constraints were proposed that describe the properties of the four-bar linkage, based on which its parameters can be solved:⁹

- The moment arm of the follower is $a_{f,flx} = 5$ cm when Skippy is fully flexed and $a_{f,ext} = 2$ cm when fully extended, see fig. 6.14. These are the most important constraints, as they are directly related to the mechanism's internal forces and its proximity to a singularity. The torque transferred through the hip is calculated as the product of the force on the follower F_f and the follower's moment arm a_f , as also explained in section 6.4. A smaller arm implies a bigger force, and the linkage mechanism is at a singularity if the moment arm equals zero. When fully flexed, F_f is approximately double that on the main spring F_s for $l_l = 10$ cm and $a_{f,flx} = 5$ cm, as explained in section 6.4. The moment arm is allowed to be shorter when Skippy is extended, because less torque is transferred through the hip, which explains why $a_{f,ext} < a_{f,flx}$.
- The overall size of the four-bar linkage is constrained by $l_f + l_k = 24$ cm, which places the torso and leg tube approximately 10 cm apart when fully flexed. A smaller linkage mechanism leads to higher internal forces and increases the effect of play in the joints. The size defined by $l_f + l_k = 24$ cm is big enough to realize the aforementioned lever moment arm of $a_{f,flx} = 5$ cm, which would not be possible for a substantially smaller linkage mechanism.
- The total hip range $\Delta\phi_{hip} = \Delta q_6 + \Delta q_7 = 180^\circ$ corresponds to a lever-torso joint range of $\Delta q_7 = 90^\circ$, which on average corresponds to the transmission ratio of

⁹ *Acknowledgements* — Dr. Featherstone coined the idea of isolating the four-bar linkage, constrained it, and implemented an optimisation routine for solving its parameters.

$\partial 6/\partial 7 = 1$ of the (simplified) 1:1 geared mechanism depicted in fig. 3.1b that had proven to work well in Driessen (2015).

- For the full motion range of the hip, the transmission ratio of the four-bar linkage (i.e. from (input) lever rotation to hip rotation $\partial q_7/\partial \phi_{\text{hip}} = \partial q_7/(\partial q_6 + \partial q_7)$) does not differ by a factor greater than 1.5 over Skippy's full work space so as to limit the non-linearity of the mechanism.

With these constraints, the four-bar linkage has been solved to have the following parameters: $l_v = 50$ mm, $l_u = 42.7$ mm, $l_k = 120.6$ mm, $l_f = 119.4$ mm and $\phi_u + \phi_v = 51.7^\circ$. Its transmission ratio ($\partial q_7/\partial \phi_{\text{hip}}$) is highest when flexed. ϕ_u is furthermore fixed at $\phi_u = 20^\circ$, mainly for reasons of mechanical design, which puts ϕ_v at $\phi_v = 31.7^\circ$.

Torso and leg The torso length and its inertial properties are largely tied to placement of the rocker joint. In the presented optimisations, neither positioning of the rocker (defined by l_r and ϕ_r) nor that of the leg (defined by l_g and ϕ_g) are optimised.

As discussed in section 6.7.4, the rocker joint is constrained to be at least 47.5 cm away from the lever-torso joint, for a lever length of $l_l = 10$ cm. To keep the robot small, to limit internal forces, and to make the robot better at jumping with zero angular momentum, positioning of the rocker joint is such that it is close to its minimum distance from the hip, at $[x_{r,x} \ x_{r,y}] = l_r[\cos \phi_r \ \sin \phi_r] = [480 \ 47.5]$ mm.

It was quickly found that allowing for the leg length to be optimised without altering inertial properties results in the optimiser favouring maximum leg length. This makes sense as it leads to an increased stroke length and it furthermore reduces build-up of positive angular momentum. However, increasing the leg length without altering its inertial properties does not necessarily correspond to a realistic design due to increased loadings, increased inertia and possible collisions—such as those between the foot and the crossbar when Skippy is fully flexed. As such, the leg is momentarily fixed so that its heel makes contact with the bumper on Skippy's head, and so that the foot does not touch the crossbar when Skippy is fully flexed. The corresponding parameters are $[x_{g,x} \ x_{g,y}] = l_g[\cos \phi_g \ \sin \phi_g] = [500 \ 17.5]$ mm.

Lever and foot The lever is defined by its length l_l and offset angle ϕ_l . Early optimisation studies have shown a desire to minimize this angle, in order to maximize the mechanical advantage when Skippy is flexed and to minimize it when extended. The mechanical design limits the lever angle at $\phi_l = 15^\circ + \phi_u = 35^\circ$, for collision would otherwise occur between the lever, main spring and torso tube for smaller angles (as visible in fig. 6.14a). The resulting transmission ratio of the hip (i.e. from coupler to hip rotation $\partial q_{10}/\partial \phi_{\text{hip}}$) is a factor of 3.11 higher when fully flexed than when fully extended.

Of the kinematic parameters, this leaves the lever length l_l , foot length l_p and toe curvature r_{tx} for optimisation, which are relatively easily changed in the mechanical design, and have been given the following bounds:

$$9 \text{ cm} \leq l_l \leq 10 \text{ cm} \quad (8.20)$$

$$15 \text{ cm} \leq l_p \leq 20 \text{ cm} \quad (8.21)$$

$$20 \text{ mm} \leq r_{\text{tx}} \leq 30 \text{ mm} \quad (8.22)$$

The lever length is estimated based on ball-park figures in section 6.2.2, and its maximum length is limited by mechanical design, as explained in section 6.7.4. The foot length is longer than in [Driessen \(2015\)](#) ($l_p = 12$ cm), as previous optimisation studies have shown that a longer foot is more desirable, and it furthermore permits housing of a bigger ankle spring.

MISCELLANEOUS PARAMETERS

Miscellaneous parameters include that related to the main actuator, ankle spring and hip end stop. The main actuator consists of the main motor, the ball screw and the main spring. As decided by preliminary design decisions, the main motor is a Maxon DCX32L motor, and the ball screw pitch is 4 mm. Parameters of the ankle spring include those defining the compression leaf spring, its attachment points (see fig. 6.16a) and the end stop.

Main spring The main spring model, as defined by eq. (8.8) and eq. (8.9), requires the three parameters $F_{s,2}$, $x_{s,0}$ and d_s to be defined. The spring behaviour is nearly identical if $F_{s,2}$ and $x_{s,0}$ are scaled identically for sufficiently high strain. As such, $x_{s,0}$ is fixed at $x_{s,0} = 8$ cm, leading up to strains of ca. $\epsilon = 1$ (since spring elongation was found not to exceed $x_{s,max} = 8$ cm), leaving only $F_{s,2}$ for optimisation, which has been given the bounds

$$1 \leq F_{s,2} \leq 2.5 \text{ kN}, \quad (8.23)$$

based on an expected maximum actuator force of ca. 2 kN. Spring damping is set to $d_s = 50 \text{ N s m}^{-1}$, corresponding to an efficiency of ca. 90% to 95% during the stance phase.

Hip end stop The hip end stop acts as a parallel spring on the torso-lever joint and is effective only when Skippy is close to its fully flexed configuration. As discussed in section 6.5, the end stop makes contact over a range of a lever-torso angle of $\Delta q_7 = 10^\circ$ from Skippy's flexed configuration, corresponding to a hip angle of $\phi_{\text{hip}} \approx q_6 + q_7 \approx 16^\circ$, as depicted in figs. 6.13 and 8.2. The end stop is assumed to exert a torque of 40 N m when Skippy is fully flexed, which has been chosen in correspondence with the requirement described in section 6.5 which states that it should not exert more than double the torque required for gravity compensation while balancing. The end stop makes contact for a torso-lever angle q_7 of $\phi_{\text{eh},o} = q_{7,o} + 10^\circ$ where $q_{7,o} \approx -8.5^\circ$ is the torso-lever angle for $\mathbf{y} = \mathbf{0}$ that satisfies the loop-closure constraint, i.e. for which $\phi_{\text{hip}} = 0$. At a torque of 40 N m and a rotation of 10° , it follows that the stiffness coefficient k in eq. (8.10) can be calculated as $1.3 \text{ kN m rad}^{-2}$, and that the total energy stored in the spring equals up to a mere 2.33 J. With a designed moment arm of 42 mm, the end-stop delivers up to ca. 0.9 kN.

Ankle spring The ankle end stop is fixed at an angle $\phi_{\text{ea},o} = 45^\circ$ by mechanical design, as mentioned in section 6.6, and offers up to $\Delta q_5 = 60^\circ$ of ankle rotation without being compressed (calculated as $\phi_{\text{ea},o} - q_{5,0}$), which was found to be high enough to not be compressed during normal operation (i.e. making hops instead of crashing).

The contribution of the compression leaf spring to the ankle torque profile can be described by only four independent parameters, such as by the triangle lengths l_b and l_c and the rest angle $\theta_{a,0}$ (depicted in fig. 6.16a) and the spring's arc angle Θ or radius R (see eqs. (7.1) and (7.2)).¹⁰ If spring failure is also taken into account, then an additional parameter is required to describe the ankle spring independently, like the spring thickness t_a .

As explained in section 6.6.3, most of the ankle design parameters have been decided based on previous simulation studies and mechanical design issues (e.g. collisions), namely $l_b = 5$ cm, $R = 180$ mm, $\theta_{a,0} = \pi/2 + \phi_b + \phi_c - q_{5,0} = 90^\circ - q_{5,0}$ and $l_c = l_p - 0.024$. Note that the latter two (l_c and $\theta_{a,0}$) are—despite being predefined and no optimisation parameters—a direct function of optimisation parameters (namely $q_{5,0}$ and l_p). As such, an update function $f_u(x, \text{parameters})$ is introduced that updates parameters that are a direct function of optimisation parameters x and other parameters, which is called at every iteration executed by the optimiser. This function is also used to update inertial properties of bodies based on their change in size, as further explained below.

The remaining ankle spring parameter that is used for optimisation is its width:

$$38 \times (2/3)^3 = 11.3 \leq w_a \leq 38 \text{ mm}, \quad (8.24)$$

with an assumed ankle spring thickness of $t_a = 3$ mm, which is one of the available GFRP thicknesses.¹¹ Remember that only the product $t_a^3 w_a$ is relevant for determining the ankle spring profile. In practice, the spring stiffness can be conserved by choosing a different thickness and width. The upper bound in eq. (8.24) corresponds to the maximum spring width that can fit in the prismatic cavity in the frame, as explained in section 6.6. The lower bound corresponds to a spring with a stiffness that would be equal to that of a spring of the same width and a thickness of 2 mm.

INERTIAL PARAMETERS

Inertial characteristics of body i are fully defined by a mass m_i , radius of gyration about the CoM r_i and coordinates of the CoM $\mathbf{c}_i = [c_{i,x} \ c_{i,y}]$, expressed in their body's local coordinate frame Ψ_i , as depicted in fig. 8.1. The model consists of 10 rigid bodies, of which 3 are fictitious, resulting in $7 \times 4 = 28$ inertial parameters. Note that due to rigid-body connectivity, these 28 parameters are not independent, and many of the parameters only have a marginal effect on the system's dynamics.

The inertial parameters are closely tied to the robot's size, as defined by the kinematic parameters, and are thus no strong candidate for optimisation parameters. This is the result from Skippy's tight mass budget and preliminary design decisions (e.g. placing both motors in the head), which leaves little leeway for moving mass around without changing the robot's kinematics.

The inertial parameters that correspond to the design depicted in fig. 6.1 are listed in table 8.2. While optimising, the parameter update function f_u is used to update inertial

¹⁰Note that the overall spring stiffness can be scaled (i.e. for all y , as defined in chapter 7) by changing the spring width w or thickness t_a (defined as t in chapter 7), but also by equally scaling l_b and l_c and R , which is why only four instead of five independent parameters are required to describe the spring profile.

¹¹The GFRP is available with thicknesses in steps of 0.3 mm (i.e. 2.4, 2.7, 3.0 mm) and 1 mm (i.e. 1, 2, 3.0 mm), which is related to the finite number of fibre layers of which it is composed.

i	body name	m_i [kg]	r_i [m]	$[c_{i,x} c_{i,y}]$ [m]
4	foot	0.175	0.065	[0.012 0.099]
5	leg	0.267	0.207	[-0.239 -0.009]
6	lever	0.248	0.079	[-0.013 0.098]
7	torso	1.481	0.189	[0.485 -0.004]
8	follower	0.057	0.053	[-0.004 -0.043]
9	rocker	0.488	0.063	[-0.010 -0.000]
10	coupler	0.2024	0.121	[0.179 0.000]
	total	2.918		

Table 8.2: Inertial parameters of Skippy, corresponding to the design depicted in fig. 6.1.

parameters of the bodies whose kinematic parameters are optimised, which include the foot and ankle spring. These updates are based on estimates on how change of body size influences change of mass and inertia, for the purpose of improving practicality of the optimisations and to prevent the optimiser from unjustifiably favouring as-large-as-possible bodies.

In particular, the change of foot and spring size require their respective inertia to be updated. This is facilitated by distinguishing between inertial data of the ankle spring and foot excluding the ankle spring. Inertial data of the foot and ankle spring are extracted separately from CAD data, and are merged as follows to calculate inertial parameters of body 4 (which entails the foot including the ankle spring):

$$m_4 = m_p + m_a \quad (8.25)$$

$$\mathbf{c}_4 = \frac{1}{m_4} \begin{bmatrix} m_p & m_a \end{bmatrix} \begin{bmatrix} \mathbf{c}_p \\ \mathbf{c}_a \end{bmatrix} \quad (8.26)$$

$$I_4 = \begin{bmatrix} m_p & m_a \end{bmatrix} \left(\begin{bmatrix} r_p^2 \\ r_a^2 \end{bmatrix} + \begin{bmatrix} \|\mathbf{c}_p\|^2 \\ \|\mathbf{c}_a\|^2 \end{bmatrix} \right) - \|\mathbf{c}_4\|^2 m_4 \rightarrow r_4 = \sqrt{I_4 / m_4}, \quad (8.27)$$

in which the subscripts p and a correspond to data belonging to the ankle spring and foot respectively. Calculations of I_4 follow from the parallel axis theorem. For example, doubling the ankle spring width would double its mass m_a , after which eqs. (8.25–8.27) can be used directly to update inertial parameters of body 4 to high degree of accuracy. Similar calculations are performed for changing the foot and spring length, which other than mass also require their coordinates of the centre of mass and radius of gyration to be scaled. In the presented design (depicted in fig. 6.1), the foot has a nominal length of $l_p = 0.184$ m.

