

UNIVERSITÉ DE LIÈGE
Faculté des Sciences Appliquées



Rhythmic Movements Control:
Parallels between Human
Behavior and Robotics

Thèse de doctorat
présentée par **Renaud RONSSE**
en vue de l'obtention du grade
de Docteur en Sciences Appliquées

Jury: Prof. P. Lefèvre (Université catholique de Louvain), co-promoter
Prof. J. Piater (Université de Liège)
Prof. R. Sepulchre (Université de Liège), promoter
Prof. M. Spong (University of Illinois at Urbana-Champaign)
Prof. D. Sternad (Pennsylvania State University)
Prof. P. Wolper (Université de Liège), president

Acknowledgments

During the last five years, I had the chance to share my time between two exciting labs: the department of Electrical Engineering and Computer Science at the Université de Liège, and the Centre for Systems Engineering and Applied Mechanics at the Université catholique de Louvain. First of all, I would like to thank all my colleagues from these labs for the working spirit and the conviviality they contribute to develop in those environments.

My most hearty acknowledgments go to my advisors, Rodolphe Sepulchre and Philippe Lefèvre, for their enthusiasm and availability in excellent scientific supervision, and their concern in strengthening the social networkings in their respective teams. Not only they have been deeply involved into all the contributions of the present thesis, but also they have learned me my job (at least, they tried...).

I am sincerely indebted to the other members of the Jury: Profs. J. Piater and P. Wolper from Liège, and Profs. D. Sternad and M. Spong who were proposed to participate as foreign members. I am honored they accepted to critically examine my work.

I am also thankful to the former and present members the “systems and modeling” research unit at Liège and the “Eyelab” group at Louvain-la-Neuve. More particularly, I would like to thank Manuel Gérard and Fabien Defays, who were deeply involved in the development of the Wiper robot; and Jean-Jacques Orban de Xivry and Olivier White for having demonstrated that professional relationship may become true friendship. My considerations go also to Tim L. and Phil K. for their fascinating charisma.

I sincerely thank the technical staff who was involved in the development of Wiper: Alain Stuber, Pascal Harmeling, Joseph Simon, Robert and Georges Gérard, and mainly Michel Dewan, whose knowhow saved the project more than once.

I would like to express my gratitude to a pool of “old-time” close friends, some of them being fresh or future PhD’s. I enjoyed to share with them our professional experiences. In particular, I acknowledge Caroline Raick and Jérôme Walmag for the unreasonably long lunch breaks we spent to rebuild the world.

Last but not least, I deeply thank my family for their support during those years. I thank my parents, who offered me to complete my studies. More than anyone else, I thank my wife Catherine for supporting me with bravery and kindness; and our son Nicolas, who has no idea of what a PhD thesis could be, but who learns me daily a much more important job. . .

Contents

Acknowledgments	ii
1 General Introduction	1
1.1 Context of the present work	1
1.2 Thesis statement	4
1.3 Major contributions of the thesis	6
1.4 Scope of the thesis	7
2 Rhythmic Movements and Juggling	9
2.1 Rhythmic movements	9
2.1.1 Rhythmic movements and the Central Pattern Generator paradigm	9
2.1.2 Rhythmic arm movements are not discrete	10
2.1.3 The coordination of rhythmic movements	12
2.1.4 Mathematical models of (coordinated) rhythmic movements .	14
2.1.5 Conclusion	16
2.2 Juggling at the crossroad	16
2.2.1 Juggling is a representative rhythmic movement	16
2.2.2 A dynamical systems perspective in the cascade juggling . . .	18
2.2.3 Juggling robots	20
2.2.4 From juggling to locomotion	22
2.2.5 Conclusion	24
2.3 Concluding remark	24
3 The Bouncing Ball	25
3.1 Introduction	25
3.2 Open-loop dynamics of a ball bouncing on a sinusoidally vibrating racket	26
3.2.1 Bouncing ball model	26
3.2.2 Sinusoidal actuation	27
3.3 Robust feedback control of bouncing robots	30
3.4 Human control of the bouncing ball	32
3.5 Conclusion	35

4	Robust Closed-Loop Control of the Bouncing Ball	37
4.1	Introduction	37
4.2	Feedback Control of the Bouncing Ball	38
4.2.1	The sinusoidal actuation vs. the mirror law	38
4.2.2	Stabilization with sensorless sinusoidal actuation	39
4.2.3	State feedback control	40
4.2.4	Deadbeat convergence	41
4.2.5	Output feedback deadbeat control	42
4.2.6	A blind mirror law based on output feedback	44
4.3	Robustness to Model Uncertainties	44
4.3.1	Uncertainty of the impact model	44
4.3.2	Robustness requires negative acceleration	45
4.3.3	Integral feedback control	47
4.4	Conclusion	49
	Appendix 4.A Linearized Equations of the 1D Noisy Bouncing Ball with Piecewise Quadratic Control	49
5	Experimental Setup: the Wiper Robot	51
5.1	Introduction	51
5.2	Wiper: an experimental setup amenable to mathematical modeling	51
5.3	Conclusion	55
6	Sensorless Stabilization of 2D Patterns of the Wiper Robot	57
6.1	Introduction	57
6.2	A bounce juggler model	58
6.3	Energy balance of the steady-state solutions	61
6.3.1	Period-one orbit	61
6.3.2	Period-two orbit	63
6.4	Periodic orbits of the sinusoidally actuated wedge-billiard	68
6.4.1	The square wedge-billiard	68
6.4.2	The general vibrating wedge-billiard	70
6.5	Stabilization of periodic orbits by rotational actuation of the edges	72
6.5.1	Stability of period-one orbits	73
6.5.2	Stability of period-two orbits in the square wedge	76
6.5.3	Stability of period-two orbits in the actuated non-square wedge	76
6.6	Experimental results	77
6.6.1	Stabilization of period-one orbits	77
6.6.2	Stabilization of period-two orbits	78
6.7	Conclusion	79
	Appendix 6.A Iteration map of the rotationally actuated wedge-billiard	80
	Appendix 6.B Jacobian matrices	81

7	Rhythmic Feedback Control of Wiper in Blindness	83
7.1	Introduction	83
7.2	Feedback Control of the Simplified Wedge-Billiard	84
7.3	Experimental Validation with the Wiper Robot	85
7.3.1	Experimental setup	85
7.3.2	From the bouncing ball model to the wiper model	85
7.3.3	Results	87
7.4	Conclusion	89
8	Human Control of Bimanual Impact Juggling: Trade-off between Performance and Robustness	91
8.1	Introduction	91
8.2	Methods	92
8.2.1	Task and experimental setup description	92
8.2.2	Data analysis	94
8.2.3	Statistics	95
8.3	Results	95
8.3.1	Tempo dictated the angular position at impact	95
8.3.2	Intermittent or rhythmic behavior	96
8.3.3	Performance vs. robustness of the sensorimotor loop	98
8.3.4	Performance vs. robustness in movement amplitude	101
8.4	Discussion	104
8.5	Conclusion	106
9	General Discussion and Perspectives	107
9.1	Impact juggling control	108
9.2	Neuroscience and robotics	110
9.3	Towards a model for the control of rhythmic movements	111
9.4	Conclusion of the thesis	112
9.5	Future perspectives	113
	Appendices	115
A	Relevant Insights in Control Theory - a Benchmark Example	117
A.1	A benchmark example	117
A.2	Feedback and feedforward control	119
A.3	Robustness vs. performance	121
B	Major Technical Characteristics of the Experimental Setup	125
B.1	Technical details about the robotics configuration of Wiper	125
B.1.1	Hardware description	125
B.1.2	Programs architecture	129
B.2	Technical details about the “human” configuration of Wiper	130

C	Computation of Gaze Under Unrestrained Head Movements	133
C.1	Introduction	133
C.2	Materials and Methods	135
C.2.1	Geometrical developments	135
C.2.2	Calibration method	140
C.2.3	Experimental validation	142
C.3	Results	144
C.3.1	Calibration	144
C.3.2	Validation tasks	147
C.4	Discussion	149
Appendix C.A	Eye-in-head orientation	152
Appendix C.B	Calibration of the points T	154
	Bibliography	157

Chapter 1

General Introduction

It would be possible to describe everything scientifically, but it would make no sense; it would be without meaning, as if you described a Beethoven symphony as a variation of wave pressure.

Albert Einstein

1.1 Context of the present work

Today's science is interdisciplinary. For example, mathematicians and economists are used to collaborate together to elaborate complex micro- or macro-economical models of the society, and to potentially anticipate the market evolution in a particular context under particular events. But the most "heterogeneous" disciplines are certainly the life sciences. People working on biological systems include physicians (with any kind of specialties), physicists, biologists, veterinarians, chemists, pharmacists, mathematicians, kinesiologists, psychologists or even engineers! They are used to collaborate — and to share their respective vocabulary — in order to disentangle the complexity of living systems, from the chemical reactions involved at the smallest molecular level, to the social and cognitive mechanisms governing the largest ecological populations. Last but not least, they have also to understand how these vastly different space- and time-scales are related to each other.

The engineering contribution to life sciences is basically twofold. First, the engineers' knowhow is mandatory in the development of dedicated technologies: bio-compatible sensors, medical imaging techniques, prosthesis design, etc... This encompasses both the "hardware" development, and the "software" management, for example by elaborating dedicated signal processing algorithms. Second, the engineering viewpoint is also emerging at the level of data interpretation. Indeed,

engineers are used to *model* the systems they deal with, and to study their behavior through the mathematical properties of those models. This is potentially relevant in living organisms, since a tremendous number of individual “agents” interact with each other to produce the global picture. In neuroscience, the use of such mathematical tools for modeling and analysis purposes refers to the discipline of *computational neuroscience*. One of the most celebrated example of computational neuroscience model has been proposed by two English physiologists, Hodgkin and Huxley (1952), to describe the ionic currents through the squid giant axon that are responsible for the propagation of action potentials through the axon membrane. Hodgkin and Huxley received the Nobel price in Physiology or Medicine in 1963 “for their discoveries concerning the ionic mechanisms involved in excitation and inhibition in the peripheral and central portions of the nerve cell membrane” (www.nobelprize.org).

In the particular field of motor control, many computational aspects are also emerging (Jordan and Wolpert, 1999; Wolpert and Ghahramani, 2000; Scott and Norman, 2003). The discipline of *computational motor control* has adapted system-theoretic concepts and related engineering computational tools to the control of movements, both on the basis of actual data sets, and under biologically plausible architectures. Some examples are (see Jordan and Wolpert, 1999, for more details):

Motor planning, which refers to the elaboration of the effector trajectory, and the related muscular command. The coordination between several joints and several limbs is programmed through motor planning, as a consequence of the redundancy in the motor system (Bernstein, 1967).

Optimal control, which refers to computational techniques used to discriminate an “optimal” trajectory in motor planning, as the one that minimizes a given cost function. Particular cost functions penalize for example the non-smoothness (see the pioneering work of Flash and Hogan, 1985, on the minimum-jerk control), the energy expenditure, the movement duration, etc... Harris and Wolpert (1998) used the theory of optimal control to propose the minimum-variance principle of motor planning for both eye and arm movements. They suggested that both eye and arm movements planning is computed to minimize the biological noise, proportional to the input command amplitude and to the duration of the movement.

Internal models: the internal model principle postulates the existence of internal dynamical models of the body and/or the environment dynamics in the brain (see e.g. Miall et al., 1993; Wolpert and Miall, 1996; Wolpert et al., 1998; Wolpert and Kawato, 1998; Kawato, 1999; Haruno et al., 2001; Mehta and Schaal, 2002). Internal models may be either *forward* models — to achieve motor or sensory prediction —, or *inverse* models — to compute the neural command related to a desired behavior.

State estimation, which describes the computational techniques aiming at retrieving the system state on the basis of the measured sensory inflows.