IMPLEMENTATION

With the optimisation parameters and their bounds defined in eqs. (8.12–8.24), the tuple of optimisation parameters \mathbf{x} can be defined as

$$\mathbf{x} = [t_{\text{rise}} \quad t_{\text{sat}} \quad t_{\text{fall}} \quad c_{\text{rise}} \quad c_{\text{fall}} \quad q_{4,0} \quad q_{10,0} \quad q_{5,0} \quad l_l \quad l_p \quad r_{\text{tx}} \quad F_{s,2} \quad w_a]^T.$$

Note that some parameters in \mathbf{x} are on the order of 1×10^{-2} (e.g. command times), whereas other parameters are on the order of 1×10^3 (e.g. $F_{s,2}$). To give all parameters

equal importance for optimisation, each of them have been normalized so that their respective lower and upper bounds take the values 0 and 1 respectively. The termination tolerance on the normalized parameters (StepTolerance) is set to 1×10^{-4} .

\mathbf{x} has furthermore been implemented in such a way that optimisation parameters can be easily added or removed, which has been done to facilitate locking physical parameters along the process of iterative design and to check whether hardware that is optimised for making a somersault is able to hop straight-up and vice versa. The latter is done by allowing only command signal parameters $\mathbf{x}_1 = [t_{\text{rise}} \ t_{\text{sat}} \ t_{\text{fall}} \ c_{\text{rise}} \ c_{\text{fall}} \ q_{4,0} \ q_{10,0}]^T$ to be optimised¹², while assuming the hardware parameters from a previous optimisation (i.e. by locking $\mathbf{x}_2 = [q_{5,0} \ l_1 \ l_p \ r_{\text{tx}} \ F_{s,2} \ w_a]$, which requires their removal from the optimisation parameters).

8.4. RESULTS AND REVISIONS

Since the optimisation algorithm is not global, the outcome may be dependent on the initial estimate of \mathbf{x} . The initial estimate is taken from a previous similar simulation result. Due to sparsity of feasible points (as caused by the chaotic nature of the system), the initial estimate is not necessarily feasible either, for which the optimisation was regularly observed to converge to an infeasible point as well. A solution is said to be feasible if the constraints (as described in section 8.3.1) are satisfied within a specified constraint tolerance of 1×10^{-4} . If this happened (or if there was sufficient suspicion that the optimiser converged to a small local minimum), then the optimisation was repeated for another initial estimate of \mathbf{x} , until it converged to a feasible point.

8.4.1. INITIAL RESULTS

Results and a discussion of optimisations corresponding to the design presented in chapter 6 and parameters listed in table 8.2 are described below.

8

RESULTS

In optimising for either ordinary hopping ($p_{r,\text{end}} = 0$) or somersaulting ($1 \leq p_{r,\text{end}} \leq 1.5$), the optimiser had indicated a strong preference for somersaulting, resulting in $p_{r,\text{end}} = 1.5$ and a corresponding jump height of $h = 3.50\text{m}$. The resulting system was found to be able to hop to a height of only $h = 1.77\text{m}$ for $p_{r,\text{end}} = 0$ (by optimising for only the initial conditions and command signal parameters \mathbf{x}_1 , as described in section 8.3.3). An optimised system for $p_{r,\text{end}} = 0$ was found to be able to reach a height of $h = 1.89\text{m}$, which is only marginally higher than 1.77m , and was still able to reach a height of 3.04m when somersaulting (for $p_{r,\text{end}} = 1.5$). The results of the optimisation can be found in table 8.3.

The difference between results for ordinary hopping and somersaulting is remarkably big. It is bigger than that from early simulation studies using tentative physical parameters similar to those presumed in [Driessen \(2015\)](#). In particular, the result for ordinary hopping is unsatisfactory, and retrying for different initial conditions did not improve the result.

¹²Note that $q_{5,0}$ is not a command signal parameter, as it is directly related to the ankle spring's rest configuration.

$p_{r,end}$	lb	original system		updated system		ub
		1.5	0	1.5	0	
t_{rise}	0.02	0.02	0.02	0.02	0.02	0.12
t_{sat}	0.05	0.154	0.0929	0.1532	0.0602	0.20
t_{fall}	0.01	0.01	0.0109	0.0890	0.12	0.12
c_{rise}	-1	1	1	1	1	1
c_{fall}	-1	-1	-0.9973	-0.469	0.5457	1
$q_{4,0}$	-0.47	-0.2625	-0.3552	-0.3858	-0.3827	-0.25
$q_{10,0}$	0.02	0.0357	0.0442	0.0346	0.0336	0.075
$q_{5,0}$	-0.1745	-0.1745	-0.1745	-0.1798	-0.1859	-0.3491
l_l	0.09	0.10	0.09	0.10	0.10	0.10
l_p	0.15	0.152	0.1558	0.15	0.1597	0.20
r_{tx}	0.02	0.03	0.0297	0.03	0.024	0.03
$F_{s,2}$	1000	2000	1648	2000	1999	2000
w_a	0.0113	0.0113	0.0113	0.0153	0.0197	0.038
h for $p_{r,end} = 0$		1.77	1.89	2.19	2.33	
h for $p_{r,end} = 1.5$		3.50	3.04	3.98	3.84	

Table 8.3: Optimisation results for ordinary hopping ($p_{r,end} = 0$) and somersaulting ($1 \leq p_{r,end} \leq 1.5$). The optimisations indicate a strong preference for somersaulting. All values are in SI base units. The lower bounds (lb) and upper bounds (ub) of x are also included.

DISCUSSION

A possible explanation for the big difference in jump height is the big difference in inertial properties of Skippy's upper body and leg, which was underestimated in earlier studies: the upper body has substantially more inertia than the leg, and its CoM is located further from the hip (towards the head) than initially estimated. As explained in [Driessen \(2015\)](#), for a simple 2R linkage mechanism, an upper body with more inertia than the lower body generally leads to more build-up of positive angular momentum. The effect can be countered by early initiating steering with the ankle, albeit at the expense of jump height. A possible solution is increasing the size of the leg or reducing the size of the torso. This theory is supported by rerunning optimisations while allowing for also the leg length (l_g) or torso length (described by $c_{6,x}$) to be optimised, which have shown higher ordinary hopping heights for a longer leg or shorter torso (reduced $c_{6,x}$). However, implementing a shorter torso in the current design is not straightforward, as for instance it would lead to collisions of the main actuator's guiding mechanism (see section 6.7). In addition, for somersaulting the jump height decreases for a shorter torso. This makes sense from the points on which original design decisions were made, as explained in section 3.1: for good hopping performance, the robot's CoM is preferably as far as possible away from the hip axis to maximize the stroke length. More specifically, early design decisions specified that a majority of the torso's mass should be at the head and that another part of its mass should be at the hip, but that its middle portion be as light as possible, in order to obtain a good compromise between balancing and hopping performance. These specifications have been successfully adhered to in making the design that is presented in chapter 6, and due to Skippy's light weight budget, it is not

feasible to make significant changes to its weight distribution unless preliminary design decisions would be overhauled.

In summary, substantially changing Skippy's kinematic and inertial properties (1) possibly leads to a design with reduced balancing and overall jump performance (permitting somersaults), (2) requires an overhaul of preliminary design decisions, introducing a lot of delay, which is not in line with the trial-and-error philosophy presented in chapter 2.

A compromise is found by only applying marginal changes to the design, which mainly involves an increased leg length.

8.4.2. UPDATED DESIGN DESCRIPTION

The new design features an extended leg that is 8 cm longer. Despite early considerations against increasing the leg length, it has been made possible as follows.

- An extension of 8 cm enables usage of Skippy's toroidal bumper as the foot end stop, so that it also makes contact at the heel, and so that the original foot end stop can be removed (see fig. 8.4). The added weight of the increased leg length and required crossbar shaft fortifications are nearly compensated for by the removal of the dedicated foot end stop.
- FEA have indicated that the leg frames of the longer leg are still strong enough to withstand the increased loadings.
- The leg length was initially limited partly because the foot would otherwise collide with the crossbar. Substantially increasing its length would require a lengthy and thus feeble crossbar shaft, leading to a heavier and less robust design. However, the optimisations have shown a desire for a shorter foot and a smaller ankle rest angle, which increases the gap between the foot and crossbar. For an extension of 8 cm, the crossbar shaft still requires to be extended to prevent collisions, but only by 3.75 cm. The marginal extension of the crossbar shaft does not require significant fortification on the torso head.

The lower leg frame and foot have also undergone small changes. Both optimisations have indicated desire for a smaller foot, see table 8.3. As such, the nominal foot length has been reduced to $l_p = 16$ cm (previously $l_p = 0.184$ m). The lower leg frame has been changed so that the ankle joint is now in line with the leg tube rather than it being above it (essentially removing the described kink in section 6.6), and has thus been moved down. The design corresponds to the one depicted in figs. 6.15 and 6.16. This change allows for (1) an increase of the ankle spring rest angle $\theta_{a,0}$ (see also section 6.6.3), (2) lowering the heel in accordance with the lower end stop, and (3) lowering the now smaller foot so that the crossbar cannot make contact with the ground for $\mathbf{y} = \mathbf{0}$. The new ankle spring rest angle is $\theta_{a,0} = \pi/2 + \phi_b + \phi_c - q_{5,0} = 100^\circ - q_{5,0}$ (previously $\theta_{a,0} = 90^\circ - q_{5,0}$). The new leg size is defined by $[x_{g,x} \ x_{g,y}] = l_g[\cos\phi_g \ \sin\phi_g] = [580 \ 5]$ mm (previously [500 17.5]).

Inertial data of the updated system can be found in table 8.4. A right-side view of the original and updated design are depicted in fig. 8.4 for $\mathbf{y} = \mathbf{0}$. Note that the updated system is only 16 g heavier.

i	body name	m_i [kg]	r_i [m]	$[c_{i,x} c_{i,y}]$ [m]
4	foot	0.175	0.064	[0.012 0.097]
5	leg	0.286	0.242	[-0.269 -0.000]
6	lever	0.248	0.079	[-0.013 0.098]
7	torso	1.480	0.193	[0.491 -0.003]
8	follower	0.057	0.053	[-0.004 -0.043]
9	rocker	0.488	0.063	[-0.010 -0.000]
10	coupler	0.202	0.121	[0.179 0.000]
	total	2.936		

Table 8.4: Inertial parameters of Skippy, corresponding to the updated design depicted in fig. 8.4b. Unaffected bodies are in grey.






Figure 8.4: Skippy's final design update only features minor changes: it mainly concerns an increased leg length, smaller foot and removal of the original dedicated foot end stop. Areas with change are encircled in red.

8.4.3. UPDATED DESIGN RESULTS


Optimisation results of the new design are also listed in table 8.3. For ordinary hopping and somersaulting, the updated design was found to hop to heights of $h = 2.19\text{m}$ and $h = 3.98\text{m}$ respectively when optimised for somersaulting, and to heights of $h = 2.33\text{m}$ and $h = 3.84\text{m}$ respectively when optimised for ordinary hopping. As can be seen in table 8.3, these results are all better than the previous results. The most prominent increase in performance is that for ordinary hopping, whose results are now satisfactory as the jump height exceeds 2 m. In favour of better performance for ordinary hopping, the results that are adopted for Skippy's final design are those of the latter. These results correspond also to the design depicted in fig. 8.4, comprising a foot length $l_f = 16\text{cm}$ and toe radius $r_{\text{tx}} = 24\text{mm}$. Results of corresponding simulations are plotted in fig. 8.5. For jumping a somersault ($p_r = 1.5$), it can be observed from fig. 8.5c that voltage is in saturation for most of the stance phase, which contributes to the high jump height. In the beginning of the stance phase of both simulations, the voltage is clipped so as to prevent the nut velocity from exceeding its maximum velocity of 0.4m s^{-1} , which can be seen in fig. 8.5d. The current has not saturated, but has reached a maximum of 49.3A , which is




(a) The GRF is shown to be within a 90° friction cone ($|\tau_1/\tau_2| < 1$), except shortly after landing and shortly before take-off.



(b) Angular momentum build-up of Skippy depicted for jumping straight up and for making a somersault.



(c) Voltage command signal U_c as defined by eq. (8.11) and the actual input voltage U , clipped by the limiting nut velocity \dot{x}_n (shown in fig. 8.5d).



(d) Nut velocity \dot{x}_n , causing the input voltage U to be saturated so as to not exceed the maximum permissible speed of 0.4 m s^{-1} .

Figure 8.5: Simulation results of the updated design, corresponding to hardware optimisation for jumping straight up ($p_r = 0$). Results are shown for Skippy jumping straight-up ($p_r = 0$) and for Skippy making a somersault ($p_r = 1.5$).

close to the 50 A limit. Furthermore, the maximum force on the actuator is 1899 N, which is close to the anticipated maximum of 2 kN for which the actuator has been designed.

CHECKS FOR CORRECTNESS

Various measures have been taken to check for correctness and validity of simulations. This includes calculation of the energy audit to check for energy flow and leakage in the system, following the approach in [Driessen \(2015\)](#). At all times, the potential and kinetic energy of the system plus work done by the motors minus the heat loss should be constant. For the full duration of the simulations, the maximum deviation of energy was observed to be only 1 mJ, which is attributed to numerical integration, and suggests that the system has been solved with sufficient precision. Other checks are related to physical constraints of the system that are not included in the constraint function discussed in section [8.3.1](#). These are the following:

- It is checked for the full duration that the GRF does not go out of a 90° friction cone, i.e. that $|\tau_1/\tau_2| < 1$, as the robot would have otherwise slipped. This is true for most of the stance phase, as can be seen in [fig. 8.5a](#). The ratio is exceeded only shortly after landing and shortly before take-off, which is due to the low vertical reaction force. As explained in section [8.3.1](#), this is likely not problematic in reality.
- The generated voltage command signal (see [eq. \(8.11\)](#)) is defined to be continuous so that it is generatable by the motor in reality, as motor induction is not included in the model. However, bear in mind that the actual command signal is also the result of current and velocity saturation. It is checked that these saturations do not cause the command voltage signal to be infeasible.
- It is checked for the full duration that $q_{4,0} \leq 75^\circ$, as the foot otherwise start rolling on its back.
- It is checked for the full duration that the robot does not collide with the ground (e.g. as possible at the hip).
- It is checked for the full duration that the foot end stop is not compressed by more than several degrees. The latest results have shown that it is not compressed at all.
- It is checked that the internal strain of the ankle curved compression leaf spring is not unrealistically high, as it would cause the ankle spring to break. [Chapter 7](#) assumes a compressive strain limit of 2.1%. Note however that real experiments will have to be done to identify the actual strain limit, as it is a function of the fibre quality, which could be decisive for the choice of the final spring thickness and the choice between cheaper E-glass fibres or the more expensive S-glass fibres.