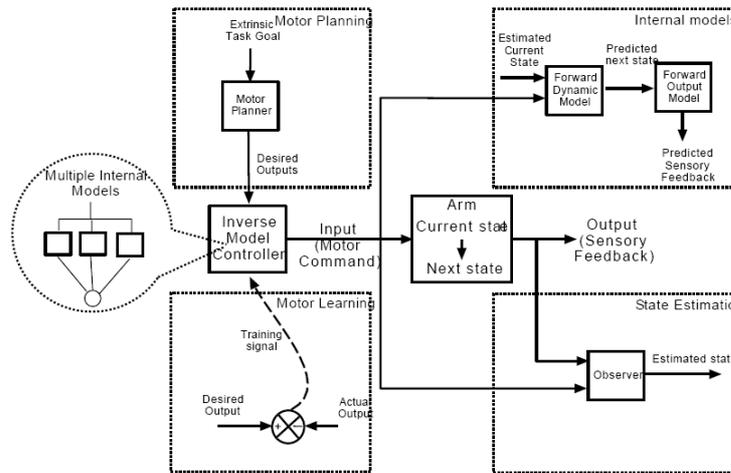


Figure 1.1: The motor system is shown schematically along with related themes of computational motor control. The motor system (center) has inputs — the motor commands — which cause it to change its states and produce an output — the sensory feedback. For clarity not all lines are shown. Reprinted from Jordan and Wolpert (1999).

Motor learning, encompassing the techniques for learning the controlled dynamics. Several *machine learning* algorithms have been explored as potentially relevant for explaining the learning effects in biological data sets.

Modularity: since the behavior is rapidly adapting in changing environments, one may suppose that several internal models of the environment are stored in parallel. These models are recruited depending on the context, under a modular architecture (see the MOSAIC model, in Wolpert and Kawato, 1998; Jordan and Wolpert, 1999; Wolpert and Ghahramani, 2000; Haruno et al., 2001).

These computational concepts, and a global picture of their relationships, are summarized in Fig. 1.1.

Computational questions in human motor control often parallel questions in robotics. Indeed, Schaal and Schweighofer (2005) pointed several fields of convergence between recent research directions in computational motor control, and well-established theories in robotics and artificial intelligence: motor control with internal models and in the presence of noise, motor learning, coordinate transformation, movement planning with motor primitives and probabilistic inference in sensorimotor control (the “Bayesian” brain). Consequently, the global picture we presented as diagram of the motor control architecture (Fig. 1.1) can also be interpreted as a sketch of the control architecture of skilled robots. The highest level box is a trajectory planner, which has to program the desired movement. This

movement is executed and closed-loop controlled, by a controller which can include internal models of the task to increase the bandwidth. Finally, the sensory feedback is used to estimate the system state and to close the loop both with the trajectory planner and the controller (Schaal and Schweighofer, 2005). Usually, these three “black-boxes” (planner, controller, estimator) are designed separately, since it is basically assumed that their bandwidths differ by several orders of magnitude.

Due to their similar global architecture, analysis and design investigations cross-fertilize between robotics and human motor control. One can say that this interaction is bidirectional in the following sense: first, the complex control strategies used in the human brain provide a source of inspiration in robotics designs. Secondly, the computational and system-theoretic models in robotics provide useful insights into the interpretation of the high-dimensional behavioral and neurophysiological data sets. In particular, the computational and mathematical tools available for design purpose are relevant to understand the human control strategies, and to better understand why and how a particular movement trajectory has been adopted in a particular context.

1.2 Thesis statement

In this thesis, we address both the design of robotics control and the analysis of human behavior in the particular context of rhythmic movements. Moreover, the designed control laws have been implemented on a robot, which executed the *same* task as the human subjects with the *same* experimental setup. The analysis of both data sets (robot and human) led to fruitful comparisons.

Two central system-theoretic concepts, ubiquitous in control design (see e.g. Franklin et al., 2005; Astrom and Murray, 2005), are considered throughout the thesis: the balance between feedback and feedforward and the trade-off between performance and robustness. Feedback and feedforward are indeed complementary:

- Feedforward control is cheap since it relaxes the need of sensor design, and is potentially of large bandwidth since it does not have to cope with error propagation and/or delays. The major drawback of feedforward is that both stability and robustness depend on the open-loop properties. A classical feedforward control scheme is shown in Fig. 1.2(a).
- Feedback control is robust, since it exploits the measured system state to adapt the control. Basically, feedback can then be used to achieve a design whose both performance and robustness are close to desired levels. The major drawback of feedback is that it rests on the sensors accuracy, and is sensitive to the delays inherent in sensory processing. A classical feedback loop is shown in Fig. 1.2(b).

To exploit their respective advantages, the ideal option is to combine those two actions, as pictured in Fig. 1.2(c). The subsequent question is thus: what is the

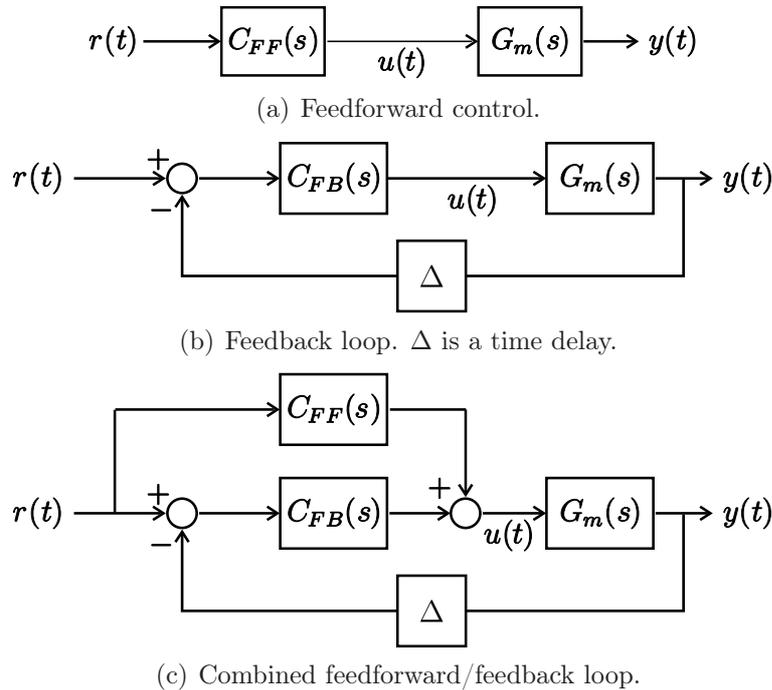


Figure 1.2: Feedforward, feedback and combined loops. C_{FF} and C_{FB} are the feedforward and feedback controllers, while G_m is the controlled open-loop system.

minimum quantity of feedback which is necessary to maintain the system robustness, while relaxing as much as possible the need of sensor design and sensed signals processing? It is of central interest to investigate this question in the context of rhythmic movements, since the rhythmicity may significantly simplify the task execution. Indeed, the control target, the actuation profile and (potentially) the sensory feedback are repeated throughout the cycles, and the mismatch between one cycle and the steady-state could consequently be rapidly identified.

The second trade-off concerns the robustness and the performance of any closed-loop system. These two criteria are always traded, such that they cannot be arbitrarily improved simultaneously. The robustness refers to the system ability to maintain its stability despite changing or noisy environments, while the performance quantifies how well it performs: how small is the static error? ; how fast is the reactivity? ; how damped is the overshoot? ; how large is the bandwidth? ; how cheap is the energy expenditure? ; how bounded is the variability? ; etc... Once again, the thesis investigates this compromise in the framework of rhythmic movements. Note that an appendix has been added to this thesis (Appendix A) to illustrate the compromise between feedback and feedforward, and between performance and robustness in a benchmark example of linear time-invariant system (a DC electrical motor).

Another major contribution of the thesis was to build an original experimental setup, both for the validation of the designed robotics control laws and for the

acquisition of human data. This setup is based on a simplified juggling paradigm, and is extensively described in Chapter 5.

1.3 Major contributions of the thesis

A central postulate of the present manuscript is that the planning of rhythmic movements is different from the planning of individual discrete movements in motor control, and should consequently be designed as such in robotics. As a central contribution, we show the advantage of rhythmic movements to increase the robustness of the system and to consequently relax the need for sensor design and sensory processing. Moreover, the dynamical systems which are considered in the present manuscript belong to the particular class of *hybrid* systems, resulting from the combination of continuous and discrete dynamics. Hybrid systems are currently a very active area of research in the control community, since the control problems are considerably more difficult than purely continuous or purely discrete systems.

Our main contributions are listed as follows:

- We propose a new control law for controlling the simplest periodic orbit of the bouncing ball, a prototype of rhythmic tasks where the actuator interacts with its environment. The proposed control law is somehow reconciling the “feedforward” approach (a sensorless actuation of the actuator) with the “feedback” approach (an actuation based on permanent tracking of the ball). We design a hybrid scheme (Fig. 1.2(c)) which minimizes the sensor design and sensory processing but maintains the closed-loop robustness. This approach contributes to the general knowledgebase in hybrid control.
- We generalize the one-dimensional bouncing ball dynamics to a planar wedgebilliard. The stabilization of its periodic orbits requires the actuation of *two* actuators, under bimanual coordination patterns.
- We analyze the sensorless stability properties of these periodic orbits. We demonstrate experimentally that the actual basins of attraction of some of these orbits are much smaller than predicted by the model. Thus the sensorless strategy is not *robust* enough.
- We generalize the hybrid scheme designed for the bouncing ball to the 2D juggler to achieve the stabilization of the complex periodic orbits. To the best of our knowledge, this constitutes the first realization of a 2D juggling robot which is able to fluently switch between different juggling patterns.
- We study the human behavior in performing the same juggling task with the same experimental setup. We report different strategy planning depending on the sensing capabilities, revealing a trade-off between performance and robustness at the level of human sensorimotor processing.

- We provide an extended description of our juggling setup.

1.4 Scope of the thesis

The rest of the thesis is organized as follows.

In Chapter 2, we propose an overview of the literature about rhythmic movements. We focus on what makes rhythmic movements particular, and why we chose to focus on them in this thesis. Section 2.2 particularly stresses the features of juggling as a representative rhythmic movement, both for the motor control and the robotics communities. A state-of-the-art overview in juggling robotics is also proposed.

In Chapter 3, we present the 1D dynamics of a ball bouncing on an actuator (or impactor). These very simple dynamics have been proved to be both intriguingly rich in their state-space description, and illustrative for analysis and design in underactuated systems. Consequently, both the robotics and motor control communities used bouncing ball experiments to investigate the particular mechanisms of trajectory planning in rhythmic environments.

In Chapter 4, we present a novel strategy for controlling the bouncing ball. This strategy was designed aiming at *minimizing* the need for sensory feedback (and consequently the need for sensor design) while maintaining a control which is both *robust* and rapidly converging. Convergence to time-varying reference is achieved in one impact, while the sole measured feedback information is the impact times. The chapter material has been published in Ronsse and Sepulchre (2006) and Ronsse et al. (2007a), Sections II to IV.

In Chapter 5, we describe the experimental juggling setup we used both to validate our control strategies in robotics experiments, and to acquire human behavioral data on similar paradigms. The construction of this setup was an important part of the project. The chapter material has been partly published in Ronsse et al. (2007a, 2006, 2007b).

In Chapter 6, we describe a set of periodic patterns corresponding to limit cycles of our juggling model. These periodic orbits are shown to be unstable, but a sinusoidal (i.e. open-loop) actuation of the juggler's arms stabilizes them in broad regions of the parameter space. Experimental results of open-loop stabilization of juggling patterns are also reported. The chapter material has been published in Ronsse et al. (2004, 2006).

Due to large discrepancies between the model and the actual setup, the experimental results are not completely convincing for complex periodic orbits. In Chapter 7, we generalize the minimum-feedback strategy of Chapter 4 to the model of our 2D planar juggler. The proposed strategy enlarges the basins of attraction of the open-loop control, just by requiring to measure the impact times. Illustrative experimental results are also provided. The chapter material has been published in Ronsse et al. (2007a), Sections II and V.

In Chapter 8, we analyze the behavior of human subjects when juggling the simplest periodic orbit with the same setup. We study the task performance under different experimental conditions, by changing the imposed task tempo, and by manipulating the visual feedback. The chapter reports the different control strategies which are adopted depending on these contexts. A publication about the chapter material is submitted (Ronsse et al., 2007c).

Finally, the thesis ends with a general discussion and raises some perspectives (Chapter 9). Parallels between the strategy adopted by the subjects in the degraded conditions, and the robust closed-loop design based on limited sensing are particularly emphasized.