Of these checks, only the latter did not pass. At an ankle rotation of $\Delta q_5 = 60^\circ$ (corresponding to end stop compression of $15^\circ + q_{5,0} \approx 5^\circ$), the maximum compressive strain can be calculated as 2.62% using [eq. \(7.13\)](#), which is higher than the reported permissible maximum of 2.1%. However, this value can be reduced by choosing a thinner and wider spring, since the optimiser had returned an optimal spring width of only $w_a = 19.7$ mm whereas the frames offer space for a spring of up to $w_a = 38$ mm. Reducing the spring thickness from $t_a = 3$ mm to $t_a = 2.4$ mm (which is an available thickness), requires the spring to have a width of $w_a = 19.7 \times (3/2.4)^3 \approx 38$ mm¹³ to conserve the spring profile, and decreases the maximum strain to $\epsilon_{c,max} = 2.4/3 = 2.09\% < \epsilon_{xc} = 2.1\%$. Note that

¹³Coincidentally, this value corresponds to the maximum permissible width.

this does not leave a safety margin, but the presumed value of $\epsilon_{xc} = 2.1\%$ is highly uncertain, as explained in chapter 7: ϵ_{xc} might be higher or lower depending on the quality of the GFRP and its layering. Experiments will have to point this out. If it is lower, then Skippy can either be equipped with a weaker spring (which is slightly sub-optimal, but still functional), or additional research could be done in optimising the shape of the spring, which can boost the energy density of the spring by up to an estimated factor of two, as explained in section 7.5.

8.5. DISCUSSION AND CONCLUSION

This chapter has presented single-objective (i.e. maximizing jump height) optimisation studies and an accompanying planar model of Skippy, and has resulted in finalisation of Skippy's design. The optimisation studies focus on the stance phase preceding Skippy's highest hop, which is Skippy's most demanding behaviour, and have shown that Skippy can realistic straight-up hop to a height of $h = 2.33$ m and to somersault up to a height of $h = 3.84$ m, which are higher than the jump height of any other continuous-hopping robot or animal. The simulations were furthermore based on a target height of at least 3 m, by assuming a previous hop to be 1.5 m high. Whereas it is not a necessity, the jump height can likely be further increased by landing from a higher previous hop, which could be a subject of future studies.

To reduce the complexity of the design process, the optimisation problem has been simplified and the design is done iteratively. In iterative design, components and parameters are finalized successively, rather than all being globally optimised at once. The solution is not guaranteed to be globally optimal for the fact that many parameters (or their bounds) have already been fixed based on previous iterative design steps. However, referring to the design philosophy as presented in chapter 2, this is not necessarily a bad thing, as the design should be good enough rather than optimal.

In addition, the optimisation problem has been simplified to being a single-objective optimisation and is only 2D. Since the cost function is defined to maximize jump height, the solution does not guarantee that Skippy is also good at balancing, or at any behaviour other than making high jumps. It is nevertheless advised to build the presented design of this thesis. In line with the design philosophy governed by trial-and-error, this will assist in finding Skippy's shortcomings in reality.

Follow-up studies are needed to investigate the topic of Skippy's combined capabilities of hopping and balancing, which are currently carried out by other PhD students in the Skippy Project. One possibility is performing multi-objective optimisations, as they are targeted to optimise a repertoire of behaviours (e.g. jumping, balancing, etc.). However, the implementation and optimisations require more time as they are more complex. The optimisation time can nevertheless be improved significantly by switching to a more efficient computational environment than Matlab and Simulink, and by switching to a more efficient optimisation algorithm (that is based on a global, non-gradient-based method). These tasks could be carried out in parallel to experimentation on the first built version of Skippy.

Defining, implementing and understanding the optimisations have required several design iterations, which have spanned several months. The studies presented in this chapter have focused only on the final iteration that leads to the final design, and have

been included partly to be exemplary for the encountered difficulties, despite the relative simplicity of the optimisation problem. The problem is relatively simple because it concerns (1) a robot with few (actuated) joints, conveyed by (2) a simplified model of that robot (i.e. it is 2D), covering (3) only one of its behaviours in isolation (single-objective optimisation), assuming (4) most of its parameters to be fixed. It illustrates how much more complex an optimisation problem can be of a full-fledged humanoid if it would also be designed for high performance by machine and behaviour co-design, design at the performance envelope and holistic design. A problem like that could be no longer tractable, which could explain why a lack of high performance in more complex robots is observed.

9

DISCUSSION, CONCLUSION AND FUTURE DEVELOPMENT

9.1. DISCUSSION

HIGH PERFORMANCE

This thesis has focused on the design of the hopping and balancing robot Skippy, as a case study of a design philosophy that promotes simple but physically high performance design. The project has been triggered by the notion that legged robots are underperforming compared to their biological counterparts, despite the more powerful building blocks technology has had to offer for already multiple decades. The design philosophy claims that this paradox is best explained by over-complexity in robot design, and is furthermore governed by the ten key points listed in section 2.4.2.

In particular, high performance can be achieved by aiming for (1) holistic design, (2) design at the performance envelope and (3) machine and behaviour co-design. Abiding these strategies is a daunting task for designing complex robots; the Skippy project has illustrated that this can be a challenging task even for a robot that is intentionally simple. Nevertheless, it has led to results that indicate that legged robots can indeed be designed to outperform animals in terms of physical skills. In particular, the thesis has resulted in a realistic design of Skippy, capable of reaching jump heights of 3.8 m while somersaulting or 2.3 m when performing straight up hops.

The design is holistic, because all of its individual parts are intimately interconnected and explicable only by reference to the whole. This has been facilitated by applying a method of iterative design. Whereas this implies that some of its components have been chosen on a presumptive basis, revisions have been done where necessary, such as that of the main motor and ball screw pitch in chapter 6.

The design is targeted on behaviour at the performance envelope, as its natural dynamics are designed to coincide with the most extreme behaviour it is expected to deliver. This has been tackled by analysing its behaviour in physical saturation (i.e. saturation defined by current limits, voltage limits, speed limits, force limits, etc.), rather than

by analysing its closed-loop behaviour as affected by control laws and imposed torque limits. Both the actuator design criterion introduced in chapter 4 and the optimisations in chapter 8 are an example of this.

The design is also an example of machine- and behaviour co-design, because only its final action has been specified in advance—namely, reaching a specified jump height—not the action sequence to achieve it. Both the system and action sequence are found coherently. Co-optimisation of both motion control parameters and hardware parameters in chapter 8 is a direct example of this. More general, the iterative design strategy between mechanical design and simulations are also an example of machine- and behaviour co-design. A disadvantage of step-wise and iterative design is that it is unlikely that a global optimum is being tracked. However, this should not be required if the results are good enough.

Many of the design decisions are technology-inspired rather than bio-inspired solutions, which have also contributed to the high performance design. An example of this is the preliminary design choice of using a crossbar, which does not have a joint rotation limit and therefore allows for high balance-offset angles. The use of a four-bar linkage to obtain a desired non-constant transmission ratio is another example of technology-inspired design.¹ Another example is the search for spring materials that excel in high gravimetric energy storage, which have led to the use of rubber and GFRP.

PROGRESS

The design philosophy described in this thesis also focuses on improving progress in the development of high-performance robots. The strategies of holistic design, design at the performance envelope and machine and behaviour co-design are counter-productive in this regard, as they slow the design process down. The optimisation problem described by chapter 8 is exemplary of this, as it requires a lot of time despite its relative simplicity: the cost-function is single-objective, the model describes a simplified (2D) robot that only has few DoFs/actuators, and it focuses only on optimising a reduced number of parameters.

The design philosophy promotes iterative design by trial-and-error, designing for reduced risk management, physical robustness, and for the robot to be good enough rather than optimal, which are all believed to improve progress in robot design and to facilitate design for high performance. Tippy—the balancing-only precursor of Skippy—has been a direct consequence of designing for trial-and error, and both robots have been designed with reduced risk management, to be physically robust, and good enough rather than optimal. They have furthermore been designed not to be bigger than necessary, by having a size and weight limited by the chosen technology and performance target. Compared to for instance humanoids, both robots are small and lightweight at ca. 3 kg, which have reduced the need for risk management and improved their physical robustness. An even smaller robot would have arguably been more fruitful regarding progress, judging from the progress of Salto (the hopping robot designed by UC Berkeley). Nevertheless, the Salto project has been ongoing for a longer time and has a greater history of building small robots. In addition, their different scale uses technologies that are not well scalable to robots of bigger (e.g. humanoid) size, whereas Skippy is based on the

¹Coincidentally, the kinematics of the four-bar linkage were found to be similar to that of a biological knee.

type of technologies (e.g. harmonic drives, milled frames, motor and bearing types) that are also used in those robots.

Design choices have typically been made using ball-park figures and simplified simulations up to the point of obtaining fulfilling expectations of their feasibility, in order to accelerate decision making. Design decisions whose feasibility is deemed less certain have been accompanied for by suggestions for possible alternatives or easy alterations. This procedure can be described as an educated-risk design strategy, which blends in well with the trial-and-error design philosophy. For example, it is worth noting that a likely possible culprit of poor balancing performance is too low a stiffness of either of the springs in Skippy, which can be fixed in practice by increasing regressiveness of respective spring profile. For the ankle spring, this only requires its replacement for a spring with a bigger radius. For the main spring instead, a possibility is modifying the main actuator to include an end stop, as explained in section 6.7.3.

Despite the proposed measures to improve the design progress, Skippy does not yet exist. The development of Tippy has—regardless of the intention of improving progress by learning through experience—significantly contributed to delays. Nevertheless, the Tippy project has led to substantial contributions and experience, and as such has been very useful as a pre-design phase in preparation for the design of Skippy. For some design decisions it has been challenging to find a productive balance between spending time on exercising rounds of trial-and-error and spending time on making elaborate studies to make the better choices for one final design. Similar trade-offs have been experienced in the exercising of iterative optimisations. The optimisation problem becomes progressively more time consuming when focusing on a broader problem. It can be accelerated by staging it in iterations and by making design assumptions, at the greater risk of having to fall back to major design revisions or to give in on performance.

Another possible explanation for the encountered delays is that the project was so ambitious. Despite the intention of settling for a design that is merely good enough, the high objectives and expectations of Skippy have pushed to new theories and technologies and a design that is likely better than necessary. For instance, it has been a struggle to adhere to a mass budget that is significantly lower than 3 kg; the original target mass of 2 kg in combination with a 4 m jump height have proven to be intractable. If, for instance, the target was set to reach only half the jump height (2 m), the problem would have been much more tractable and the project would have likely accelerated. A continuous-hopping robot reaching heights of 2 m is still impressive as it outperforms other robots and animals. Another ambitious requirement of Skippy is that it was not allowed to have more than two actuators, which has required the development of new design solutions to synergize hopping with steering and balancing (i.e. the passive ankle joint). Significant amounts of time have been spent in an attempt to improve its limited ability to steer, which might have been avoidable if it had been considered to add an extra actuated joint and body. Ironically, the argument for using only two actuators was that of simplicity, as it is the minimum number of actuators required to control a robot in full 3D. However, it has been proven by Salto-1P that a high-performance robot with more actuators can be simpler, which uses additional actuators (i.e. the thrusters) to substantially ease the balancing and in-flight control problem.

Nevertheless, it should be noted that if it were not for Skippy's ambitious objectives,

many of the theories and ideas useful for high-performance robot design would have likely never been found. Such contributions include (1) the actuator design criterion (chapter 4), (2) a generalization of the centre of percussion (chapter 6) and (3) curved compression leaf modelling (chapter 7), which are a direct product of the thesis. Previous such contributions of the project include the invention of the ring screw and new balancing algorithms.

9.2. CONCLUSION

Through the case study design project of the hopping and balancing robot Skippy, it is shown that proper design for high physical performance (1) can indeed lead to a robot design that is capable of physically outperforming humans and animals and (2) is already very challenging for a robot that is intended to be very simple. In particular, the robot design can reach a record jump height of 3.8 m in realistic simulations. The endeavour for holistic design, design at the performance envelope and machine and behaviour co-design have contributed to the high performance, but have also challenged the design procedure, and the overly ambitious objectives have further complicated it. During the development, many complications were encountered in the decision-making progress of the iterative design strategy, involving trade-offs between exercising trial-and-error and elaborate simulation studies, but they have also lead to the development of new theories. These new theories include an actuator design criterion, a generalization of the centre of percussion and modelling of curved compression leaf springs. These are contributions of the thesis in addition to the realistic design of Skippy—exemplary for light-weight, physically robust and affordable design—and the design, creation and demonstrations of Tippy—designed to be customizable and educational. As such, it could be concluded that the thesis project has been productive, and that the design philosophy has paid off.

9.3. FUTURE DEVELOPMENTS

As discussed in section 8.5, it is advised to proceed with building Skippy, despite the fact that optimisation studies have only been single-objective. This will assist in finding Skippy's shortcomings in reality, which is in line with the design philosophy governed by trial-and-error. Possible limiting factors on the performance, such as spring stiffness, are easily altered in practice.

The design of Skippy, as presented by chapter 6 and finalized in chapter 8 is nearly ready to be sent off to the mechanical workshop for production, lacking mainly the required technical drawings. Exact dimensioning of the ankle spring, bumper sponge (regarding the cutting pattern and filling density), their caps and main spring are best finalized through experimentation, including (primarily) material identification and (secondarily) jumping experiments, abiding the design strategy of trial-and-error.

To validate the results that claim jump heights of up to 3.8 m in chapter 8, the first jumping experiments should focus on the stance phase in preparation for the highest jump. Note that the simulations do not use feedback control, whereas the system is highly chaotic in reality: during stance, the robot experiences ca. 10g downward acceleration, making it susceptible to tipping over. Nevertheless, the stance phase lasts

only 0.2 s, so the robot does not have that much time to tip over. It is expected that with sufficient trials, the stance phase can be reproduced in reality with only feed forward control, by dropping it from a predetermined height corresponding to Skippy's initial conditions (ca. $h = 1.5$ m). In parallel, control strategies that combine 3D hopping and balancing should be worked out in greater detail, as required for making a sequence of hops.

For the longer term, revisions for a future version of Skippy are to be made based on the experimental results and new developments of control theories. Some considerations are already listed below.

- The resulting design is heavier (3 kg) and slightly bigger than originally anticipated in [Driessen \(2015\)](#); [Featherstone \(2014\)](#) (2 kg). The weight of a possible future versions is preferably reduced, for which the most obvious solution is reducing Skippy's overall size. However, Skippy's original sizing was also based on minimizing size and weight, and was largely limited by preliminary choices of the used technologies in the robot. A major influence is the speed limit of the ball screw. As explained in section 6.2 and appendix A, a robot that is subject to a speed constraint suffers from substantial performance reduction if it is downscaled. The recent developments of the ring screw and availability of small roller screws rid us of the speed constraint, and subsequently open doors to downscaling the robot. Nevertheless, downscaling the presented design remains difficult: many of the selected off-the-shelf component types already have (near-)minimum size, including some of the bearing types, electronics, batteries and the harmonic drive, none of which can be simply scaled down. Some of the wall thickness of milled structural components are minimal, which might imply that downscaling requires resorting to different design style or machining technology. Skippy's weight could also be reduced using brushless motors.
- As discussed in section 6.2, using brushless instead of brushed motors (both for the main and swivel motor) could further boost the performance of Skippy. In the presented work, brushed motors were chosen mainly due to time constraints. Brushless motors generally have a better power-to-weight ratio and are furthermore available as rotor/stator units, so that the motor shaft and housing can be designed to be an integral part of the frame and their respective transmission shafts, thereby saving significant weight. In addition, brushless motors are available with a higher torque density, which offers more solutions for transmissions with lower gear ratios, which is of especial importance for the swivel actuator.
- Skippy makes use of a parallel spring in the hip, which boosts its jump height. In the presented design the spring stores 2.3 J when Skippy is fully flexed. Energy storage is limited because it covers only a small portion of the hip range $\Delta\phi_{\text{hip}}$. In a future design, this end stop could possibly be upgraded to a spring that covers Skippy's full hip range, providing it does not significantly hinder Skippy's other behaviours, such as in-flight behaviour or balancing. Simulation studies will have to point this out. However, even if the spring provides only the force required for gravity compensation for a balanced configuration, it would store up to ca. 15 J, which is a significant improvement over the currently installed end stop. A spring

that covers the full range cannot be implemented by making use of a (compressive) rubber end stop. A possible design solution is the use of a rubber tensile spring that is spanned between the lever-leg joint and follower-torso joint.

- Like in Salto-1P, it could be considered to add an extra actuated joint to Skippy, providing it indeed leads to substantially simplified design and control. It should also be investigated whether such an actuator would not excessively spoil the performance by adding the required additional weight. An example of such an actuated joint could be a control-moment gyroscope, a tail like in Salto-1P (essentially a crossbar in the sagittal plane), or thrusters, each of which would be able to alter Skippy's orientation and rotational velocity in flight.
- Skippy's GFRP ankle spring is now defined by simple shapes, comprising a constant width and thickness. For the purpose of further reducing weight and offering possibilities for housing a stronger spring in Skippy's ankle (if so desired), its shape can be optimised by introducing a width and thickness that vary over the length of the spring, as explained in more detail in section 7.5.
- Optimisations can be expanded to include more optimisation parameters and to make fewer assumptions. In particular, multi-objective optimisations can be used to come to a design that offers a desired trade-off between different types of hopping and balancing behaviours. These optimisations will be especially fruitful if Skippy's 3D balancing and gait switching will be better established.

A

JUMPING AND SCALING




Figure A.1: A simple jumper's actuator scaling visually explained.

This appendix includes derivations on how robot and animal scaling affects the jump height using simple assumptions on scaling laws, based on different actuator limits. The actuator limits are either a force limit, velocity limit or a power limit (a combined force and velocity limit). The actuator here refers to the actuator (or muscle) that is responsible for preparing the jump, and exerts a vertical force F (either through some type of transmission or on its own) which allows the robot or animal to gain height and jump. The jump height of the robot is h and the total height is $h_{\text{tot}} = h + s$, following the definitions in table 2.1.

The simple scaling laws assume that if a robot or animal scales with size S , its stride length s will scale with S too, and its mass m with S^3 . It also assumes that the actuator scales equal to the size of the robot, so that its maximum output force F_{max} scales with S^2 and its maximum velocity \dot{z}_{max} with S . A visual representation of these assumptions on scaling can be seen in fig. A.1.

Variables of the scaled robot are denoted by an apostrophe (i.e. \square'), e.g. $m' = mS^3$.

FORCE LIMIT

In the first case, it is assumed that the maximum output force $F = F_{\max}$ is the limiting factor. The actuator pushes with maximum force $F = F_{\max} > G = mg$ over a stride length s , after which the robot lifts off, so that the work done equals $W = \int_0^s F_{\max} dz = F_{\max}s$, and the resulting potential energy at apex is $P_a = mg(h + s) = Gh_{\text{tot}}$. By solving $W = P_a$ for the total jump height we find

$$h_{\text{tot}} = \frac{W}{G} = \frac{F_{\max}s}{G},$$

of which we find that it does not scale:

$$h'_{\text{tot}} = \frac{F_{\max}S^2sS}{GS^3} = h_{\text{tot}}S^0 = h_{\text{tot}}.$$

It follows that the jump height decreases for increased robot size only due to the increased stride length s :

$$h' = h'_{\text{tot}} - s' = h_{\text{tot}} - sS.$$

VELOCITY LIMIT

If we instead assume actuator velocity \dot{z}_{\max} to be the limiting factor, which would be the case for small jumpers under the assumption that $\dot{z}'_{\max} = \dot{z}_{\max}S$, then the jump height follows directly from the energy balance at take-off and apex. We find:

$$h = \frac{\dot{z}_{\max}^2}{2g} \tag{A.1}$$

$$h' = \frac{\dot{z}_{\max}^2S^2}{2g} = hS^2 \tag{A.2}$$

POWER LIMIT

Suppose now that, by using appropriate and perfect transmissions and/or power modulating springs, the mechanical power output $\dot{W} = F\dot{z}$ of the actuator is the only limiting factor, so that the actuator power output equals $\dot{W}_{\max} = F_{\max}\dot{z}_{\max}$ at all time. It follows that $\dot{W}'_{\max} = S^3\dot{W}_{\max}$.

The energy balance at any point in time t during the thrust phase reads

$$W = K + P$$

$$\dot{W}_{\max}t = \frac{1}{2}m\dot{z}(t)^2 + mgz(t)$$

If we neglect the gravitational potential energy (mgz) during the thrust phase¹, we have:

$$\dot{W}_{\max}t \approx \frac{1}{2}m\dot{z}^2 \tag{A.3}$$

¹This essentially ignores the difference between h_{tot} and h , which is acceptable for small jumpers since the stride length is significantly shorter than the jump height.

Rewriting gives

$$\dot{z} \simeq \sqrt{\frac{2\dot{W}_{\max} t}{m}} \quad (\text{A.4})$$

Integrating gives

$$\begin{aligned} \int_0^z dz &\simeq \sqrt{\frac{2\dot{W}_{\max}}{m}} \int_0^t \sqrt{t} dt \\ z &\simeq \sqrt{\frac{2\dot{W}_{\max}}{m}} \frac{2}{3} t \sqrt{t} \\ \rightarrow t &\simeq \left(\frac{3}{2} z \sqrt{\frac{m}{2\dot{W}_{\max}}} \right)^{2/3} \end{aligned} \quad (\text{A.5})$$

By filling eq. (A.5) in eq. (A.4) we find

$$\dot{z} \simeq \sqrt{\frac{2\dot{W}_{\max}}{m}} \left(\frac{3}{2} z \sqrt{\frac{m}{2\dot{W}_{\max}}} \right)^{1/3} = \left(\frac{3}{2} z \frac{2\dot{W}_{\max}}{m} \right)^{1/3} = \left(\frac{3z\dot{W}_{\max}}{m} \right)^{1/3} \quad (\text{A.6})$$

At take-off we thus have

$$\dot{z}_{\text{end}} \simeq \left(\frac{3s\dot{W}_{\max}}{m} \right)^{1/3} \quad (\text{A.7})$$

The jump height can now be calculated as follows:²

$$h = \frac{\dot{z}_{\text{end}}^2}{2g} \simeq \left(\frac{3s\dot{W}_{\max}}{m} \right)^{2/3} \frac{1}{2g}$$

and scales as follows:

$$h' = \frac{\dot{z}'_{\text{end}}{}^2}{2g} \simeq \left(\frac{3SsS^3\dot{W}_{\max}}{S^3m} \right)^{2/3} \frac{1}{2g} \simeq S^{2/3} h.$$

This suggests that, even for proper transmissions, smaller jumpers jump less high.

²Note that this equation is a factor of $3/2^{2/3}$ off from the one derived in [Armour et al. \(2007\)](#), in which the incorrect assumption is made that acceleration is constant. However, the scaling factor $S^{2/3}$ is identical, which is the only relevant result.

B

TRIANGLE FUNCTIONS

The functions $\gamma_{\bar{\Gamma}}$ and $\gamma_{\hat{\Gamma}}$, referred to as “triangle functions”, are loop closure functions of a simple triangle mechanism that has one edge of variable length and three revolute joints, as depicted in fig. B.1. The triangle has edge lengths l_1 , l_2 and l_3 and corresponding opposite joint angles θ_1 , θ_2 and θ_3 , with respective joint angle velocities $\dot{\theta}_1$, $\dot{\theta}_2$ and $\dot{\theta}_3$. The edge length l_1 is variable with velocity \dot{l}_1 .




Figure B.1: A simple triangle mechanism that has one edge of variable length (l_1), used for calculation of triangle functions $\gamma_{\bar{\Gamma}}$ and $\gamma_{\hat{\Gamma}}$.

$\gamma_{\bar{\Gamma}}$ takes the variable side length l_1 as independent position variable, and calculates the opposite angle and the two other angles, the Jacobian of those angles $\mathbf{G}_{\bar{\Gamma}}$ and the corresponding convective accelerations $\mathbf{g}_{\bar{\Gamma}}$.

$\gamma_{\hat{\Gamma}}$ takes the variable angle θ_1 as independent position variable, and calculates the corresponding opposite length and the two other angles, the Jacobian of those variables $\mathbf{G}_{\hat{\Gamma}}$ and the corresponding convective accelerations $\mathbf{g}_{\hat{\Gamma}}$.

Both functions require knowledge of the fixed lengths l_2 and l_3 , and calculation of the convective accelerations also requires an independent velocity variable. We have:

The functions that are described in this attachment have been originally derived by Dr. Featherstone.

$$\{\theta_1, \theta_2, \theta_3, \mathbf{G}_{\hat{\Gamma}}, \mathbf{g}_{\hat{\Gamma}}\} = \gamma_{\hat{\Gamma}}(l_1, l_2, l_3, \dot{l}_1), \quad (\text{B.1})$$

$$\{l_1, \theta_2, \theta_3, \mathbf{G}_{\hat{\Gamma}}, \mathbf{g}_{\hat{\Gamma}}\} = \gamma_{\hat{\Gamma}}(\theta_1, l_2, l_3, \dot{\theta}_1). \quad (\text{B.2})$$

B

 $\gamma_{\hat{\Gamma}}$

For $\gamma_{\hat{\Gamma}}$, angles θ_1, θ_2 and θ_3 opposite to l_1, l_2 and l_3 are calculated using the cosine rule:

$$\theta_1 = \arccos \frac{l_2^2 + l_3^2 - l_1^2}{2l_2l_3}$$

$$\theta_2 = \arccos \frac{l_3^2 + l_1^2 - l_2^2}{2l_3l_1}$$

$$\theta_3 = \arccos \frac{l_1^2 + l_2^2 - l_3^2}{2l_1l_2}$$

The Jacobian $\mathbf{G}_{\hat{\Gamma}}$ is then calculated as

$$\mathbf{G}_{\hat{\Gamma}} = \begin{bmatrix} G_{\hat{\Gamma},1} \\ G_{\hat{\Gamma},2} \\ G_{\hat{\Gamma},3} \end{bmatrix} = \begin{bmatrix} \partial\theta_1/\partial l_1 \\ \partial\theta_2/\partial l_1 \\ \partial\theta_3/\partial l_1 \end{bmatrix} = \begin{bmatrix} l_1/(l_2l_3 \sin\theta_1) \\ -\cos\theta_3/(c \sin\theta_1) \\ -\cos\theta_2/(c \sin\theta_1) \end{bmatrix},$$

and the convective accelerations $\mathbf{g}_{\hat{\Gamma}}$ are

$$\mathbf{g}_{\hat{\Gamma}} = \begin{bmatrix} g_{\hat{\Gamma},1} \\ g_{\hat{\Gamma},2} \\ g_{\hat{\Gamma},3} \end{bmatrix} = \dot{l}_1^2 \begin{bmatrix} \partial^2\theta_1/\partial l_1^2 \\ \partial^2\theta_2/\partial l_1^2 \\ \partial^2\theta_3/\partial l_1^2 \end{bmatrix} = \begin{bmatrix} \dot{l}_1^2(1/l_2/l_3 - \cos\theta_1 G_{\hat{\Gamma},1}^2)/\sin\theta_1 \\ \dot{l}_1^2((l_2^2 - l_3^2)/(l_3 l_1^3) - \cos\theta_2 G_{\hat{\Gamma},2}^2)/\sin\theta_2 \\ \dot{l}_1^2((l_3^2 - l_2^2)/(l_2 l_1^3) - \cos\theta_3 G_{\hat{\Gamma},3}^2)/\sin\theta_3 \end{bmatrix}.$$

 $\gamma_{\hat{\Gamma}}$

Likewise, for $\gamma_{\hat{\Gamma}}$ we find

$$l_1 = \sqrt{l_2^2 + l_3^2 - 2l_2l_3 \cos\theta_1},$$

$$\mathbf{G}_{\hat{\Gamma}} = \begin{bmatrix} G_{\hat{\Gamma},1} \\ G_{\hat{\Gamma},2} \\ G_{\hat{\Gamma},3} \end{bmatrix} = \begin{bmatrix} \partial l_1/\partial\theta_1 \\ \partial\theta_2/\partial\theta_1 \\ \partial\theta_3/\partial\theta_1 \end{bmatrix} = \begin{bmatrix} l_2l_3 \sin\theta_1/l_1 \\ (l_3^2 - l_2^2 - l_1^2)/(2l_1^2) \\ (l_2^2 - l_3^2 + l_1^2)/(2l_1^2) - 1 \end{bmatrix},$$

$$\mathbf{g}_{\hat{\Gamma}} = \begin{bmatrix} g_{\hat{\Gamma},1} \\ g_{\hat{\Gamma},2} \\ g_{\hat{\Gamma},3} \end{bmatrix} = \dot{\theta}_1^2 \begin{bmatrix} \partial^2 l_1/\partial\theta_1^2 \\ \partial^2\theta_2/\partial\theta_1^2 \\ \partial^2\theta_3/\partial\theta_1^2 \end{bmatrix} = \begin{bmatrix} \dot{\theta}_1^2(l_2l_3 \cos\theta_1 - G_{\hat{\Gamma},1}^2)/l_1 \\ \dot{\theta}_1^2 G_{\hat{\Gamma},1}(l_2^2 - l_3^2)/l_1^3 \\ -\dot{\theta}_1^2 G_{\hat{\Gamma},1}(l_2^2 - l_3^2)/l_1^3 \end{bmatrix}.$$

C

SKIPPY'S LOOP CLOSURE

This appendix describes loop closure of Skippy's hip, including both triangle kinematics of the actuated lever and that of the four-bar linkage.