Three appendices are added: Appendix A describes the trade-off between feedback and feedforward and between performance and robustness within a benchmark example. Through this example, it gives relevant insights in general control theory. Appendix B.1 provides the main technical details about the experimental setup described in Chapter 5, both for the robotics configuration (B.1) and the “human” configuration (B.2). Appendix C describes the computational technique we used to calculate the subjects’ gaze orientation in Chapter 8. This appendix material has been published in Ronsse et al. (2007d).

Chapter 2

Rhythmic Movements and Juggling

The trick to juggling is determining which balls are made of rubber and which ones are made of glass.

anonymous

This chapter explains the specifics of *rhythmic* movements, in particular juggling, and surveys important related contributions in neuroscience and robotics. Section 2.1 explains what makes rhythmic movements particular, and overviews the major contributions of the neuroscience literature in that field. Section 2.2 particularly stresses the advantages of juggling as a representative rhythmic movement, both for the motor control and the robotics communities. The juggling scientific literature is subsequently reviewed. In Section 2.2.4, we describe some connections of juggling to other rhythmic movements, in particular to locomotion.

2.1 Rhythmic movements

The aim of this section is to describe the specifics of rhythmic movements. The mechanisms of rhythmic movements, and the underlying coordination principles, make them different from discrete movements — reaching, aiming or pointing — for which a lot more computational models have been developed.

2.1.1 Rhythmic movements and the Central Pattern Generator paradigm

Rhythmic movements are phylogenetically old motor behaviors found in many organisms, ranging from insects to primates (Schaal et al., 2004). Indeed, rhythmic movements are involved both in locomotion (walking, running, hopping) and feeding activities (scratching, chewing), and are consequently necessary to life. Rhythmic

movements are also ubiquitous in human daily-life, ranging from basic functions to skillful abilities, like juggling or dancing.

In many species, rhythmic movements have been proved to be the output of dedicated neural circuitries, the Central Pattern Generator(s) (CPGs¹) (Marder, 2000; Marder and Bucher, 2001). More particularly in vertebrates, those CPGs are located in the spinal cord and the brainstem (see e.g. Cohen et al., 1988; Duysens and Van de Crommert, 1998; Swinnen, 2002). In higher vertebrates (including humans), CPGs have been more difficult to locate because the corresponding nervous structures are more complex, and the movements are supposed to be modulated by higher brain centers. However, evidences of CPGs structures have been proposed for human (Duysens and Van de Crommert, 1998; Marder, 2000), initially for locomotion while recent insights suggest that the concept of CPG applies to the upper limb as well (see e.g. Dietz, 2002; Zehr et al., 2004; White et al., 2007).

2.1.2 Rhythmic arm movements are not discrete

In contrast with rhythmic movements, *discrete* movements (such as reaching, grasping, pointing or kicking) have reached sophistication primarily in younger species, particularly primates (Schaal et al., 2004). Discrete movements are delimited in time, beginning and ending with pose periods. Moreover they are supposed to be sequenced with no clear periodicity in time and/or space. Discrete movements are programmed by a complex brain network, involving the cortex to a large extent (see e.g. Kalaska et al., 1997; Sabes, 2000; Desmurget et al., 2001).

For about a decade, researchers have been tracking the differences between discrete and rhythmic movements both at behavioral and imaging levels. Three distinct hypotheses on their relationship may be distinguished (van Mourik and Beek, 2004):

1. rhythmic movements are concatenated discrete movements;
2. discrete movements are a limit of rhythmic movements, aborted after a half-cycle; and
3. discrete and rhythmic movements are motor primitives that may be combined but are irreducible to each other.

Schaal, Sternad, Osu and Kawato (2004) recently conducted an imaging study to differentiate the brain areas involved in the production of similar rhythmic and discrete wrist movements (flexion – extension). They reported that similar areas are activated in the discrete movements as in complex reaching or pointing experiments. These are basically high-level cortical planning areas (see Fig. 2.1 and Schaal et al., 2004). In contrast, rhythmic movements show much less cortical and cerebellar activity: mostly motor areas are activated. These results strongly contradict the

¹A central pattern generator is a neural circuit that produces self-sustaining patterns of behavior independently of any sensory input (Swinnen, 2002).

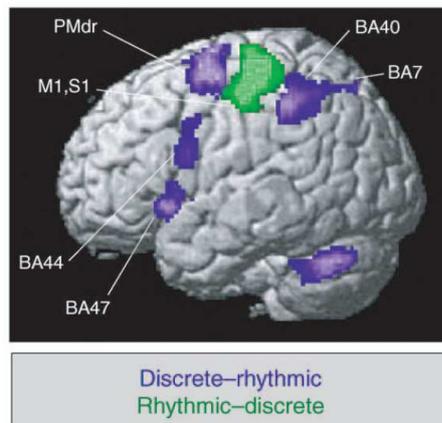


Figure 2.1: In blue, the brain areas involved in the production of a discrete wrist movement, but not in the production of a rhythmic wrist movement: the rostral part of the dorsal premotor cortex (PMdr), Broca's area (BA44), parietal cortex (BA7 and BA40) and the area B47. Widespread activation was also reported in the cerebellum. In green, the brain areas that were more involved in the production of the rhythmic wrist movement than in the production of the discrete wrist movement, i.e. the primary sensorimotor and premotor cortices (S1 and M1). Left hemisphere. Reprinted from Schaal et al. (2004).

first hypothesis: rhythmic movements are not concatenated discrete movements, since they do not recruit the high-level cortical areas involved in those movements production. However, it leaves the door open for choosing among the two remaining perspectives.

At the behavioral level, recent studies investigated to what extent discrete and rhythmic movements are related with each other (Sternad et al., 2000; de Rugy and Sternad, 2003; van Mourik and Beek, 2004; Buchanan et al., 2006). Sternad et al. (2000) and de Rugy and Sternad (2003) studied the interaction between discrete and rhythmic forearm movements in a combined experiment, requiring the production of both movements. Their major findings were: (1) The onset of the discrete movement was confined into a limited phase-window in the rhythmic cycle. (2) The duration of the discrete movement was influenced by the period of the oscillation. (3) The phase of the rhythmic oscillation was reset after a discrete stroke. They elaborated a mathematical model of this task, where the two movements were viewed as distinct primitives (Schaal and Schweighofer, 2005), described by two stable dynamical regimes of the model.

Differences between rhythmic and discrete arm movements exist also in their kinematic profiles (position, velocity, acceleration). Van Mourik and Beek (2004) compared such profiles in a reaching paradigm. Due to large differences between the two profiles, they further confirmed that rhythmic movements cannot be under-

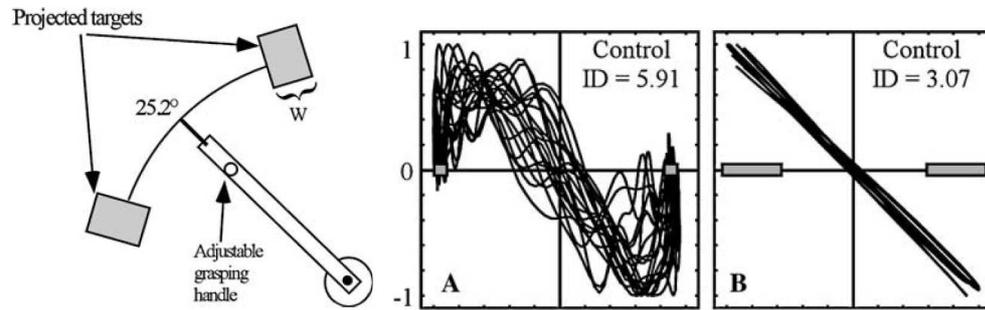


Figure 2.2: The left-hand side picture represents the experimental setup, showing the rotating handle and the targets of width W . The right-hand side panels show two typical phase portraits (position vs. acceleration) of the handle between narrow targets (large index of difficulty — $ID = 5.91$) and between wide targets ($ID = 3.07$). Reprinted from Buchanan et al. (2006).

stood in terms of concatenated series of discrete movements while the two others theories provide more plausible perspectives. More recently, Buchanan et al. (2006) investigated a task where a *transition* was forced between rhythmic and repeated discrete movements. The task consisted of repeated aiming actions between two targets of variable width. Narrow targets (i.e. high index of difficulty) were aimed with discrete movements, with pose intervals (zero acceleration on targets in Fig. 2.2.A). In contrast, the subjects switched back and forth in a continuous rhythmic movement when the targets were wider (i.e. low index of difficulty — acceleration in anti-phase with the position in Fig. 2.2.B). These data tend to support the third hypothesis — i.e. that continuous actions may be composed from either discrete or rhythmic units of action (or motor primitives) and that the discrete and rhythmic units of action are irreducible to each other.

Even in the absence of consensus, all the aforementioned studies agree to reject the first hypothesis: rhythmic movements are not concatenations of discrete strokes. Consequently, all the computational models described in the introduction must be handled with care in the context of rhythmic movements, since they have been elaborated in the discrete framework, mainly for trajectory planning. Our thesis studies the specifics of planning and controlling rhythmic movements.

2.1.3 The coordination of rhythmic movements

By definition, rhythmic movements are not sequential and are continuous in time. However, rhythmic tasks often require the recruitment of many degrees of freedom in parallel, hence requiring movement *coordination* (Bernstein, 1967; Kelso et al., 1979; Turvey, 1990). Coordination may be *intra-limb* (between several segments of a single limb) or *inter-limb* (between several limbs). Coordination is obviously mandatory in the production of some dynamical patterns: for example, locomotion

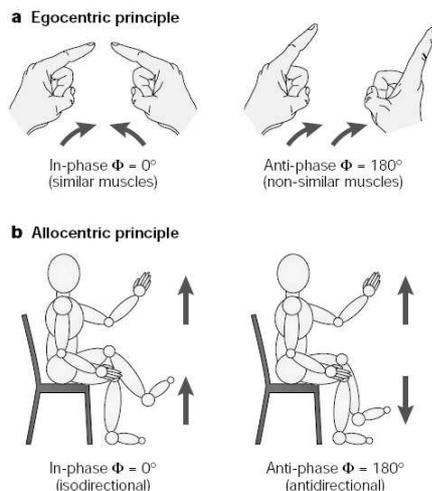


Figure 2.3: Basic coordination constraints: the egocentric and allocentric principles. The egocentric principle refers to a preference for moving according to mirror symmetry, which involves activating similar muscle groups simultaneously (a). The allocentric principle refers to a preference for moving the limbs or joints in the same direction in extrinsic space (b). Reprinted from Swinnen (2002).

patterns are stabilized through particular coordinated steps.

Limbs — or the individual limb segments — cannot be controlled arbitrarily. Coordination rules underly the possible movement patterns, and those principles cannot be inferred from the laws of single-joint or single-limb movements (Swinnen, 2002; Swinnen and Wenderoth, 2004). At the frequency level, the default mode of coordination is *synchronization*, ubiquitous in biological systems (the recent book by Strogatz, 2003, abounds with such examples). The *phase* relationships between the oscillating “agents” (joints and/or limbs) are also governed by coordination rules. The egocentric and allocentric principles are such basic coordination constraints, governing the preference for moving either in-phase, or in anti-phase (see Fig. 2.3 and Swinnen, 2002).

We do not aim at covering all the contributions provided in the domain of animal or human coordination. The interested reader is referred to the aforementioned reviews. Nevertheless, we aim at mentioning the seminal work by Kelso and coworkers (see Kelso, 1995) in this domain. Their fundamental ideas are appealing for the next section: the derivation of mathematical models.

The *dynamic pattern* theory (DPT) aims to show that “it is possible to understand behavioral pattern generation at several levels of description (kinematic, electromyographic, neuronal) by means of the concepts and tools of stochastic nonlinear dynamics” (Schoner and Kelso, 1988; Kelso, 1995). As a motivating example, Kelso (1984) reported phase transitions in a bimanual “index tapping” task. The task consisted of synchronized bimanual index tapping. The in-phase mode corresponds

to synchronized impacts, while the anti-phase mode corresponds to alternated impacts: one finger impacting when the other is at the apex of the trajectory. When the subject starts the task in the anti-phase mode, increasing movement frequency undoubtedly causes an abrupt transition toward the in-phase, mirror-symmetrical mode, which is more stable and less attentional demanding in the egocentric frame (see Fig. 2.3 — see also Kelso, 1995). The transition is viewed as a bifurcation in the stability diagram of cyclical patterns. This coordination task is recognized as a benchmark example in the motor control literature, and the related theoretical aspects are thought to apply to other coordinated movements. Switching between different gaits in animal behavior, as the movement frequency changes, is one of them (see e.g. Collins and Richmond, 1994).