PROBLEM DESCRIPTION

Figure C.1 shows a planar kinematic model of Skippy's hip joint in its zero position, consisting of 5 joints and 6 bodies, which are all numbered. Bodies and joints 0 to 5 correspond to bodies 5 to 10 in Skippy's model described in chapter 6. As such, bodies 0 to 5 are called leg, lever, torso, follower rocker and coupler respectively, and the leg is considered as the ground. Joints 1 to 4 are revolute joints (described in radians) and joint 5 is a prismatic joint on the x axis (described in meters). The five joint positions are described by joint variables and are elements of the tuple $\mathbf{q} = [q_1 \ q_2 \ q_3 \ q_4 \ q_5]^T$. Joint i connects body i to λi , where $\lambda = [0 \ 1 \ 2 \ 2 \ 3]$ describes the connectivity tree.

The loop closure constraint is satisfied when (1) the coupler at A' (at Ψ_4) connects to the lever at A , and (2) the follower at S' connects to the leg at S . In fig. C.1 this constraint is not satisfied. The mechanism has 1 DoF, for which we define one independent joint variable $y = q_5$. The objective is to find the loop closure function $\mathbf{q} = \gamma(y)$, which expresses \mathbf{q} as function of y , and to obtain the velocity mapping or Jacobian $\mathbf{G} = \mathbf{J}(\mathbf{q}, y) = \partial \mathbf{q} / \partial y$ and convective acceleration tuple $\mathbf{g} = \dot{\mathbf{G}} \dot{y} = \mathbf{J}(\mathbf{G}, y) \dot{y}^2 = \dot{y}^2 \partial \mathbf{G} / \partial y = \dot{y}^2 \partial^2 \mathbf{q} / \partial y^2$, so that the joint velocities can be calculated as $\dot{\mathbf{q}} = \mathbf{G} \dot{y}$ and joint accelerations as $\ddot{\mathbf{q}} = \mathbf{G} \ddot{y} + \mathbf{g}$. Derivations are done by making use of the predefined loop closure functions $\gamma_{\bar{\Gamma}}$ and $\gamma_{\hat{\Gamma}}$, which calculate the kinematics of a triangle mechanism that has one variable edge length. See appendix B for an explanation and derivation of both functions.

THE FIRST TRIANGLE

To define the loop closure constraint of the actuated lever (making A and A' coincide) only one triangle is required, drawn in red in fig. C.1.

We define the variable input length $l_m = q_{5,o} - q_5 = l_r \cos \phi_r - q_5$, with time derivatives $\dot{l}_m = -\dot{q}_5$ and $\ddot{l}_m = -\ddot{q}_5$. $q_{5,o}$ is an offset length that describes horizontal translation from coordinate frames Ψ_5 to Ψ_4 for $q_5 = 0$, and has been introduced so that q_5 can be set




Figure C.1: Full kinematic model of Skippy's hip joint in its zero position ($q = \mathbf{0}$). The loop closure constraints are not satisfied. Not to scale.

to a desired zero value for a specific configuration.¹ The static lengths that complete the triangle are the lever arm length l_l and the rocker joint length l_r , see fig. C.1. We then calculate the corresponding opposing angles of the triangle using eq. (B.1):

$$\{\theta_m, \theta_l, \theta_r, \mathbf{G}_1, \mathbf{g}_1\} = \gamma_{\bar{\Gamma}}(l_m, l_l, l_r, \dot{l}_m)$$

By analysing the geometry in fig. C.1 we find for q_2 and q_4

$$q_2 = -\theta_m - \phi_r + \phi_1 + \pi/2 \quad (\text{C.1})$$

$$q_4 = -\theta_l + \phi_r. \quad (\text{C.2})$$

From eqs. (C.1) and (C.2) we find $\ddot{q}_2 = -\ddot{\theta}_m$ and $\ddot{q}_4 = -\ddot{\theta}_l$, from which we can find the velocity mappings and convective accelerations:

$$\begin{aligned} \ddot{q}_2 &= G_2 \ddot{q}_5 + g_2 \\ &= -\ddot{\theta}_m = -(G_{1,1} \ddot{l}_m + g_{1,1}) \\ &= -(G_{1,1} (-\ddot{q}_5) + g_{1,1}) \\ &= \underbrace{G_{1,1}}_{G_2} \ddot{q}_5 - \underbrace{g_{1,1}}_{g_2} \end{aligned} \quad \begin{aligned} \ddot{q}_4 &= G_4 \ddot{q}_5 + g_4 \\ &= -\ddot{\theta}_l = -(G_{1,2} \ddot{l}_m + g_{1,2}) \\ &= -(G_{1,2} (-\ddot{q}_5) + g_{1,2}) \\ &= \underbrace{G_{1,2}}_{G_4} \ddot{q}_5 - \underbrace{g_{1,2}}_{g_4} \end{aligned}$$

¹In Skippy, for practical purposes, $q_{5,o}$ is defined so that $q_5 = 0$ when Skippy is fully flexed (i.e. when the torso and leg tubes are parallel to each other, like in fig. C.1).

THE SECOND AND THIRD TRIANGLE

Calculation of the remaining entries (q_1, q_3, G_1, G_3, g_1 and g_3) requires the second and third triangle, which are drawn blue and green in fig. C.1 respectively. The third (green) triangle is drawn twice, to demonstrate that two of its sides have fixed length if the loop closure constraint is satisfied (i.e. $l_f = l'_f$ and $l_v = l'_v$ are static if S coincides with S').

For the second triangle we use the input angle $\theta_{h1} = \pi/2 - q_2$, with time derivative $\dot{\theta}_{h1} = -\dot{q}_2$. With the static lengths l_k (lever centre length) and l_u (hip upper joint distance), we calculate

$$\{l_h, \theta_k, \theta_u, \mathbf{G}_{II}, \mathbf{g}_{II}\} = \gamma_{\hat{T}}(\theta_{h1}, l_k, l_u, \dot{\theta}_{h1})$$

For the third triangle we use the variable input length l_h , with time derivative $\dot{l}_h = G_{II,1}\dot{\theta}_{h1}$. The static lengths are l_v (hip support joint length) and l_f (follower link length). We calculate

$$\{\theta_{h2}, \theta_v, \theta_f, \mathbf{G}_{III}, \mathbf{g}_{III}\} = \gamma_{\hat{T}}(l_h, l_v, l_f, \dot{l}_h)$$

COMPLETE LOOP CLOSURE FORMULAE

We now have all the ingredients to complete $\mathbf{q} = \gamma(y)$, \mathbf{G} and \mathbf{g} . For $\mathbf{q} = \gamma(y)$ we find:

$$\mathbf{q} = \gamma(y) = \begin{bmatrix} \theta_f - \theta_u + \phi_v + \phi_u - \pi/2 \\ -\theta_m - \phi_r + \phi_l - \phi_u + \pi/2 \\ \theta_v - \theta_k + \pi/2 \\ -\theta_l + \phi_r \\ q_5 \end{bmatrix} \quad (\text{C.3})$$

Derivation of G_1, g_1, G_3 and g_3 is slightly less trivial. A derivation of G_3 and g_3 follows below.

$$\begin{aligned} \ddot{q}_3 &= G_3 \ddot{q}_5 + g_3 \\ &= \ddot{\theta} - \ddot{\theta}_k \\ &= \ddot{l}_h G_{III,2} + g_{III,2} - \ddot{\theta}_k \\ &= (\ddot{\theta}_{h1} G_{II,1} + g_{II,1}) G_{III,2} + g_{III,2} - (\ddot{\theta}_{h1} G_{II,2} + g_{II,2}) \\ &= (\ddot{\theta}_m G_{II,1} + g_{II,1}) G_{III,2} + g_{III,2} - (\ddot{\theta}_m G_{II,2} + g_{II,2}) \\ &= ((\ddot{l}_m G_{I,1} + g_{I,1}) G_{II,1} + g_{II,1}) G_{III,2} + g_{III,2} - ((\ddot{l}_m G_{I,1} + g_{I,1}) G_{II,2} + g_{II,2}) \\ &= ((-\ddot{q}_5 G_{I,1} + g_{I,1}) G_{II,1} + g_{II,1}) G_{III,2} + g_{III,2} - ((-\ddot{q}_5 G_{I,1} + g_{I,1}) G_{II,2} + g_{II,2}) \\ &= ((-\ddot{q}_5 G_{I,1} + g_{I,1}) G_{II,1} + g_{II,1}) G_{III,2} + g_{III,2} - ((-\ddot{q}_5 G_{I,1} + g_{I,1}) G_{II,2} + g_{II,2}) \\ &= \underbrace{G_{I,1} (G_{II,2} - G_{II,1} G_{III,2})}_{G_3} \ddot{q}_5 + \underbrace{g_{I,1} (G_{II,1} G_{III,2} - G_{II,2}) + g_{II,1} G_{III,2} - g_{II,2} + g_{III,2}}_{g_3} \end{aligned}$$

A similar derivation can be done for G_1 and g_1 . Starting from $\ddot{q}_1 = \ddot{q}_5 G_1 + g_1 = \ddot{q}_1 = \ddot{\theta}_f - \ddot{\theta}_u$, we find

$$\begin{aligned} G_1 &= G_{I,1} (G_{II,3} - G_{II,1} G_{III,3}) \\ g_1 &= g_{I,1} (G_{II,1} G_{III,3} - G_{II,3}) + g_{II,1} G_{III,3} - g_{II,3} + g_{III,3} \end{aligned}$$

REFERENCES

- Agility Robotics. (2018). *Cassie*. Retrieved 2019-01-17, from <http://www.agilityrobotics.com/robots/#cassie>
- Aguilar-Avelar, C., Rodríguez-Calderón, R., Puga-Guzmán, S., & Moreno-Valenzuela, J. (2017). Effects of nonlinear friction compensation in the inertia wheel pendulum. *Journal of Mechanical Science and Technology*, 31(9), 4425–4433. doi: 10.1007/s12206-017-0843-4
- Akano, T. T., & Fakinlede, O. A. (2013). An Effective Finite Element Method for Beam Based Compliant Mechanism. *International Journal of Engineering & Technology (IJET)*, 3(4), 407–419.
- Alexander, R. M. (1990). Three Uses for Springs in Legged Locomotion. *The International Journal of Robotics Research (IJRR)*, 9(2), 53–61. Retrieved from <http://journals.sagepub.com/doi/10.1177/027836499000900205> doi: 10.1177/027836499000900205
- Altendorfer, R., Moore, N., Komsuoglu, H., Buehler, M., Brown, H. B., McMordie, D., ... Koditschek, D. E. (2001). RHex: A biologically inspired hexapod runner. *Autonomous Robots*, 11(3), 207–213. doi: 10.1023/A:1012426720699
- Andrievsky, B. R. (2011). Global stabilization of the unstable Reaction-Wheel Pendulum. *Automation and Remote Control*, 72(9), 1981–1993. doi: 10.1134/S0005117911090189
- Armour, R., Paskins, K., Bowyer, A., Vincent, J., & Megill, W. (2007). Jumping robots: a biomimetic solution to locomotion across rough terrain. *Bioinspiration & Biomimetics*, 2(3), S65–S82. doi: 10.1088/1748-3182/2/3/S01
- Atkeson, C. G., Benzun, P. W., Banerjee, N., Berenson, D., Bove, C. P., Cui, X., ... Xinjilefu, X. (2016). *What happened at the DARPA robotics challenge finals* (Vol. 121). (submitted to DRC Finals Special Issue of the Journal of Field Robotics) doi: 10.1007/978-3-319-74666-1_17
- Azad, M. (2014). *Balancing and hopping motion control algorithms for an under-actuated robot* (Unpublished doctoral dissertation). School of Engineering, The Australian National University, Canberra, Australia.
- Azad, M., & Featherstone, R. (2013). Balancing and hopping motion of a planar hopper with one actuator. In *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 2027–2032). doi: 10.1109/ICRA.2013.6630848
- Azad, M., & Featherstone, R. (2014a). Balancing control algorithm for a 3D under-actuated robot. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 3233–3238). Retrieved from <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=6943011> doi: 10.1109/IROS.2014.6943011
- Azad, M., & Featherstone, R. (2014b, June). A New Nonlinear Model of Contact Normal Force. *IEEE Transactions on Robotics*, 30(3), 736–739. doi: 10.1109/TRO.2013