The production of rhythmic movements must fulfill the intra- and inter-limb coordination rules. Since this thesis claims for an integration of rhythmic primitives within the computational tools for trajectory planning, the coordination rules must also be embedded into their internal representations. For robotics design, a robust implementation of the coordination mechanisms is thought to be of high relevance to achieve the stabilization of patterns whose stability depends on the accuracy of the coordination (e.g. locomotion).

2.1.4 Mathematical models of (coordinated) rhythmic movements

The transition between coordination modes can be described as bifurcations in the parameter space. Using the concepts of dynamical systems theory, the coordination patterns are stable limit cycles, whose stability is impaired as some parameters (e.g. the movement frequency) are modified. This is the central viewpoint of the Haken-Kelso-Bunz (HKB) model of coordination (Haken et al., 1985). This model captures the system behavior with a potential function V , which depends on the so-called order parameters ϕ (the relative phase between index fingers in the tapping task). Their dynamics are governed by $\dot{\phi} = \partial V / \partial \phi$, such that the system behavior can be described by identifying ϕ with the coordinate of a particle which moves in an overdamped fashion in the potential V (Haken et al., 1985). In the state space, the stable limit cycles thus correspond to the minima of the function V .

In the HKB model, the particular potential function takes the form:

$$V = -a \cos \phi - b \cos (2\phi) \quad (2.1)$$

where a and b are the parameters governing the transitions. The potential V/a is depicted in Fig. 2.4. As the ratio b/a decreases, the anti-phase pattern loses stability, and becomes unstable at $b/a = 0.25$ (see the black disks in Fig. 2.4). Below that value, the behavior switches to the in-phase pattern and remains in this coordinated mode, even if b/a increases back above 0.25. This corresponds to the observed data in the “index tapping” task, since the fingers remained in-phase as the frequency decreased.

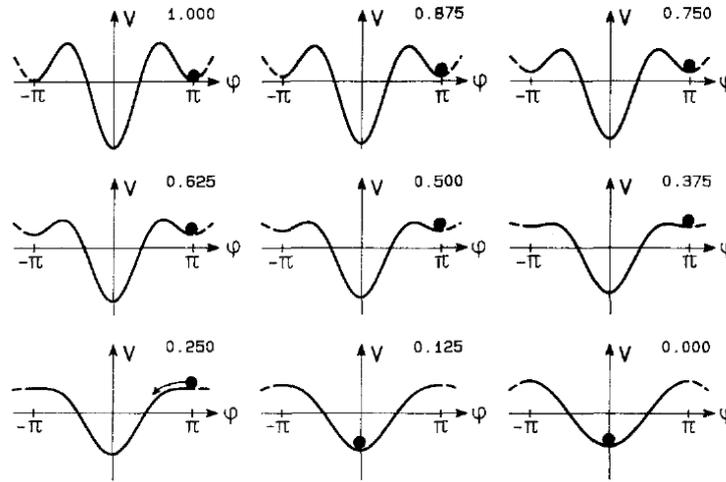


Figure 2.4: The potential V/a for the varying values of b/a , as referred by the numbers. The bottom-left panel corresponds to the critical value, where the anti-phase ($\phi = \pi$ or $-\pi$) loses stability and where the system is “forever” attracted into the in-phase coordination mode. Reprinted from Haken et al. (1985).

One still has to establish how the potential function (2.1) could emerge from the dynamics of the individual agents, that are the finger, muscular and neural dynamics. The HKB model proposes a set of two non-linearly coupled non-linear oscillators to capture the bifurcation into the potential function (2.1) (Haken et al., 1985).

Other contributions have focused on the particular structure of the non-linear oscillating system which causes sustained rhythmic movements. In the papers by Sternad et al. (2000) and de Rugy and Sternad (2003), the rhythmic pattern generator is based on a half-centered oscillator model, formalized through a set of leaky integrator equations by Matsuoka (1985, 1987).

Kuo (2002b) provided another interesting viewpoint on the generation of rhythmic activities. This contribution is particularly appealing in this thesis context, since it establishes the relative roles of feedforward and feedback in the control of rhythmic movements, at the level of the neural CPG. While a purely feedforward CPG is highly sensitive to unexpected disturbances, pure feedback control — analogous to reflex pathways — can compensate for disturbances, but is sensitive to imperfect sensors. The balance between both control mechanisms appears since the “optimal” trade-off between robustness to noise and imperfect sensors is reached through a proper combination of feedforward and feedback control. Moreover, with this combined mechanism, the CPGs can still produce rhythmic trajectory through the feedforward path when sensory output is removed, as observed biologically. Kuo’s model is both biologically plausible, and provides behaviorally consistent simulated data on a pendulum model of the limb.

2.1.5 Conclusion

The control of rhythmic movements is different from the control of discrete movements, for which the computational tools for optimal planning provide useful insights. Both at the modeling level and through the analysis of biological data sets, rhythmic movements are supposed to be produced by lower-level Central Pattern Generators (CPGs). Thus, rhythmic trajectory planning is not achieved through a segmentation of the movement, but as the asymptotically stable limit cycles of the corresponding oscillating circuit. Nonetheless, rhythmic movements can be controlled through, for example, appropriate modulation of their cycle phase and/or amplitude. In this case, however, the movement planning does not encompass the whole trajectory, but only the desired timing and amplitude.

This thesis perspective is to establish how these rhythmic movements are actually controlled, in the context of a particular task; and how the available sensory feedback involved in the loop influences the mode of control, and the related coordination rules. This twofold perspective is our general guideline throughout the manuscript, both when studying human behavior and robotics designs.

2.2 Juggling at the crossroad

This section introduces the particular rhythmic movement which is considered in the present manuscript, i.e. juggling.

2.2.1 Juggling is a representative rhythmic movement

The paper by Beek and Lewbel (1995), that vulgarizes some scientific aspects of juggling, opens with the following funny anecdote: “To complete a delivery of munition, a 148-pound man must traverse a high, creaking bridge that can support only 150 pounds. The problem is, he has three, one-pound cannonballs and time for only one trip across. The solution to this old riddle is that the man juggled the cannonballs while crossing. In reality, juggling would not have helped, for catching a tossed cannonball would exert a force on the bridge that would exceed the weight limit. The courier would in fact end up at the bottom of the gorge.”

While not always the ultimate solution of mechanical problems, juggling has nevertheless been recognized as a skillful art, requiring the production of rhythmic movements in a highly coordinated manner. The earliest known depiction of toss juggling is Egyptian, from the 15th Beni Hassan tomb of an unknown prince, dating from the middle kingdom period of about 1994-1781 B.C. (see Fig. 2.5 and Lewbel, 2002). From that time, mainly three fields of scientific investigations have benefited from the intriguing properties of juggling as a benchmark rhythmic movement (Beek and Lewbel, 1995):

1. the study of human movements and coordination;

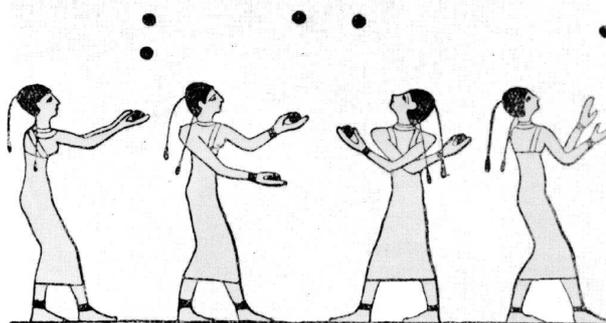


Figure 2.5: One of the earliest representations of juggling, at the ancient Egyptian age. The image source comes from Lewbel (2002).

2. the development of juggling machines in robotics, useful to catch the real-time necessity of the mechanical control of such underactuated systems;
3. the mathematics, through the surprising numerical properties of juggling patterns.

The two first of these fields are exactly within the scope of this thesis, in which juggling is again recognized as a useful benchmark to investigate the role of sensory feedback in the subsequent control strategy. The mathematical aspects of juggling patterns are not covered in this thesis. We nevertheless mention the initial “constructive” theorem of mathematical relationship in juggling patterns, since it has been proposed by an engineer and founder of information-theory, Claude E. Shannon² (Shannon, 1993):

Theorem 1 (Shannon, 1993) *Given N the number of objects, and H the number of hands involved in a juggling pattern, the following equation must be fulfilled during steady-state juggling cycles:*

$$(F + D)H = (V + D)N \quad (2.2)$$

where F is the time a ball spends in the air, D is the time a ball spends in a hand and V is the time a hand is vacant.

This theorem is illustrated in Fig. 2.6 for the three-ball cascade: a figure-eight pattern (see Fig. 2.7). The three-balls cascade is certainly the most fundamental juggling pattern, by which many juggling neophytes start their learning. Interestingly, Shannon’s theorem is nevertheless valid for any juggling pattern in which no

²Claude Elwood Shannon (1916-2001) was extremely influent in the early development of computers and digital communication. In 1990, *Scientific American* called his paper on information theory, “The Magna Carta of the Information Age”, from A. Lewbel’s personal tribute to Claude Shannon, www2.bc.edu/~lewbel/Shannon.html.

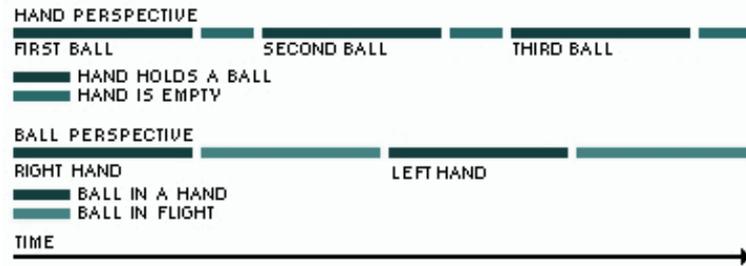


Figure 2.6: The juggling theorem proposed by Shannon (1993) is schematically represented for the three-ball cascade. The theorem is proved by following one complete cycle of the juggle from the point of view of the hand and of the ball and then equating the two. Reprinted from Beek and Lewbel (1995), adapted from Shannon (1993).

hand holds more than one ball at any one instant of time, regardless of the sizes and shapes of the juggled objects, the postures and limb configurations of the juggler, and the species of the juggler (human or robot) (Beek and Turvey, 1992)!

Shannon’s theorem provides a useful clarification on the exact coordination requirements in juggling patterns: not only the hands (or the limbs) have to be synchronized and coordinated together, but also they have to be coordinated with the juggled objects. Indeed, the objects dynamics during the flying phases cannot be influenced by the juggler(s), while their flying periods influence the hand trajectory, via Shannon’s equation (2.2). Juggling is then really a closed-loop process: the juggled objects dynamics are obviously influenced by the hands via the catching phases, and influence also the hands trajectory through the requested coordination rule captured by (2.2).

Since juggling requires the stabilization of the interlimb pattern (in-phase?, anti-phase?, others?) and the stabilization of the external environment (the juggled objects), it legitimately serves as an illustrative framework for considering the trade-off between efficient and robust control. The control efficiency (or performance) is understood both in terms of (1) the expended *energy* and (2) the trajectory variability (i.e. the extent to which the trajectory varies around the steady-state cycle). The control robustness refers to the controller ability to maintain the juggling pattern stable despite uncertainties or perturbations in the environment.

2.2.2 A dynamical systems perspective in the cascade juggling

Beek and colleagues deeply investigated the learning mechanisms involved in the cascade juggling, and formalized a so-called *dynamical systems perspective* for this analysis (see e.g. Beek and van Santvoord, 1992; Beek and Turvey, 1992; van Santvo-

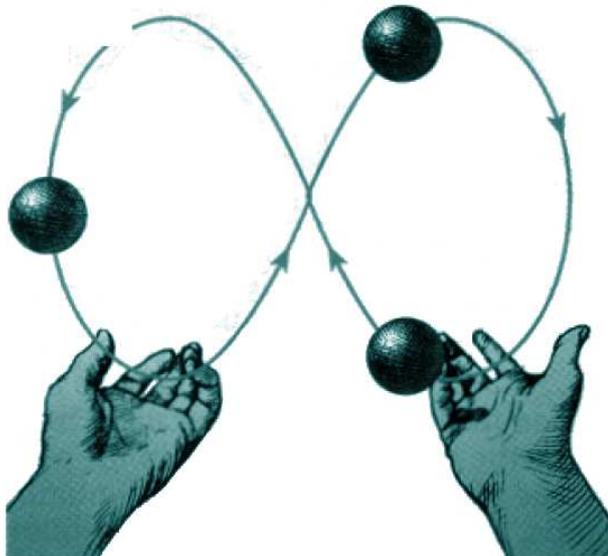


Figure 2.7: The three-balls cascade juggling pattern. Reprinted from Beek and Lewbel (1995).

ord and Beek, 1996; Post et al., 2000; Huys and Beek, 2002; Huys et al., 2003; Huys, 2004).

Given Shannon's theorem (2.2), for a fixed number of hands H and balls N , one of the remaining time quantities is constrained by the other two. Moreover, assuming that the global period of the complete cycle is fixed — and dictated for example by a metronome or by the desired juggling height — one degree of freedom is still remaining in the juggler's strategy. Beek and van Santvoord (1992) proposed consequently a three-stage model of the learning process of the metronome-paced three-balls cascade ($H = 2$, $N = 3$):

- The first stage consists in learning to accommodate the real-time requirements of juggling, as expressed in Shannon's equation of juggling (2.2).
- The second stage of learning consists in discovering the primary frequency lock of 0.75 between the shorter term dynamical regime underlying the repetitive subtask of transporting a ball (D in (2.2)) and the longer term dynamical regime underlying the total hand loop cycle ($V + D$ in (2.2)).
- The third and last stage of learning consists in discovering the principles of frequency modulation from 0.75 to lower (averaged) values of the proportion of time that a hand carries a ball during the total hand cycle time.

From the perspective of the trade-off between performance and robustness, these three learning stages could be stated differently. The first stage consists simply in fulfilling the task, by adopting a coordinated behavior which is both a limit

cycle solution of the system, and stable. The second stage is definitely a matter of robustness, since the “dwell-ratio” $D/(V + D) = 0.75$ has been shown to be the more robust frequency lock for juggling the cascade with three (Beek and van Santvoord, 1992) or more than three (Beek and Turvey, 1992) balls. Finally, the third stage relaxes the need for robustness, since the juggling pattern is assumed to be properly mastered. The juggler adopts smaller frequency locks than $D/(V + D) = 0.75$, in which the average number of airborne balls is consequently larger. This can be viewed as a performance improvement, increasing the control flexibility.

Later, Huys and colleagues investigated how the learning and expertise in the cascade juggling could also affect the coupling with other functional subsystems, such as the point-of-gaze, the respiration and the body sway (Huys and Beek, 2002; Huys et al., 2003; Huys, 2004). Their results indicated that dissimilar learning dynamics may arise in the functional embedding of subsystems into such a task-specific organization (Huys et al., 2003). More particularly, Huys and Beek (2002) revealed an strong coupling between the balls trajectories and the point-of-gaze around balls’ apex, i.e. the highest point of their trajectories (see also Amazeen et al., 1999). In that region, the subjects made not only position but also velocity tracking of the balls, through appropriate frequency locks between the balls and the point-of-gaze.

In conclusion, juggling has been used for many years in a system-theoretic perspective to illustrate the functional organization of a complex and coordinated rhythmic movement. Learning mechanisms have consequently been emphasized, as an illustration of the acquisition of task-related ability and *flexibility*. While focusing not primarily on learning issues, this thesis is investigating a similar trade-off between ability (performance) and flexibility (robustness) in a simplified juggling paradigm. Our viewpoint is to assess whether this trade-off is influential for trajectory planning.

2.2.3 Juggling robots

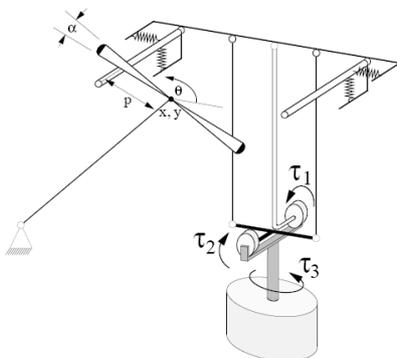
In the robotics community, juggling has also been an intriguing source of inspiration for the development of skilled robots (see Fig. 2.8). The first known juggling robot has also been manufactured by Claude E. Shannon (see Fig. 2.8(a)). He built a machine essentially consisting of a motor attached at the center of a rod which has two catchers mounted at each end. By driving the motor sinusoidally and adjusting the distance of the catchers, the motor frequency and amplitude, and the height of the setup above the floor, it is possible to find a configuration in which the balls are juggled in a stable fashion, without need of feedback from their current state (open-loop control). A drum was used to provide an elastic floor. Juggling three balls requires one full oscillation during the flight of a ball (Schaal and Atkeson, 1993). Shannon’s other contribution to juggling robotics concerns his famous diorama (see Fig. 2.8(b)). This is obviously not really a juggling robot, but the balls, rings, and clubs, and clowns hands all moved realistically. A movie of Shannon’s juggling machines can be found on A. Lewbel’s homepage at www2.bc.edu.



(a) 1970s: Shannon's juggling robot, from Shannon (1993)



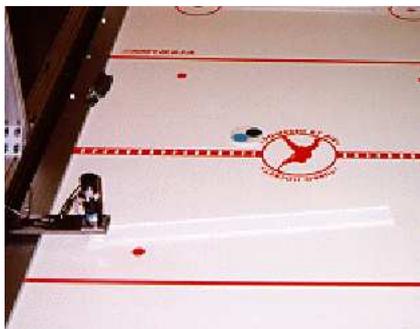
(b) 1982: Shannon's diorama, from Shannon (1993)



(c) 1991: van Zil's devil sticking robot, from Schaal and Atkeson (1993)



(d) 1992: Rizzi and Koditschek's juggling robot, from Rizzi and Koditschek (1993)



(e) 2000s: Flatland, from Lynch and Black (2001)



(f) 2000s: Sarcos "DB" robot juggling the 3-balls cascade, from Atkeson et al. (2000)

Figure 2.8: Juggling robots.

`edu/~lewbel/Shannon.html`.

A decade latter, Schaal and Atkeson (1993) reported the existence of a “devil sticking” juggling robot. Devil sticking requires manipulating a center stick with two hand sticks by hitting the center stick back and forth between the hand sticks. No picture of the robot has been found, but it is sketched on Fig. 2.8(c). Pioneering work investigating robotic tasks in rhythmic contexts has been done by Buehler, Koditschek and Kindlmann (1988, 1990, 1994). They developed the famous *mirror law algorithms*, in which tracking feedback of the juggled objects is used to robustly synchronize the robot with the juggling pattern. The simplest version of these algorithms is described in Section 3.3 for the 1D bouncing ball dynamics. However, they have also been adapted to complex environments. Fig. 2.8(d) depicts a 3D juggling robot developed by Rizzi and Koditschek (Rizzi et al., 1992; Rizzi and Koditschek, 1992, 1993), which implemented the mirror law algorithms to vertically bounce two ping-pong balls in 3D space.

Flatland is a planar (2D) robot built by Lynch and colleagues (see e.g. Lynch and Black, 2001), see Fig. 2.8(e). It is also based on a vision-system to extract relevant state feedback information from the objects dynamics. This robot architecture — based on a tilted air-hockey table providing frictionless motion of the juggled pucks — is appealing in this thesis context since it directly inspired the design of our own juggling robot, presented in Chapter 5. Lynch and Black’s control strategy is based on the real-time extraction of the puck state, in order to anticipate its trajectory, and to produce adapted control actions in consequence.

The most developed juggling robot constructed so far is certainly the Sarcos “DB” robot, since it can juggle fluently the three-balls cascade (see Fig. 2.8(f)). This 30 degrees-of-freedom robot has been built by the ERATO brain project in Japan (www.cns.atr.jp) and has been widely used to reproduce and analyze complex human behaviors in a broad set of tasks (see e.g. Atkeson et al., 2000).

2.2.4 From juggling to locomotion

Juggling served as benchmark for investigations in a broad set of other rhythmic tasks, both in the motor control and in the robotics literature. Juggling, bipedal locomotion, robot gymnastics, and robot air hockey are fundamentally related to the control of redundant and underactuated systems and share indeed some interesting common features (Spong, 1999). Locomotion is certainly of particular importance since common to many animal species and humans, while a direct analogy can be established with juggling:

- the locomotor limbs corresponding to the juggler hands;
- the body corresponding to a single juggled object.

With $H = 2$ and $N = 1$, Shannon’s theorem (2.2) may be consequently adapted to bipedal locomotion, the “dwell-ratio” $D/(V + D)$ referring now to the fraction of the

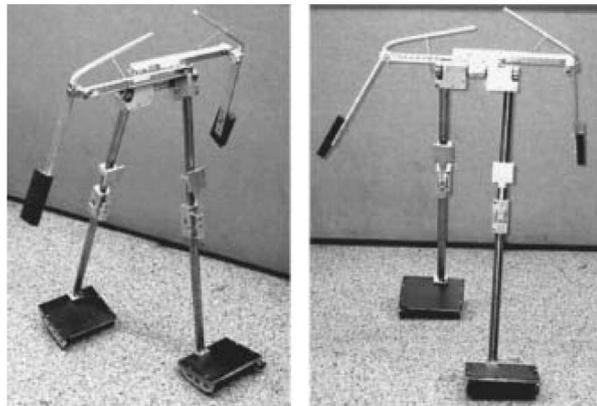


Figure 2.9: Collins, Wisse and Ruina’s bipedal passive-walking robot. Reprinted from Collins et al. (2001).

leg time period during which the leg is contacting the ground (the so-called stance phase). This ratio tuning, such as the phase relationship between the coordinated limbs, defines obviously the different gait patterns. Gait transitions are thus similar to transitions in juggling patterns, and the switching strategies could be studied in parallel.

Locomotion is a major field of investigation in robotics. Examples include both multipod robots (see e.g. Saranli et al. (2001), or the recent review by Holmes et al. (2006)) or biped walking (see e.g. the RABBIT project, as described by Plestan et al., 2003; Westervelt et al., 2004). Historically, the first attempts to tackle the problem of biped locomotion synthesis exploited the concept of *passive walking* (McGeer, 1990; Goswami et al., 1998): passive walkers travel down a gentle slope and walk in a stable, passive, three-dimensional gait, without any source of external energy. Stability analysis of these orbital gaits show that they are asymptotically stable, resulting from an optimal balance of the energies involved in the system: the potential energy is transformed into kinetic energy, which is lost in turn at impacts. A 3D passive-dynamic walking robot with two legs and knees has been studied by Collins et al. (2001) (see Fig. 2.9). Later, Collins et al. (2005) studied the relevance of passive-based architectures for the design of actively powered walking robots. Their paper describes three robots based on passive-dynamics, with small active power sources substituting for gravity, which can consequently walk on ground level (see also Kuo, 2002a). Due to their passive-based architectures, these robots use less control and less energy than other powered robots, yet walk more similarly to bipeds. This further suggests the importance of passive-dynamics in human locomotion, and places consequently the study of passively-based locomotion as another source of cross-fertilization between robotics and human (or animal) behavior.

Passive-based locomotion designs — and earlier studies on hopping systems (Raibert, 1986) — can be interestingly paralleled with a particular class of jug-

gling, namely *impact* juggling. In impact juggling, the contact between the hand and the object is supposed to be instantaneous³, such that the object energy is potentially not completely dissipated through the impacts. An academic example of impact (or bounce) juggling system has been widely investigated in both the robotics and motor control literature. It refers to the 1D motion of a bouncing ball, and the related literature is overviewed in Chapter 3. Impact juggling nicely connects with passive-based locomotion since a broad set of impact juggling patterns can be stabilized through passive control. Here, passive control is not understood in the sense that no energy supply is provided to the system (obviously, the impactor is actuated), but refers to control strategies that are *sensorless* stable: i.e. no feedback is needed from the state of the juggled objects to maintain the pattern stability. Open-loop asymptotic stability of bounce juggling patterns is obtained through a simple sinusoidal actuation of the impactor(s). This has been studied in 1D (Holmes, 1982; Guckenheimer and Holmes, 1986) and 2D (Schaal and Atkeson, 1993) juggling movements. The present manuscript focuses also on the 2D impact juggler and studies how to stabilize *several* impact juggling patterns through actuation of the arms. More particularly, Chapter 6 describes sensorless (i.e. passive) strategies which stabilize these patterns.