- .2293833
- Azad, M., & Featherstone, R. (2016). Angular momentum based balance controller for an under-actuated planar robot. *Autonomous Robots*, 40(1), 93–107. doi: 10.1007/s10514-015-9446-z
- Berkemeier, M. D., & Fearing, R. S. (1999). Tracking fast inverted trajectories of the underactuated Acrobot. *IEEE Transactions on Robotics and Automation*, 15(4), 740–750. doi: 10.1109/70.782028
- Bhounsule, P. A., Cortell, J., & Ruina, A. (2012). Design and control of Ranger: an energy-efficient, dynamic walking robot. In *International Conference on Climbing and Walking Robots (CLAWAR)* (pp. 441–448). doi: 10.1142/9789814415958
- Bjelonic, M., Bellicoso, C. D., Viragh, Y. D., Sako, D., Tresoldi, F. D., Jenelten, F., & Hutter, M. (2019). Keep Rollin' – Whole-Body Motion Control and Planning for Wheeled Quadrupedal Robots. *IEEE Robotics and Automation Letters*, 1, 1–8.
- Blickhan, R. (1989). The Spring-Mass Model for Running and Hopping. *Journal of Biomechanics*, 22(11/12), 1217–1227.
- Block, D. J., Astrom, K. J., & Spong, M. W. (2007). *The Reaction Wheel Pendulum*. Morgan & Claypool Publishers.
- Bogue, R. (2012). Robots for space exploration. *Industrial Robot: An International Journal*, 39(4), 323–328. Retrieved from <http://www.emeraldinsight.com/doi/10.1108/01439911211227872> doi: 10.1108/01439911211227872
- Boston Dynamics. (n.d.). *Changing your idea of what robots can do*. Retrieved 2019-06-26, from <https://www.bostondynamics.com>
- Boston Dynamics. (2012). *SandFlea – leaps small buildings in a single bound*. Retrieved 2018-12-20, from <https://www.bostondynamics.com/sandflea>
- Boston Dynamics. (2017). *Handle*. Retrieved 2019-01-18, from <https://www.bostondynamics.com/handle>
- Boston Dynamics. (2018). *Spot Mini*. Retrieved 2019-01-17, from <https://www.bostondynamics.com/spot-mini>
- Bosworth, W., Kim, S., & Hogan, N. (2015, Oct). The mit super mini cheetah: A small, low-cost quadrupedal robot for dynamic locomotion. In *Ieee international symposium on safety, security, and rescue robotics (ssrr)* (p. 1-8). doi: 10.1109/SSRR.2015.7443018
- Boyle, C., Howell, L. L., Magleby, S. P., & Evans, M. S. (2003). Dynamic modeling of compliant constant-force compression mechanisms. *Mechanism and Machine Theory*, 38(12), 1469–1487. doi: 10.1016/S0094-114X(03)00098-3
- Briod, A., Klapotcz, A., Zufferey, J. C., & Floreano, D. (2012). The AirBurr: A flying robot that can exploit collisions. In *International conference on complex medical engineering (icme)* (pp. 569–574). doi: 10.1109/ICCME.2012.6275674
- Briod, A., Kornatowski, P., Zufferey, J.-C., & Floreano, D. (2014). A Collision-resilient Flying Robot. *Journal of Field Robotics*, 31(4), 496–509. doi: 10.1002/rob
- Canaday, B., Zapolsky, S., & Drumwright, E. (2017). Interactive, iterative robot design. In *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 1188–1195). IEEE. doi: 10.1109/ICRA.2017.7989142
- Chen, D.-Z., & Tsai, L.-W. (1993). Kinematic and Dynamic Synthesis of Geared Robotic Mechanisms. *Journal of Mechanical Design*, 115(2), 241. doi: 10.1115/1.2919183

- Cherouvim, N., & Papadopoulos, E. (2005). Single actuator control analysis of a planar 3DOF hopping robot. In *Robotics: Science and Systems (RSS)* (pp. 145–152).
- Chevallereau, C., Abba, G., Aoustin, Y., Plestan, F., Wit, C. C. D., Grizzle, J. W., ... Canudas-de wit, C. (2003). RABBIT: A Testbed for Advanced Control Theory. *IEEE Control Systems Magazine*, 23(5), 57–79.
- Chiu, J., & Goswami, A. (2014). Design of a Wearable Scissored-Pair Control Moment Gyroscope (SP-CMG) For Human Balance Assist. In *ASME International Design Engineering Technical Conferences & Computers and Information in Engineering Conference* (pp. 1–10). doi: 10.1115/DETC2014-35539
- Collins, S. (2005). Efficient Bipedal Robots Based on Passive-Dynamic Walkers. *Science*, 307(5712), 1082–1085. Retrieved from <http://www.sciencemag.org/cgi/doi/10.1126/science.1107799> doi: 10.1126/science.1107799
- De Boer, T. (2012). *Foot placement in robotic bipedal locomotion* (Doctoral dissertation, Delft University of Technology). Retrieved from <https://repository.tudelft.nl/islandora/object/uuid%3A795fa8f5-84a0-4673-810c-a8265e29791c>
- Degani, A., Feng, S., Benjamin Brown, H., Lynch, K. M., Choset, H., & Mason, M. T. (2011). The ParkourBot - A dynamic BowLeg climbing robot. In *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 795–801). IEEE. doi: 10.1109/ICRA.2011.5979937
- Dholakiya, D., Bhattacharya, S., Gunalan, A., Singla, A., Bhatnagar, S., Amrutur, B., ... Kolathaya, S. (2019). *Design, Development and Experimental Realization of a Quadrupedal Research Platform: Stoch*. Retrieved from <http://arxiv.org/abs/1901.00697> (Submitted to International Conference on Control, Automation and Robotics (ICCAR))
- Digumarti, K. M., Gehring, C., Coros, S., Hwangbo, J., & Siegwart, R. (2014). Concurrent Optimization of Mechanical Design and Locomotion Control of a Legged Robot. *Mobile Service Robotics*, 315–323. Retrieved from http://www.worldscientific.com/doi/abs/10.1142/9789814623353_{_}0037 doi: 10.1142/9789814623353_{_}0037
- Driessen, J. J. M. (2015). *Machine and behaviour co-design of a powerful minimally actuated hopping robot* (MSc thesis). Delft University of Technology.
- Driessen, J. J. M. (2017, October). *Tippy – Mechanical Reference Manual* [Manual].
- Driessen, J. J. M., Featherstone, R., & Gkikakis, A. E. (2018, October). An actuator design criterion to maximize physical balance recovery. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Madrid, Spain.
- Driessen, J. J. M., Gkikakis, A. E., Featherstone, R., & Singh, B. R. P. (2019, May). Experimental demonstration of high-performance robotic balancing. In *IEEE International Conference on Robotics and Automation (ICRA)*. Montreal, Canada.
- Egan, P., Sinko, R., Leduc, P. R., & Keten, S. (2015). The role of mechanics in biological and bio-inspired systems. *Nature Communications*, 6(May), 1–12. Retrieved from <http://dx.doi.org/10.1038/ncomms8418> doi: 10.1038/ncomms8418
- Engelsberger, J., Werner, A., Ott, C., Henze, B., Roa, M. A., Garofalo, G., ... Albu-Schäffer, A. (2015). Overview of the torque-controlled humanoid robot TORO. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)* (Vol. 2015-Febru, pp.

- 916–923). doi: 10.1109/HUMANOIDS.2014.7041473
- Fanni, S., & Jweda, F. (2016). *Design of carbon fibre composite driveshaft end fittings and adhesive joint for motorsport applications* (Degree Project). KTH Royal Institute of Technology – School of Industrial Engineering and Management.
- Featherstone, R. (2008). *Rigid body dynamics algorithms*. Springer.
- Featherstone, R. (2014, September). *The Skippy Project*. Retrieved 2018-12-20, from <http://royfeatherstone.org/skippy/>
- Featherstone, R. (2015a, September). A new simple model of balancing in the plane. In *International Symposium on Robotics Research*. Sestri Levante, Italy.
- Featherstone, R. (2015b, July). Quantitative Measures of a Robot's Ability to Balance. In *Proceedings of robotics: Science and systems*. Rome, Italy. Retrieved from <http://www.roboticsproceedings.org/rss11/p26.html> doi: 10.1177/0278364916669599
- Featherstone, R. (2016a). *Linear Drive Mechanism of the Screw and Nut Type with Perfect Rolling Contact*. WIPO Patent No. WO 2015/181341 A1.
- Featherstone, R. (2016b). Quantitative measures of a robot's physical ability to balance. *The International Journal of Robotics Research*, 35(14), 1681–1696. doi: 10.1177/0278364916669599
- Featherstone, R. (2017). A simple model of balancing in the plane and a simple preview balance controller. *The International Journal of Robotics Research*, 36(13–14), 1489–1507. doi: 10.1177/0278364917691114
- Flyability. (2018). *Elios – Inspect & Explore Indoor and Confined Spaces*. Retrieved 2019-01-18, from <https://www.flyability.com/elios/>
- Fujimoto, Y. (2004). Minimum Energy Biped Running Gait and Development of Energy Regeneration Leg. In *IEEE International Workshop on Advanced Motion Control (AMC)* (pp. 415–420).
- Gajamohan, M., Merz, M., Thommen, I., & D'Andrea, R. (2012). The Cubli: A cube that can jump up and balance. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 3722–3727). doi: 10.1109/IROS.2012.6385896
- Galloway, K. C., Haynes, G. C., Ilhan, B. D., Johnson, A. M., Knopf, R., Lynch, G. A., ... Koditschek, D. E. (2010). *X-RHex: A Highly Mobile Hexapedal Robot for Sensorimotor Tasks* (Tech. Rep. No. November). doi: 10.1177/02783640122067570
- Gaponov, I., Palli, G., May, C., & Ryu, J.-H. (2016, October). *Twisted String Actuation: State of the Art, Challenges and New Applications*. (Workshop at IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS))
- Ghost Robotics. (2016, September). *Ghost Minitaur Intro*. Retrieved 2019-01-28, from https://www.youtube.com/watch?v=_YrWX9ez3jM
- Goswami, A., & Kalleem, V. (2004). Rate of change of angular momentum and balance maintenance of biped robots. In *Ieee international conference on robotics and automation (icra)* (pp. 3785–3790). doi: 10.1109/ROBOT.2004.1308858
- Grimes, J., & Hurst, J. w. (2012). The Design of ATRIAS 1.0 – a Unique Monopod, Hopping Robot. In *International conference on climbing and walking robots (clawar)* (pp. 548–554). Retrieved from [http://mime.oregonstate.edu/research/drl/publications/{_}documents/grimes{_\]2012a.pdf](http://mime.oregonstate.edu/research/drl/publications/{_}documents/grimes{_]2012a.pdf) doi: 10.1142/9789814415958_0071

- Grizzle, J. W., Hurst, J. W., Morris, B., Park, H. W., & Sreenath, K. (2009). MABEL, a new robotic bipedal walker and runner. In *American control conference (acc)* (pp. 2030–2036). doi: 10.1109/ACC.2009.5160550
- Grizzle, J. W., Moog, C. H., & Chevallereau, C. (2005). Nonlinear control of mechanical systems with an unactuated cyclic variable. *IEEE Transactions on Automatic Control*, 50(5), 559–576. doi: 10.1109/TAC.2005.847057
- Groothuis, S. S., Rusticelli, G., Zucchelli, A., Stramigioli, S., & Carloni, R. (2014). The variable stiffness actuator vsaUT-II: Mechanical design, modeling, and identification. *IEEE/ASME Transactions on Mechatronics*, 19(2), 589–597. doi: 10.1109/TMECH.2013.2251894
- Guenther, F., & Iida, F. (2017). Energy-Efficient Monopod Running with a Large Payload Based on Open-Loop Parallel Elastic Actuation. *IEEE Transactions on Robotics*, 33(1), 102–113. doi: 10.1109/TRO.2016.2623342
- Guizzo, E., & Ackerman, E. (2014, Feb). *Who Is SCHAFT, the Robot Company Bought by Google and Winner of the DRC?* Retrieved 2019-01-18, from <https://spectrum.ieee.org/automaton/robotics/humanoids/schaft-robot-company-bought-by-google-darpa-robotics-challenge-winner>
- Ha, S., Coros, S., Alspach, A., Kim, J., & Yamane, K. (2016). Task-based limb optimization for legged robots. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2016-Novem*, 2062–2068. doi: 10.1109/IROS.2016.7759324
- Haeufle, D. F. B., Taylor, M. D., Schmitt, S., & Geyer, H. (2012). A clutched parallel elastic actuator concept: Towards energy efficient powered legs in prosthetics and robotics. In *Ieee raslembis international conference on biomedical robotics and biomechatronics* (pp. 1614–1619). doi: 10.1109/BioRob.2012.6290722
- Haldane, D. W., Plecnik, M., Yim, J. K., & Fearing, R. S. (2016a). A power modulating leg mechanism for monopodal hopping. In (Vol. 2016-Novem, pp. 4757–4764). doi: 10.1109/IROS.2016.7759699
- Haldane, D. W., Plecnik, M. M., Yim, J. K., & Fearing, R. S. (2016b, dec). Robotic vertical jumping agility via Series-Elastic power modulation. *Science Robotics*, 1(1), 1–9. doi: 10.1126/scirobotics.aag2048
- Haldane, D. W., Yim, J. K., & Fearing, R. S. (2017). Repetitive extreme-acceleration (14-g) spatial jumping with Salto-1P. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 3345–3351). doi: 10.1109/IROS.2017.8206172
- Harmonic Drive AG. (2014, December). *Engineering Data CSF-Mini Gearboxes* [Datasheet].
- Hashimoto, K., Kimura, S., Sakai, N., Hamamoto, S., Koizumi, A., Sun, X., ... Takanishi, A. (2017). WAREC-1 - A four-limbed robot having high locomotion ability with versatility in locomotion styles. In *IEEE International Symposium on Safety, Security and Rescue Robotics, Conference (SSRR)* (pp. 172–178). doi: 10.1109/SSRR.2017.8088159
- Hauser, J., & Murray, R. M. (1990). Nonlinear controllers for non-integrable systems: the Acrobot example. In *American Control Conference (ACC)* (pp. 669–671).
- Hebert, P., Bajracharya, M., Ma, J., Hudson, N., Aydemir, A., Reid, J., ... Burdick, J. (2015). Mobile manipulation and mobility as manipulation - Design and algorithms of RoboSimian. *Journal of Field Robotics*, 32(2), 255–274. doi: 10.1002/rob.21566

- Heijmink, E. (2018). *The ring screw: Modelling, development and evaluation of a novel screw transmission* (MSc thesis). Delft University of Technology.
- Hill, A. V. (1950). The dimensions of animals and their muscular dynamics. *Science Progress*, 38(150), 209–230. Retrieved from <http://www.jstor.org/stable/43422833>
- Hirose, M., & Ogawa, K. (2007). Honda humanoid robots development. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 365(1850), 11–19. Retrieved from <http://rsta.royalsocietypublishing.org/cgi/doi/10.1098/rsta.2006.1917> doi: 10.1098/rsta.2006.1917
- Hobbelen, D. G. E., & Wisse, M. (2007). A disturbance rejection measure for limit cycle walkers: The gait sensitivity norm. *IEEE Transactions on Robotics*, 23(6), 1213–1224. doi: 10.1109/TRO.2007.904908
- Hodgins, J. K., & Raibert, M. H. (1990). Biped gymnastics. *The International Journal of Robotics Research*, 9(2), 115–128. Retrieved from <https://doi.org/10.1177/027836499000900209> doi: 10.1177/027836499000900209
- Hof, A. L. (2008). The 'extrapolated center of mass' concept suggests a simple control of balance in walking. *Human Movement Science*, 27(1), 112–125. doi: 10.1016/j.humov.2007.08.003
- Hubicki, C., Abate, A., Clary, P., Rezazadeh, S., Jones, M., Peekema, A., ... Hurst, J. W. (2018). Walking and Running With Passive Compliance – Lessons from Engineering a Live Demonstration of the ATRIAS Biped. *IEEE Robotics and Automation Magazine*, 25(3), 23–39. doi: 10.1109/MRA.2017.2783922
- Hubicki, C., Grimes, J., Jones, M., Renjewski, D., Spröwitz, A., Abate, A., & Hurst, J. W. (2016). ATRIAS: Design and validation of a tether-free 3D-capable spring-mass bipedal robot. *The International Journal of Robotics Research (IJRR)*, 35(12), 1497–1521. Retrieved from <http://journals.sagepub.com/doi/10.1177/0278364916648388> doi: 10.1177/0278364916648388
- Hurst, J. W., & Rizzi, A. A. (2008). Series compliance for an efficient running gait. *IEEE Robotics and Automation Magazine (RAM)*, 15(3), 42–51. doi: 10.1109/MRA.2008.927693
- Hutter, M., Gehring, C., Jud, D., Lauber, A., Bellicoso, C. D., Tsounis, V., ... Hoepflinger, M. (2016). ANYmal – A Highly Mobile and Dynamic Quadrupedal Robot. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 38–44). doi: 10.3929/ethz-a-010782581
- Hyon, S. H., Suewaka, D., Torii, Y., & Oku, N. (2017). Design and Experimental Evaluation of a Fast Torque-Controlled Hydraulic Humanoid Robot. *IEEE/ASME Transactions on Mechatronics*, 22(2), 623–634. doi: 10.1109/TMECH.2016.2628870
- iC-Haus. (2018). *iC-MU150*. Retrieved 2018-12-20, from www.ichaus.de/ic-mu150
- Jung, G. P., Casarez, C. S., Jung, S. P., Fearing, R. S., & Cho, K. J. (2016). An integrated jumping-crawling robot using height-adjustable jumping module. In (Vol. 2016-June, pp. 4680–4685). doi: 10.1109/ICRA.2016.7487668
- Kaminaga, H., Ko, T., Masumura, R., Komagata, M., Sato, S., Yorita, S., & Nakamura, Y. (2016). Mechanism and Control of Whole-Body Electro-Hydrostatic Actuator Driven Humanoid Robot Hydra. In D. Kulic, Y. Nakamura, O. Khatib, & G. Venture

- (Eds.), *International Symposium on Experimental Robotics (ISER)* (pp. 656–665). Springer. Retrieved from <http://arxiv.org/abs/1608.00627> doi: 10.1007/978-3-319-50115-4
- Kaneko, T., Ogata, K., Sakaino, S., & Tsuji, T. (2015). Impact force control based on stiffness ellipse method using biped robot equipped with biarticular muscles. *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS), 2015-Decem*, 2246–2251. doi: 10.1109/IROS.2015.7353678
- Kenneally, G. (2016). Design Principles for a Family of Direct-Drive Legged Robots. *IEEE Robotics and Automation Letters*, 1(2), 900–907.
- Kim, D., Campbell, O., Ahn, J., Sentis, L., & Paine, N. (2017). Investigations of viscoelastic liquid cooled actuators applied for dynamic motion control of legged systems. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)* (pp. 710–717). doi: 10.1109/HUMANOIDS.2017.8246950
- Klaptocz, A., Briod, A., Daler, L., Zufferey, J. C., & Floreano, D. (2013). Euler spring collision protection for flying robots. In *Ieee/rsj international conference on intelligent robots and systems (iros)* (pp. 1886–1892). doi: 10.1109/IROS.2013.6696606
- Ko, T., Kaminaga, H., & Nakamura, Y. (2018). Key design parameters of a few types of electro-hydrostatic actuators for humanoid robots. *Advanced Robotics*, 32(23), 1241–1252. Retrieved from <https://doi.org/10.1080/01691864.2018.1545604> doi: 10.1080/01691864.2018.1545604
- Kojima, K., Karasawa, T., Kozuki, T., Kuroiwa, E., Yukizaki, S., Iwaishi, S., ... Inaba, M. (2015). Development of life-sized high-power humanoid robot JAXON for real-world use. In *IEEE-RAS International Conference on Humanoid Robots (Humanoids)* (pp. 838–843). IEEE. doi: 10.1109/HUMANOIDS.2015.7363459
- Koolen, T., Bertrand, S., Thomas, G., de Boer, T., Wu, T., Smith, J., ... Pratt, J. (2016). Design of a Momentum-Based Control Framework and Application to the Humanoid Robot Atlas. *International Journal of Humanoid Robotics*, 13(1), 1650007. doi: 10.1142/S0219843616500079
- Koolen, T., de Boer, T., Rebula, J., Goswami, A., & Pratt, J. (2012). Capturability-based analysis and control of legged locomotion, Part 1: Theory and application to three simple gait models. *The International Journal of Robotics Research*, 31(9), 1094–1113. doi: 10.1177/0278364912452673
- Kovač, M., Fuchs, M., Guignard, A., Zufferey, J. C., & Floreano, D. (2008). A miniature 7g jumping robot. In *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 373–378). doi: 10.1109/ROBOT.2008.4543236
- Kovač, M., Schlegel, M., Zufferey, J. C., & Floreano, D. (2009). A miniature jumping robot with self-recovery capabilities. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 583–588). doi: 10.1109/IROS.2009.5354005
- Kussmaul, R. (2016). *Static and Buckling Analyses of Fiber Reinforced Composites in Siemens NX 10* (Tech. Rep.). Zurich: ETH.
- Lee, S. H., & Goswami, A. (2007). Reaction Mass Pendulum (RMP): An explicit model for centroidal angular momentum of humanoid robots. In *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 4667–4672). doi: 10.1109/ROBOT.2007.364198
- Matsuzaki, R., & Fujimoto, Y. (2013). Walking assist device using Control Moment Gy-

- rosopes. *IECON 2013 - 39th Annual Conference of the IEEE Industrial Electronics Society*, 6581–6586. doi: 10.1109/IECON.2013.6700220
- Maus, H.-M., Lipfert, S., Gross, M., Rummel, J., & Seyfarth, A. (2010). Upright human gait did not provide a major mechanical challenge for our ancestors. *Nature Communications*, 1(6), 1–6. doi: 10.1038/ncomms1073
- Maxon Motor. (2017). *Program 2017/18 – High Precision Drives and Systems* [Catalogue]. Retrieved from <http://www.maxonmotor.nl/maxon/view/catalog/>
- Meyer, J., Delson, N., & De Callafon, R. A. (2009). Design, modeling and stabilization of a moment exchange based inverted pendulum. In *IFAC Symposium on System Identification* (Vol. 15, pp. 462–467). IFAC. doi: 10.3182/20090706-3-FR-2004.0326
- M. Gabriel, J. (1984, 12). The effect of animal design on jumping performance. *Journal of Zoology*, 204, 533–539. doi: 10.1111/j.1469-7998.1984.tb02385.x
- Mills, N. J., & Gilchrist, A. (1991). The effectiveness of foams in bicycle and motorcycle helmets. *Accident Analysis and Prevention*, 23(2-3), 153–163. doi: 10.1016/0001-4575(91)90045-7
- Mintchev, S., de Rivaz, S., & Floreano, D. (2017). Insect-Inspired Mechanical Resilience for Multicopters. *IEEE Robotics and Automation Letters*, 2(3), 1248–1255. Retrieved from <http://ieeexplore.ieee.org/document/7833174/> doi: 10.1109/LRA.2017.2658946
- Mizuuchi, I., Nakanishi, Y., Sodeyama, Y., Namiki, Y., Nishino, T., Muramatsu, N., ... Inaba, M. (2008). An advanced musculoskeletal humanoid kojiro. In *IEEE/RAS International Conference on Humanoid Robots (Humanoids)* (pp. 294–299). doi: 10.1109/ICHR.2007.4813883
- Muehlebach, M., & D'Andrea, R. (2017). Nonlinear analysis and control of a reaction wheel-based 3D inverted pendulum. *IEEE Transactions on Control Systems Technology*, 25(1), 235–246.
- muRata. (2015). *Murata Boy: Information about Murata Boy's capabilities*. Retrieved 2018-12-20, from <https://www.murata.com/support/library/video/robotsvideo/robotvideo/murataboyvideo01>
- Nagatani, K., Kiribayashi, S., Okada, Y., Tadokoro, S., Nishimura, T., Yoshida, T., ... Hada, Y. (2011). Redesign of rescue mobile robot Quince – Toward emergency response to the nuclear accident at Fukushima Daiichi Nuclear Power Station on March 2011. In *IEEE International Symposium on Safety, Security, and Rescue Robotics (SSRR)* (pp. 13–18). IEEE. doi: 10.1109/SSRR.2011.6106794
- Negrello, F., Garabini, M., Catalano, M. G., Kryczka, P., Choi, W., Caldwell, D. G., ... Tsagarakis, N. G. (2016). WALK-MAN humanoid lower body design optimization for enhanced physical performance. In *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 1817–1824). doi: 10.1109/ICRA.2016.7487327
- Nelson, G., Saunders, A., Neville, N., Swilling, B., Bondaryk, J., Billings, D., ... Raibert, M. (2012). PETMAN: A Humanoid Robot for Testing Chemical Protective Clothing. *Journal of the Robotics Society of Japan (JRSJ)*, 30(4), 372–377. Retrieved from <http://japanlinkcenter.org/DN/JST.JSTAGE/jrsj/30.372?lang=en&from=CrossRef&type=abstract> doi: 10.7210/jrsj.30.372
- Nguyen, Q. V., & Park, H. C. (2012). Design and Demonstration of a Locust-Like Jumping Mechanism for Small-Scale Robots. *Journal of Bionic Engineering*, 9(3), 271–281.

- Retrieved from [http://dx.doi.org/10.1016/S1672-6529\(11\)60121-2](http://dx.doi.org/10.1016/S1672-6529(11)60121-2) doi: 10.1016/S1672-6529(11)60121-2
- Niiyama, R., & Kuniyoshi, Y. (2010). Design principle based on maximum output force profile for a musculoskeletal robot. *Industrial Robot: An International Journal*, 37(3), 250–255. Retrieved from <http://www.emeraldinsight.com/doi/10.1108/01439911011037640> doi: 10.1108/01439911011037640
- Niiyama, R., Nishikawa, S., & Kuniyoshi, Y. (2010). Athlete robot with applied human muscle activation patterns for bipedal running. In *IEEE/RAS International Conference on Humanoid Robots (Humanoids)* (pp. 498–503). doi: 10.1109/ICHR.2010.5686316
- Ogura, Y., Aikawa, H., Shimomura, K., Kondo, H., Morishima, A., Lim, H. O., & Takanishi, A. (2006). Development of a New Humanoid Robot WABIAN-2. In *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 76–81). doi: 10.1109/ROBOT.2006.1641164
- Orsolino, R., Focchi, M., Mastalli, C., Dai, H., Caldwell, D. G., & Semini, C. (2018). Application of Wrench based Feasibility Analysis to the Online Trajectory Optimization of Legged Robots. *IEEE Robotics and Automation Letters*, 3(4), 3363–3370. Retrieved from <http://arxiv.org/abs/1712.06833> doi: 10.1109/LRA.2018.2836441
- Ott, C., Baumgärtner, C., Mayr, J., Fuchs, M., Burger, R., Lee, D., ... Hirzinger, G. (2010, dec). Development of a biped robot with torque controlled joints. *IEEE-RAS International Conference on Humanoid Robots (Humanoids)*, 167–173. doi: 10.1109/ICHR.2010.5686340
- Paine, N., & Sentis, L. (2012). A new prismatic series elastic actuator with compact size and high performance. In *IEEE International Conference on Robotics and Biomimetics (ROBIO)* (pp. 1759–1766). doi: 10.1109/ROBIO.2012.6491222
- Paine, N., & Sentis, L. (2015). Design and Comparative Analysis of a Retrofitted Liquid Cooling System for High-Power Actuators. *Actuators*, 4(3), 182–202. Retrieved from <http://www.mdpi.com/2076-0825/4/3/182/> doi: 10.3390/act4030182
- Papadopoulos, E., & Cherouvim, N. (2004). On increasing energy autonomy for a one-legged hopping robot. In *IEEE International Conference on Robotics and Automation (ICRA)* (Vol. 5, pp. 4645–4650). doi: 10.1109/ROBOT.2004.1302450
- Park, H. W., Park, S., & Kim, S. (2015). Variable-speed quadrupedal bounding using impulse planning: Untethered high-speed 3D Running of MIT Cheetah 2. In *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 5163–5170). IEEE. doi: 10.1109/ICRA.2015.7139918
- Park, H.-W., Wensing, P., & Kim, S. (2015). Online Planning for Autonomous Running Jumps Over Obstacles in High-Speed Quadrupeds. In *Robotics: Science and systems conference (rss)*. doi: 10.15607/rss.2015.xi.047
- Pasch, K. A., & Seering, W. P. (1984). On the Drive Systems for High- Performance Machines. *ASME Journal of Mechanisms, Transmissions, and Automation in Design*, 106(83), 102–108. doi: 10.1115/1.3258545
- Patel, N. M., Kang, B.-S., Renaud, J. E., & Tovar, A. (2009). Crashworthiness Design Using Topology Optimization. *Journal of Mechanical Design*, 131(6), 061013. Retrieved from <http://mechanicaldesign.asmedigitalcollection>