2.2.5 Conclusion

Juggling is a benchmark for the study of rhythmic movements, requiring both the stabilization of a particular bimanual coordination pattern, and the stabilization of external object(s). It has been used both for investigations in motor learning and control, and for the design of robots performing in rhythmic environments. Juggling is connected to other rhythmic movements, including locomotion.

2.3 Concluding remark

The planning of rhythmic movements in general, and juggling movements in particular, is different from the planning of discrete movements. This is due to (1) the difference in the neural circuitries involved in the production of both movements; (2) the underlying coordination principles governing rhythmic movements; and (3) the extent to which these movements are influenced by the sensory inflows (passive / active control). We focus on juggling experiments, claiming that they are representative of the whole class of rhythmic movements.

The guideline for the rest of this thesis is to investigate how the planning of juggling movements is achieved with respect to the trade-off between performance and robustness, and how the available sensory feedback influences the control strategy.

³In Shannon’s equation (2.2), this means that D equals 0. This juggling “strategy” consists in trying to maximize the average number of juggled objects in the air, and is consequently often referred as “hot potatoes juggling” by expert jugglers.

Chapter 3

The Bouncing Ball

Success is how high you bounce
when you hit the bottom.

General George S. Patton

3.1 Introduction

This thesis highlights parallels between robotics and motor control (Schaal and Schweighofer, 2005) in the particular context of rhythmic tasks. One of such task has been widely investigated by both communities in the two last decades and is consequently introduced in this chapter as an illustrative benchmark. The *bouncing ball* model describes the movement of a ball that periodically bounces on an actuated impactor, e.g. a racket. This task is illustrative of situations where an effector (i.e. either a human or a robot) interacts with an object in the environment (de Rugy et al., 2003). The control is underactuated, since the robot degrees of freedom are fewer than the object degrees of freedom (Lynch and Black, 2001).

The bouncing ball dynamics have been initially studied by Holmes (1982) under a particular actuator trajectory: a simple sinusoidal motion. These dynamics turned out to become one of the simplest example in non-linear dynamics, which exhibits deterministic chaos in a given range of the sinusoidal amplitude. The bouncing ball dynamics indeed produce a bifurcation route that is similar to the well-known logistic map (see e.g. Tufillaro and Albano, 1985; Tufillaro et al., 1992). The main features of the bouncing ball dynamics are reviewed in Section 3.2.

The most illustrative problem when considering stabilization of ball-bouncing patterns is to stabilize its elementary periodic orbit, i.e. a succession of bounces at a constant height (see Fig. 3.2). In Section 3.3, we briefly review the major contribution by Buehler, Koditschek and Kindlmann in the design of the so-called *mirror law algorithms*. These designs have long been recognized as pioneering investigations in the context of rhythmic robotics. The mirror law robustly stabilizes

the ball bouncing at a constant height. Moreover, wide basins of attraction have been empirically observed in a broad range of experimental contexts.

The behavior of humans when “juggling” the bouncing ball has been studied by Sternad, Schaal and coworkers (Schaal et al., 1996; Sternad, 1999; Sternad et al., 2001a,b; Katsumata et al., 2003; de Rugy et al., 2003; Dijkstra et al., 2004; Wei et al., 2006). Their major contributions are overviewed in Section 3.4.

Juggling has been mentioned as a relevant benchmark for rhythmic motor control tasks (Section 2.2.1), requiring different levels of coordination. The bouncing ball paradigm, even if unimanual, is a good example of a juggling task that shares a lot of commonalities with “regular” juggling (Sternad, 1999): spatial and temporal constraints, sensorimotor processing, coordination, etc... More generally, both lines of research conducted theoretical analysis to address questions of movement control, perception, and learning; while both of them have been investigated from a dynamical systems perspective (Sternad, 1999; Sternad et al., 2001a).

This chapter objective is not to cover all the modeling and design investigations that have been made on the bouncing ball dynamics¹. Instead, we aim at reviewing its basic properties, and the major contributions from the robotic and the motor control communities. Many of the core results of this thesis can be understood on this simple benchmark.

3.2 Open-loop dynamics of a ball bouncing on a sinusoidally vibrating racket

3.2.1 Bouncing ball model

The dynamics of a ball bouncing on an actuated racket is hybrid (Holmes, 1982; Guckenheimer and Holmes, 1986). During the flight times, the ball follows a ballistic parabolic flight (see Fig. 3.1). The position of impact therefore obeys the following discrete-time flight map, derived from Newton’s law:

$$s(t[k+1]) = s(t[k]) + v^+(t[k])(t[k+1] - t[k]) - \frac{g}{2}(t[k+1] - t[k])^2 \quad (3.1)$$

where $s(t)$ denotes the continuous trajectory of the racket and $v(t)$ is the ball velocity. The time of two successive impacts, namely the k^{th} and $(k+1)^{\text{th}}$, are denoted $t[k]$ and $t[k+1]$ and the $+$ superscript in (3.1) denotes the post-impact velocity accordingly, since the ball velocity is discontinuous at impact. g is the constant of gravity. Similarly, the pre-impact velocity, $v^-(t[k+1])$ is equal to:

$$v^-(t[k+1]) = v^+(t[k]) - g(t[k+1] - t[k]). \quad (3.2)$$

¹On January 3, 2007, GOOGLE SCHOLAR[®] pointed out 2,620 contributions for the tag “bouncing ball”.

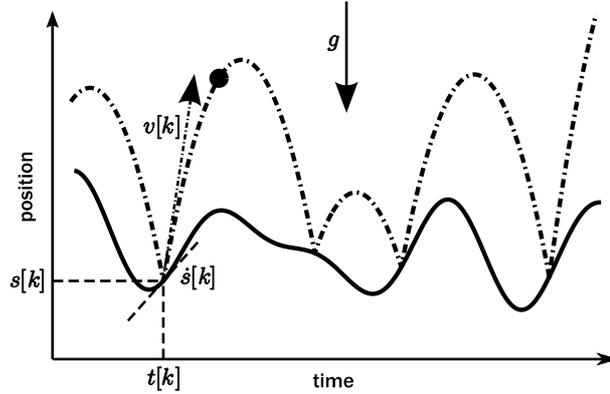


Figure 3.1: 1D bouncing ball. The racket (respectively the ball) trajectory is depicted with solid (respectively dash-dotted) lines over time. At time $t[k]$ (k^{th} impact), the actuator (and ball) position is $s(t[k])$, the actuator velocity is $\dot{s}(t[k])$ and the ball post-impact velocity is $v[k] = v^+(t[k])$.

Based on Newton's law, the relative velocity of the ball with respect to the actuator is reversed at impact and multiplied by the *coefficient of restitution* $0 \leq e \leq 1$ that models the energy dissipation:

$$v^+(t[k+1]) - \dot{s}(t[k+1]) = -e(v^-(t[k+1]) - \dot{s}(t[k+1])). \quad (3.3)$$

Equation (3.3) assumes that the actuator motion is unaffected by the impacts. This assumption is valid if the actuator is largely heavier than the ball (the inertia of the actuator is much larger than the inertia of the ball).

The complete bouncing ball dynamics are therefore described by the discrete *Poincaré map*, whose state is the impact position $s[k] = s(t[k])$ and post-impact velocity $v[k] = v^+(t[k])$, see Fig. 3.1:

$$s[k+1] = s[k] + v[k](t[k+1] - t[k]) - \frac{g}{2}(t[k+1] - t[k])^2, \quad (3.4)$$

$$v[k+1] = -e v[k] + e g (t[k+1] - t[k]) + (1 + e)\dot{s}[k+1] \quad (3.5)$$

where $\dot{s}[k] = \dot{s}(t[k])$. Equation (3.4) is the flight map and (3.5) is the impact rule, derived from (3.2) and (3.3).

The flight time, i.e. the time elapsed during two consecutive impacts, is deduced from (3.4):

$$t[k+1] - t[k] = \frac{v[k] + \sqrt{v[k]^2 - 2g(s[k+1] - s[k])}}{g}. \quad (3.6)$$

3.2.2 Sinusoidal actuation

Holmes (1982) studied the bouncing ball dynamics under a special racket trajectory,

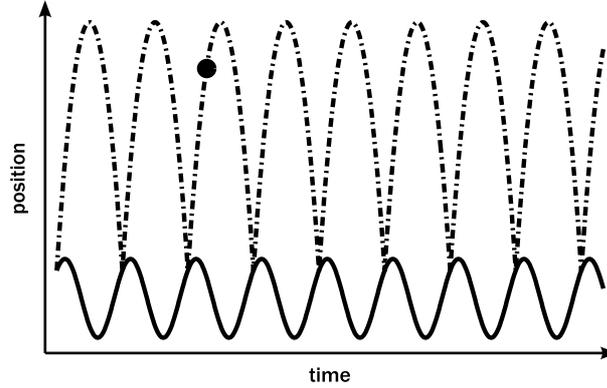


Figure 3.2: Period-one orbit of the 1D bouncing ball under sinusoidal actuation. The racket (respectively the ball) trajectory is depicted with solid (respectively dash-dotted) lines over time.

i.e. a simple sinusoidal motion:

$$s(t) = A \sin(\omega t) \tag{3.7}$$

where A and ω denote the movement amplitude and frequency (pulsation), respectively.

The ball dynamics under this sinusoidal actuation are astonishingly rich. Depending on the amplitude and frequency, several periodic orbits of the model are stable. The parametric stability regions of these periodic orbits are mutually exclusive, such that the ball steady-state trajectory follows a bifurcation route of period doubling as the amplitude (or the frequency) increases.

The simplest periodic orbit is the fixed point of (3.4) and (3.5), i.e. constant impact position and post-impact velocity. It corresponds to a train of bounces at constant height (see Fig. 3.2):

$$v^* = \frac{g}{2} \Delta t^*, \tag{3.8}$$

$$\dot{s}^* = \frac{1 - e}{1 + e} \frac{g}{2} \Delta t^*. \tag{3.9}$$

It is called the period-one, and forces obviously the steady-state flight time Δt^* to be equal to a multiple $n \in \mathbb{N}$ of the racket period:

$$\Delta t^* = n \frac{2\pi}{\omega} \tag{3.10}$$

where $n = 1$ when there is one racket period between two impacts. Assuming (3.10),

the fixed point of (3.4) and (3.5) is given by:

$$v^* = \frac{n\pi g}{\omega}, \quad (3.11)$$

$$\begin{aligned} s^* &= A\omega \cos \phi^* \\ &= \frac{1 - e n\pi g}{1 + e \omega} \end{aligned} \quad (3.12)$$

and the steady-state impact phase is equal to:

$$\phi^* = \arccos \left(\frac{1 - e n\pi g}{1 + e A\omega^2} \right). \quad (3.13)$$

The steady-state impact position is $s^* = A \sin \phi^*$.

Local stability of this periodic motion is established through the linearization of (3.4) and (3.5) around the steady-state (3.11) and (3.13). It gives the following equations:

$$\delta t[k + 1] = \delta t[k] + \frac{1 + e}{g} \delta v[k], \quad (3.14)$$

$$\delta v[k + 1] = e^2 \delta v[k] - (1 + e) A \omega^2 \sqrt{1 - \left(\frac{(1 - e) \pi n g}{(1 + e) A \omega^2} \right)^2} \delta t[k + 1] \quad (3.15)$$

where δt and δv denote the first-order small perturbations on the impact time, and post-impact velocity, respectively. Injecting (3.14) into (3.15), one obtains the following linearized non-dimensional matrix form:

$$\begin{pmatrix} \frac{\omega}{g} \delta v[k + 1] \\ \omega \delta t[k + 1] \end{pmatrix} = \underbrace{\begin{pmatrix} e^2 - \frac{(1 + e) A \omega^2}{g} \sqrt{1 - \left(\frac{(1 - e) \pi n g}{(1 + e) A \omega^2} \right)^2} & - \frac{(1 + e)^2 A \omega^2}{g} \sqrt{1 - \left(\frac{(1 - e) \pi n g}{(1 + e) A \omega^2} \right)^2} \\ 1 & 1 + e \end{pmatrix}}_{\mathbf{A}_{BB}} \begin{pmatrix} \frac{\omega}{g} \delta v[k] \\ \omega \delta t[k] \end{pmatrix}. \quad (3.16)$$

It can be shown that the eigenvalues of \mathbf{A}_{BB} lie into the unitary circle if and only if the following condition holds (Bapat et al., 1986):

$$\pi n \frac{1 - e}{1 + e} < \frac{A \omega^2}{g} < \sqrt{\pi^2 n^2 \left(\frac{1 - e}{1 + e} \right)^2 + \frac{4(1 + e^2)^2}{(1 + e)^4}}. \quad (3.17)$$

This corresponds to the amplitude and frequency range of stability for the period-one motion, depending on the coefficient of restitution e .

The solution and the parametric stability region of the period-two (the bounces alternate at two different heights) could be found from the fixed points of a double iteration of (3.4) and (3.5), according to similar derivations (Bapat et al., 1986). The lower limit of the parametric stability region is given by the upper bound in (3.17), i.e. the right-hand side term. By further increasing the racket amplitude (or frequency), the period-two loses stability for a period-four trajectory, then a period-eight, etc... along a route of period-doubling to chaos.

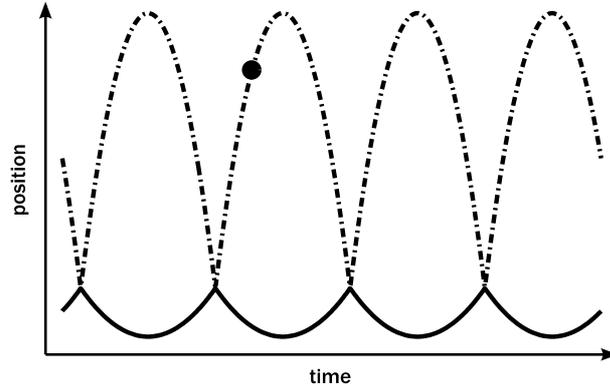


Figure 3.3: Period-one orbit of the 1D bouncing ball under mirror law control. The racket (respectively the ball) trajectory is depicted with solid (respectively dash-dotted) lines over time.

3.3 Robust feedback control of bouncing robots

As already mentioned in Section 2.2.3, early juggling robots were constructed to implement one-dimensional bounce juggling. These robotic developments are due to the seminal work of Buehler, Koditschek and Kindlmann (1988, 1990, 1994). These authors designed the so-called *mirror law algorithms* that turned out to be robust feedback control laws to stabilize sustained period-one bouncing trajectories in 1D, 2D and even 3D environments.

For simplicity, we present only the 1D version of the mirror law, assuming therefore that the ball motion is restricted to one dimension. This mirror law is based on permanent tracking of the ball trajectory $\beta(t)$, since the racket trajectory is computed to mirror the ball:

$$s(t) = \frac{-(1-e)}{1+e}\beta(t) - \kappa_1 (E_\rho^* - E(t))\beta(t). \quad (3.18)$$

The first term of (3.18) is just mirroring the ball trajectory. The second term is a proportional feedback that is used to isolate a particular period-one pattern, characterized by its energy level: $E_\rho^* = gs_\rho^* + 0.5(v_\rho^*)^2$, through permanent comparison with the ball energy $E(t)$. The gain κ_1 will determine the dynamics of the closed-loop system. In steady-state, that is when $E(t) = E_\rho^*$, the mirror law behavior is depicted in Figure 3.3.

The mirror law, as defined by (3.18), sharply contrasts with the sinusoidal law defined in (3.7) in term of feedback requirement. On the first hand, the sinusoidal law was purely sensorless and stabilized periodic orbits thanks to their open-loop stability properties. On the other hand, the mirror law requires permanent tracking of the juggled objects to compute their position and energy, and consequently led to robust implementations in various environments.

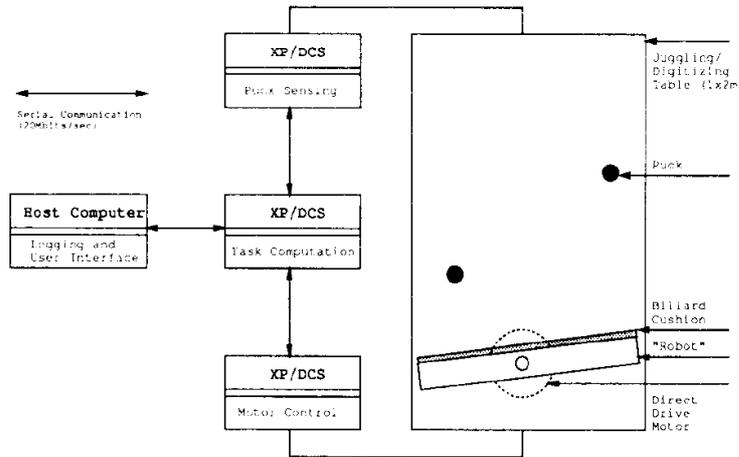


Figure 3.4: The planar juggler with simultaneous juggling of two pucks. Reprinted from Buehler et al. (1994).

Further developments of the mirror law algorithms led to the stabilization of more complex juggling patterns. For example, simultaneous vertical juggling of *two* pucks in anti-phase has been realized with a 2D planar juggling robot, depicted in Fig. 3.4 (Buehler et al., 1994). In Section 2.2.3, we also described the 3D adaptation of the mirror law for the juggling robot designed by Rizzi and Koditschek (1993) (see Fig. 2.8(d)).

Related publications and an illustrative movie can be found on Martin Buehler’s web page at www.martinbuehler.net.

An alternative feedback method to control periodic motions of the bouncing ball has been developed by Vincent and Mees (2000). This method is based on the sinusoidal trajectory (3.7): the controller output is the motion frequency $u = \omega - \omega^*$ (where ω^* denotes the steady-state value) and is computed on the basis of the measured quantities, i.e. the impact phase $x_1 = \phi - \phi^*$ and the post-impact velocity $x_2 = v - v^*$. The resulting control system is hybrid, since the system input is a continuous-time actuation while the measured outputs are discrete-time quantities. Two controller designs, based on a classical LQR approach and a variant of a “greedy” method respectively, led to good closed-loop performance for control of the period-one orbit inside and *outside* the open-loop stability region (3.17), with simulated data.

The bouncing ball served as a motivating example for further studies on controllability properties and feedback control design of impact systems (see e.g. Tornambe, 1999; Menini and Tornambe, 2003), with several contributions focusing directly on juggling dynamics (Zavala-Rio and Brogliato, 1999; Brogliato and Zavala-Rio, 2000; Lynch and Black, 2001; Zavala-Rio and Brogliato, 2001; Brogliato et al., 2006). These authors provided a general framework for studying the controllability and

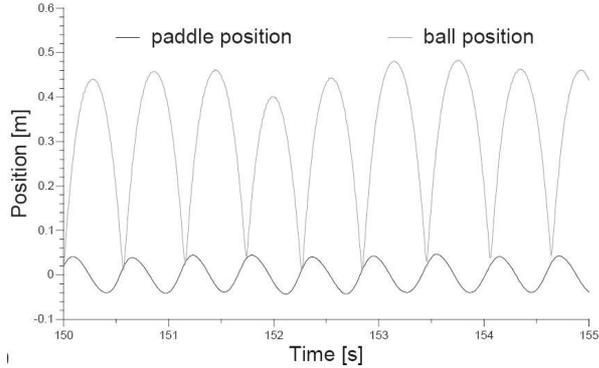


Figure 3.5: Typical strategy of a human subject playing the 1D bouncing ball task with a racket (paddle). Reprinted from Schaal et al. (1996).

stabilization of mechanical systems with impact, and used juggling robots as representative examples.

3.4 Human control of the bouncing ball

The bouncing ball has also motivated several studies in the motor control community. Most of the contributions are due to the seminal work of Sternad, Schaal and coworkers (Schaal et al., 1996; Sternad, 1999; Sternad et al., 2001a,b; Katsumata et al., 2003; de Rugy et al., 2003; Dijkstra et al., 2004; Wei et al., 2006). These authors have investigated human behavior when playing the bouncing ball task with a racket, being asked to stabilize the period-one pattern.

A typical plot of human behavior in this 1D bouncing ball task is depicted in Fig. 3.5. A first observation of this figure clearly reveals that human behavior is much more similar to the sinusoidal actuation (3.7) (see Fig. 3.2) than the mirror law control (3.18) (see Fig. 3.3). Consequently, the central question raised by Sternad, Schaal and their coworkers was to address whether the bouncing ball task was performed by human subjects with or without sensory feedback processing, i.e. in closed- or open-loop. An alternative option is that humans, when performing ball-bouncing, exploit the stability properties of the sinusoidally actuated model (Sternad et al., 2001a).

These authors first observed that the parametric stability region of the period-one motion (3.17) scales either with the movement amplitude or with the square of its frequency. It corresponds moreover to the following range of steady-state acceleration $\ddot{s}^* = -A\omega^2 \sin \phi^*$:

$$\frac{-2(1+e^2)}{(1+e)^2}g < \ddot{s}^* < 0. \quad (3.19)$$

That is, a necessary condition to produce an open-loop stable period-one motion

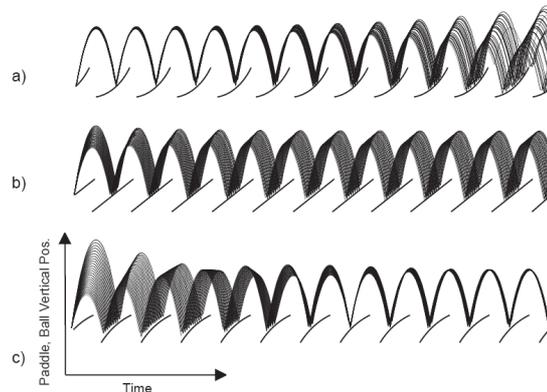


Figure 3.6: Illustration of the role of the acceleration at impact for open-loop stability of the period-one bouncing ball. 25 balls started at the same position but at different initial velocity. The racket trajectory was accelerating (a), at constant velocity (b) or decelerating (c) at impact. Reprinted from Schaal et al. (1996).

is to impact the ball with a racket velocity given by (3.12) and when *decelerating* (i.e. negative acceleration, according to (3.19)) (Schaal et al., 1996; Sternad et al., 2001a,b).

The necessity of negative acceleration at impact for open-loop stability is illustrated in Fig. 3.6. Asymptotic tracking of a given period-one ball motion is only achieved in the third case (c), with negative acceleration at impact. Since the energy restored to the ball depends on the racket velocity at impacts, the acceleration at impact can be interpreted as a *gain* between the puck energy (that is, the puck flight time) and the racket velocity. Negative acceleration at impact provides a negative gain, which is necessary for stability (see Section 4.2.2).

Human subjects played the bouncing ball task with negative acceleration, a strategy which is not intuitive a priori (Schaal et al., 1996; Sternad et al., 2001a,b). On the one hand, this strategy could be guided by the planning system. This requires however a complete assimilation of the task dynamical properties in this trajectory planner, in order to be able to exploit the open-loop stability. Alternatively, the ball and racket dynamics could simply converge into the open-loop stable regime (Schaal et al., 1996). This does not exclude the presence of closed-loop mechanisms in the loop (see below), but the open-loop stable behavior dominates in steady-state.

Sternad et al. (2001a,b) also studied the influence of visual feedback during ball trajectory and haptic feedback at impact. They acquired ball-bouncing data in normal condition, by suppressing the visual feedback, and by suppressing the haptic feedback. This last condition was realized through a mechanical decoupling between the manipulator and the actual racket motion (see Sternad et al., 2001a,b). Fig. 3.7 displays three time series and the phase portraits of the three perceptual conditions.