- .asme.org/article.aspx?articleid=1450060 doi:10.1115/1.3116256
- Pedersen, C. B. W. (2003). Topology optimization for crashworthiness of frame structures. *International Journal of Crashworthiness*, 8(1), 29–39. Retrieved from <http://link.springer.com/article/10.1533/ijcr.2003.0218{%}5Cnpapers2://publication/uuid/CADC1A89-CA22-4794-AC97-06F4259D3F5B> doi: 10.1533/ijcr.2004.0288
- Pfeifer, R., Lungarella, M., & Iida, F. (2007). Self-Organization, Embodiment, and Biologically Inspired Robotics. *Science*, 318(November), 1088–1093.
- Plecnik, M. M., Haldane, D. W., Yim, J. K., & Fearing, R. S. (2016). Design Exploration and Kinematic Tuning of a Power Modulating Jumping Monopod. *Journal of Mechanisms and Robotics*, 9(1), 011009. doi: 10.1115/1.4035117
- Pololu. (2016). *Pololu G2 High-Power Motor Driver 24v21*. Retrieved 2018-12-21, from <https://www.pololu.com/product/2995>
- Poulakakis, I., & Grizzle, J. W. (2007). Monopodal running control: SLIP embedding and virtual constraint controllers. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 323–330). doi: 10.1109/IROS.2007.4399559
- Poulakakis, I., & Grizzle, J. W. (2009). Modeling and control of the monopodal robot Thumper. In *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 3327–3334). Retrieved from <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5152708> doi: 10.1109/ROBOT.2009.5152708
- Prabhu, S. G. R., Seals, R. C., Kyberd, P. J., & Wetherall, J. C. (2018, December). A survey on evolutionary-aided design in robotics. *Robotica*, 36(12), 1804–1821.
- Pratt, G. A., Williamson, M. M., Dillworth, P., Pratt, J., & Wright, A. (1995). Stiffness isn't everything. *International Symposium on Experimental Robotics (ISER)*, 253–262. Retrieved from <http://www.springerlink.com/index/10.1007/BFb0035216> doi: 10.1007/BFb0035216
- Pratt, J., & Krupp, B. (2008). Design of a bipedal walking robot. In *Proceedings of SPIE* (p. 69621F). Retrieved from <http://proceedings.spiedigitallibrary.org/proceeding.aspx?doi=10.1117/12.777973> doi: 10.1117/12.777973
- Pratt, J. E., Chew, C.-M., Torres, A., Dilworth, P., & Pratt, G. A. (2001). Virtual Model Control: An Intuitive Approach for Bipedal Locomotion. *The International Journal of Robotics Research (IJRR)*, 20(2), 129–143. Retrieved from <http://journals.sagepub.com/doi/10.1177/02783640122067309> doi: 10.1177/02783640122067309
- Project Centauro. (2014). *Centauro*. Retrieved 2019-01-17, from <https://www.centauro-project.eu/>
- Radford, N. A., Strawser, P., Hambuchen, K., Mehling, J. S., Verdeyen, W. K., Donnan, S. A., ... Akinyode, J. (2015). Valkyrie: NASA's First Bipedal Humanoid Robot. *Journal of Field Robotics*, 32(3), 397–419. doi: 10.1002/rob
- Raibert, M. H. (1986). *Legged robots that balance*. MIT Press, Cambridge, Massachusetts.
- Robinson, D. W., Pratt, J. E., Paluska, D. J., & Pratt, G. A. (1999). Series elastic actuator development for a biomimetic walking robot. In *IEEE/ASME International Conference on Advanced Intelligent Mechatronics (AIM)* (pp. 561–568). Retrieved from <http://ieeexplore.ieee.org/document/803231/> doi: 10.1109/

- AIM.1999.803231
- Schuitema, E. (2012). *Reinforcement learning on autonomous humanoid robots* (Doctoral dissertation, Delft University of Technology). Retrieved from <https://repository.tudelft.nl/islandora/object/uuid%3A986ea1c5-9e30-4aac-ab66-4f3b6b6ca002>
- Schwind, W. J., & Koditschek, D. E. (1997). Characterization of monopod equilibrium gaits. In *IEEE International Conference on Robotics and Automation (ICRA)* (Vol. 3, pp. 1986–1992). Retrieved from <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=619161> doi: 10.1109/ROBOT.1997.619161
- Segway. (2019). *Segway: Personal Transportation That Simply Moves You*. Retrieved 2019-04-11, from <http://www.segway.com>
- Semini, C. (2010). *HyQ – Design and Development of a Hydraulically Actuated Quadruped Robot* (Doctoral dissertation, University of Genova). Retrieved from http://www.semmini.ch/downloads/Semini2010PhD_HyQ_Robot.pdf
- Semini, C., Barasuol, V., Goldsmith, J., Frigerio, M., Focchi, M., Gao, Y., & Caldwell, D. G. (2017). Design of the Hydraulically Actuated, Torque-Controlled Quadruped Robot HyQ2Max. *IEEE/ASME Transactions on Mechatronics*, 22(2), 635–646. doi: 10.1109/TMECH.2016.2616284
- Seok, S., Wang, A., Yee, M., Chuah, M., Hyun, D. J., Lee, J., ... Kim, S. (2015). Design principles for energy-efficient legged locomotion and implementation on the MIT cheetah robot. *IEEE/ASME Transactions on Mechatronics*, 20(3), 1117–1129.
- Seyfarth, A., Geyer, H., Blickhan, R., Lipfert, S., Rummel, J., Minekawa, Y., & Iida, F. (2006). Running and walking with compliant legs. In M. Diehl & K. Mombaur (Eds.), *Fast motions in biomechanics and robotics* (pp. 383–402). Heidelberg: Springer.
- Seyfarth, A., Geyer, H., Günther, M., & Blickhan, R. (2002). A movement criterion for running. *Journal of Biomechanics*, 35(5), 649–655. doi: 10.1016/S0021-9290(01)00245-7
- Shao, X., Huang, Q., Wang, Z., Cai, Q., & Wang, W. (2014). Motion planning and compliant control for a quadruped robot on complicated terrains. In *IEEE International Conference on Mechatronics and Automation (ICMA)* (pp. 1587–1594). IEEE. doi: 10.1109/ICMA.2014.6885937
- Spielberg, A., Araki, B., Sung, C., Tedrake, R., & Rus, D. (2017). Functional Co-Optimization of Articulated Robots. In *IEEE International Conference on Robotics and Automation (ICRA)* (pp. 5035–5042). IEEE. doi: 10.1109/ICRA.2017.7989587
- Spong, M. W. (1995). The Swing Up Control Problem For The Acrobot. *IEEE Control Systems Magazine*, 15(1), 49–55. doi: 10.1109/37.341864
- Sreenath, K., Park, H., & Grizzle, J. (2013). Embedding Active Force Control within the Compliant Hybrid Zero Dynamics to Achieve Stable, Fast Running on MABEL. *The International Journal of Robotics Research (IJRR)*, 32(3), 324–345.
- Steinmeyer. (2015). *Rolled Ball Screws - The Alternative in Linear Technology* [Catalog].
- Stelzer, M., & von Stryk, O. (2005). Efficient Forward Dynamics Simulation and Optimization of Locomotion: From Legged Robots to Biomechanical Systems. In *Symposium on Adaptive Motion of Animals and Machines (AMAM)*.
- Toniatti, G., Schiavi, R., & Bicchi, A. (2005). Design and control of a variable stiffness actuator for safe and fast physical human/robot interaction. In *IEEE International*

- Conference on Robotics and Automation (ICRA)* (Vol. 2005, pp. 526–531). doi: 10.1109/ROBOT.2005.1570172
- Treloar, L. R. G. (1973). The elasticity and related properties of rubbers. *Reports on Progress in Physics*, 36, 755-826.
- Tuerk, D.-A., Zueger, A., Klahn, C., & Meboldt, M. (2015). Combining additive manufacturing with CFRP composites: Design potentials. In *International conference on engineering design (iced)* (Vol. 4, pp. 279–290). Retrieved from <https://www.engineeringvillage.com/share/document.url?mid=cpx{ }M7fa5c8511566b182b55M7c6210178163171{&}database=cpx>
- Unitree. (2017). *Laikago*. Retrieved 2019-01-18, from <http://www.unitree.cc/>
- Vallery, H., Bögel, A., O'Brien, C., & Riener, R. (2012, November). Cooperative Control Design for Robot-Assisted Balance During Gait. *Automatisierungstechnik*, 60(11), 715–720. doi: 10.1524/auto.2012.1041
- Vanderborght, B., Albu-Schaeffer, A., Bicchi, A., Burdet, E., Caldwell, D. G., Carloni, R., ... Wolf, S. (2013). Variable impedance actuators: A review. *Robotics and Autonomous Systems*, 61(12), 1601–1614. Retrieved from <http://dx.doi.org/10.1016/j.robot.2013.06.009> doi: 10.1016/j.robot.2013.06.009
- Vanderborght, B., Van Ham, R., Lefeber, D., Sugar, T. G., & Hollander, K. W. (2009). Comparison of Mechanical Design and Energy Consumption of Adaptable, Passive-compliant Actuators. *The International Journal of Robotics Research (IJRR)*, 28(1), 90–103. Retrieved from <http://ijr.sagepub.com/cgi/doi/10.1177/0278364908095333> doi: 10.1177/0278364908095333
- Vectornav. (2009). *VN-100 IMU/AHRS*. Retrieved 2019-01-09, from <https://www.vectornav.com/products/vn-100>
- Venkiteswaran, V. K. (2017). *Development of a Design Framework for Compliant Mechanisms using Pseudo-Rigid-Body Models* (Doctoral dissertation, The Ohio State University). Retrieved from <https://u.osu.edu/kalpathyvenkiteswaran.1/files/2017/03/Venkat-K-V-dissertation-161q49z.pdf>
- Viana, J. C. (2006). Polymeric materials for impact and energy dissipation. *Plastics, Rubber and Composites*, 35(6-7), 260–267. Retrieved from <http://www.tandfonline.com/doi/full/10.1179/174328906X146522> doi: 10.1179/174328906X146522
- Wang, W., & Yang, Y. (2011). Turning maneuvers and mediolateral reaction forces in a quadruped robot. In *IEEE International Conference on Robotics and Biomimetics (ROBIO)* (pp. 515–520). doi: 10.1109/ROBIO.2011.6181338
- Wensing, P. M., Wang, A., Seok, S., Otten, D., Lang, J., & Kim, S. (2017). Proprioceptive actuator design in the MIT Cheetah: Impact mitigation and high-bandwidth physical interaction for dynamic legged robots. *IEEE Transactions on Robotics*, 33(3), 509–522. doi: 10.1109/TRO.2016.2640183
- Wikipedia. (2019). *Super Ball*. Retrieved 2019-03-08, from https://en.wikipedia.org/wiki/Super_Ball
- Woodward, M. A., & Sitti, M. (2014). MultiMo-Bat: a biologically inspired integrated jumping-gliding robot. *The International Journal of Robotics Research*, 33(12), 1511-1529.
- Yim, J. K., & Fearing, R. S. (2018). Precision Jumping Limits from Flight-phase Control in

- Salto-1P. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (pp. 2229–2236). Madrid, Spain.
- Yu, X., & Iida, F. (2013). Minimalistic models of an energy efficient vertical hopping robot. *IEEE International Conference on Robotics and Automation (ICRA)*, 7–12. doi: 10.1109/ICRA.2013.6630549
- Yu, Y.-Q., Li, Q., & Xu, Q.-P. (2018). Pseudo-rigid-body dynamic modeling and analysis of compliant mechanisms. *Proceedings of the Institution of Mechanical Engineers, Part C: Journal of Mechanical Engineering Science*, 232(9), 1665–1678. Retrieved from <http://journals.sagepub.com/doi/10.1177/0954406217707547> doi: 10.1177/0954406217707547
- Zaitsev, V., Gvirsman, O., Ben Hanan, A., U. and Weiss, Ayali, A., & Kosa, G. (2015). Locust-inspired miniature jumping robot. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*.
- Zeglin, G., & Brown, H. B. (2002). First Hops of the 3D Bow Leg Hopper. In *International Conference on Climbing and Walking Robots (CLAWAR)* (pp. 357–364).
- Zollner Elektronik AG. (2013). *World's Largest Walking Robot Dragon*. Retrieved 2019-01-25, from <https://www.youtube.com/watch?v=qmbW5gvAX4U>

CURRICULUM VITÆ

Josephus J. M. DRIESSEN

12-07-1991 Born in Delft, The Netherlands.

EDUCATION

- 2009–2012 Undergraduate in Mechanical Engineering
Delft University of Technology
- 2012–2015 Master in Mechanical Engineering – Bio-Mechanical Design
Machine and behaviour co-design of a powerful minimally actuated hopping robot (MSc thesis, 2015).
Delft University of Technology

STUDENT SUPERVISION

- Winter 2017 Elco Heijmink
The Ring Screw: Modelling, Development and Evaluation of a Novel Screw Transmission (MSc thesis, 2018).
Delft University of Technology

DISSEMINATION

- *Driessen, J. J. M., Featherstone, R., & Gkikakis, A. E.* (2018, October). An actuator design criterion to maximize physical balance recovery. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Madrid, Spain.
- *Driessen, J. J. M., Gkikakis, A. E., Featherstone, R., & Singh, B. R. P.* (2019, May). Experimental demonstration of high-performance robotic balancing. In *IEEE International Conference on Robotics and Automation (ICRA)*. Montreal, Canada.