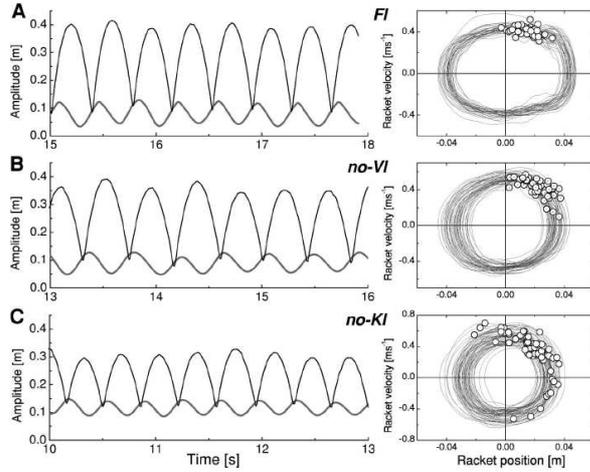


Figure 3.7: Three time series (left panels) and their respective phase portraits (right panels) of exemplary trials performed in the three perceptual conditions: both visual and haptic feedback available (top), no visual feedback (middle) and no haptic feedback (bottom). The dots in the phase portraits denote the impacts. Reprinted from Sternad et al. (2001b).

First, this figure shows that a majority of impacts occurred during the decelerating phase, i.e. the upper right quadrant of the phase portraits. This confirms that the bouncing actions were performed close to an open-loop stable regime in the three conditions. Secondly, without visual feedback (no-VI), the acceleration at impact was just more variable, revealing that visual information might help nevertheless to stabilize the task. More interestingly, when deprived from the haptic perceptual inflow (no-KI), the subjects sometimes abandoned the open-loop stable regime. This is directly visible in Fig. 3.7, since some impacts occurred in a quadrant corresponding to a positive acceleration (upper left). The haptic system may then be relevant for the tuning into open-loop stability, even if it provides only discrete-time inflows, while the visual system provides continuous-time ones. Interestingly, Chapter 4 of this thesis proposes an hypothesis to characterize the valuable role of haptic measurements, as a way to acquire *timing feedback* of the task.

The most recent contributions of Sternad and colleagues in the ball-bouncing task studied the same experiment under large perturbations (de Rugy et al., 2003; Dijkstra et al., 2004; Wei et al., 2006). Human subjects were asked to play the task in a *virtual* environment. Periodically, the racket coefficient of restitution was unexpectedly modified, such that the relaxation behavior, i.e. the interval to recover the period-one motion, was investigated. The main result is that the relaxation time is much shorter than predicted by the open-loop model, undoubtedly revealing that closed-loop mechanisms are implemented by the subjects. Moreover, the authors showed that the *frequency* of the racket movement was modulated after

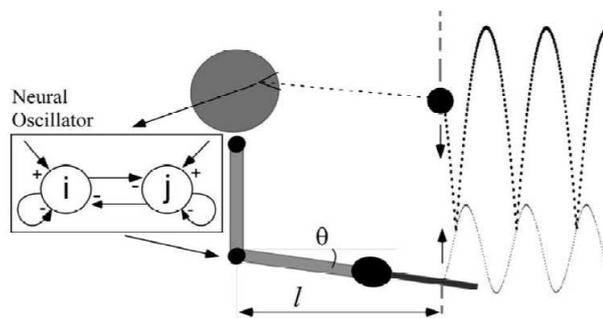


Figure 3.8: Neural-based model of a ball-bouncing controller. Reprinted from de Rugy et al. (2003).

the perturbations, while the amplitude remained roughly constant. The modulation was such that the impacts occurred at a negative acceleration (de Rugy et al., 2003).

de Rugy et al. (2003) also simulated the behavior of a “neural-based” model to control the bouncing ball (see Fig. 3.8). This controller both exploited the open-loop stability of the sinusoidal trajectory, and modulated its frequency according to the “measured” velocity of the ball (closed-loop). This model qualitatively reproduced the human data in order to “track” the open-loop stability with active control of the oscillatory period.

Recently, Tlili et al. (2004) reproduced the negative acceleration criterion for soccer juggling. They reported that human subjects, highly skilled in soccer juggling, impacted the ball with their foot in the decelerating phase. This was observed for a broad range of juggling heights, except for the smallest they tested (0.5m) where the acceleration at impact was just slightly positive. This height corresponded to an average juggling period of 435ms. The active strategy adopted in this configuration remains an open question.

3.5 Conclusion

The bouncing ball dynamics have been widely investigated in the literature. Both the robotics community and the motor control community recognized this very simple task as an illustrative benchmark for studying more complex rhythmic tasks. Indeed, while sharing the main features of regular juggling (Sternad, 1999), such simplified juggling dynamics are amenable to handy mathematical modeling.

Basically two strategies have been developed to control the bouncing ball. The sinusoidal motion of Holmes (1982) is sensorless, while the mirror law of Buehler et al. requires a permanent tracking of the ball to compute its energy and position. However, these two control schemes achieve the same performance, i.e. stabilization of the period-one motion. The mirror law has wider basins of attraction and has been generalized to challenging experimental contexts. Schaal et al. (1996) noticed an-

other important distinction between these two strategies: the sinusoidal law impacts the ball in a *decelerating* upward movement, while the mirror law is always accelerating. The acceleration of the steady-state mirror law is indeed $(1 - e)/(1 + e)g \geq 0$. Through local and non-local stability analysis, they showed that *negative acceleration at impact is a necessary condition for open-loop stability*, in the range defined in (3.19).

Studying human behavior when bouncing a ball with a racket, Sternad, Schaal and coworkers reported that human subjects juggle the 1D bouncing ball with negative acceleration at impact (Schaal et al., 1996; Sternad et al., 2001a,b). Consequently, they concluded that human subjects exploit the open-loop stability properties of the task and do not rely on complex feedback-driven mechanisms, such as the mirror law. Sensory information may nevertheless help to stabilize the task, in wider basin of attraction than predicted by the open-loop model (de Rugy et al., 2003; Dijkstra et al., 2004; Wei et al., 2006).

The scientific background on the bouncing ball dynamics is of high interest for this thesis since our original contributions are based on similar impact tasks dynamics. Chapter 4 aims at reconciling the sensorless and the closed-loop approaches within a hybrid scheme. In Chapters 5, 6, 7 and 8, we switch to a 2D version of the bouncing ball, which is viewed as an idealization of a planar juggler. However, the model-based derivations of the present chapter are still useful since this planar setup is viewed as a 2D extension of the bouncing ball.

Chapter 4

Robust Closed-Loop Control of the Bouncing Ball

If the facts don't fit the theory,
change the facts.

Albert Einstein

4.1 Introduction

As illustrated in Chapter 3, open-loop stabilization of the bouncing ball requires negative acceleration of the actuator at impact (eq. (3.19), see Schaal et al., 1996). In contrast, the mirror law algorithms have been designed in closed-loop and can consequently afford positive acceleration at impact. The aim of this chapter is to discuss the role of acceleration tuning for *robust* closed-loop design of the bouncing ball dynamics with *limited* feedback, i.e. when the complete ball trajectory cannot be sensed as in the mirror law implementation. We focus on the measurement of the *impact times* as sole discrete output. This information provides a source of “rhythmic” feedback, directly available through auditive or haptic inflows for human subjects. Such a design is also cheap and efficient in robotics, since the times of impact can be measured by cheap sensors, like accelerometers (recording the vibrations on the actuator), microphones, etc. . .

In human juggling experiments, it has long been recognized that the control of timing is ubiquitous, either for throwing and catching time in the three-balls cascade (Amazeen et al., 1999) or for 1D ball bouncing (see Section 3.4 and Sternad et al., 2001b). This has been further emphasized recently in a robotic experiment by Hirai and Miyazaki (2006). They studied a juggling-like ball-passing task in the horizontal plane (i.e. no effect of gravity) whose stabilization is based only on feedback measurement of the impact times between the balls and the robot. The objective of the present chapter is to emphasize the role of impact times as a central information

for the *model-based* feedback control of the bouncing ball and to demonstrate the existence of *robust* control schemes based on this sole discrete information. The measurement of the occurrences of timed events, like the impacts in this task, could obviously be generalized as an exclusive source of feedback for most rhythmic systems. For example, in locomotion, the initiation and/or termination of the stance phase are marked periodic events in the cycle that could be useful for feedback.

The 1D bouncing ball dynamics (Holmes, 1982; Guckenheimer and Holmes, 1986) have been reviewed in Section 3.2 of this thesis. The actuation is the continuous-time motion $s(t)$. The ball dynamics are governed by the gravitational field g and the impacts with the actuator. The complete bouncing ball dynamics are described by the discrete *Poincaré map* (3.4) (flight map) and (3.5) (impact rule), whose state is the impact position $s[k]$ and post-velocity $v[k]$. Given a reference trajectory $(s_\rho[\bullet], v_\rho[\bullet])$ for the system (3.4), (3.5), we study the design of a *continuous-time* motion $s(t)$ that achieves asymptotic tracking of this *discrete* reference.

We first derive a controller that achieves deadbeat tracking of an *arbitrary* reference trajectory (Sections 4.2.3 and 4.2.4). This controller is a *state* feedback controller, which is then turned into an *output* feedback controller (Section 4.2.5): the state of the ball is reconstructed from the sole impact times by means of a deadbeat observer. The output feedback controller assigns the impact position and the impact velocity of the actuator but the stability of the closed-loop system puts no constraints on the impact *acceleration*. In contrast, this parameter turns to be a crucial design parameter in the *robustness* analysis of the closed-loop system (Section 4.3). We focus the robustness analysis on the model uncertainty arising from the impact model. The design is based on the simple Newton's impact model and we model the uncertainty by treating the variations of the coefficient of restitution as an external disturbance. By a proper tuning of the acceleration at impact, simulation results illustrate that the uncertainty on the coefficient of restitution can be efficiently rejected.

4.2 Feedback Control of the Bouncing Ball

4.2.1 The sinusoidal actuation vs. the mirror law

In Chapter 3, we introduced two actuation laws to stabilize the period-one motion of the bouncing ball. First, the sinusoidal actuation (3.7) is purely open-loop and stabilizes the period-one motion if the amplitude and frequency are tuned according to (3.17) (Holmes, 1982; Guckenheimer and Holmes, 1986; Bapat et al., 1986). The stability of the orbit implies that the actuator acceleration at impact is *negative* (Schaal et al., 1996) in the range given by (3.19). This is illustrated in Fig. 4.1, top.

In contrast with the sensorless sinusoidal law, the *mirror law* is a feedback strategy based upon permanent tracking of the ball (Buehler, Koditschek and Kindlmann, 1988, 1990, 1994). Its simplest version has been derived in Section 3.3. In steady-state, the mirror law behavior is reproduced in Fig. 4.1, bottom.

The two control strategies that stabilize the same pattern are clearly distinct in term of feedback requirement. The first one is sensorless while the second requires a permanent tracking of the ball. Comparing the sinusoidal law and the mirror law in Fig. 4.1, we see that both of them stabilize the bouncing ball period-one at the same steady-state (s^*, v^*) , but with significant differences in the underlying continuous-time control law. We observe that the sinusoidal trajectory is decelerating at impact while the mirror law is accelerating. It suggests that the acceleration at impact possibly influences the feedback requirements of the control law.

4.2.2 Stabilization with sensorless sinusoidal actuation

Even if there is no control loop with the sensorless sinusoidal actuation $s(t) = A \sin(\omega t)$, it is of interest to interpret its stabilizing feedback mechanism in the system dynamics.

The linearized state-space equations (3.14) and (3.15) can be written under the following matrix form:

$$\begin{aligned} \begin{pmatrix} \frac{\omega}{g} \delta v[k+1] \\ \omega \delta t[k+1] \end{pmatrix} &= \begin{pmatrix} e^2 & 0 \\ 1 & 1+e \end{pmatrix} \begin{pmatrix} \frac{\omega}{g} \delta v[k] \\ \omega \delta t[k] \end{pmatrix} \\ &\quad + \begin{pmatrix} -(1+e) \frac{A\omega^2}{g} \sqrt{1 - \left(\frac{(1-e)\pi n g}{(1+e)A\omega^2} \right)^2} \\ 0 \end{pmatrix} \omega \delta t[k+1] \\ &= \left(\underbrace{\begin{pmatrix} e^2 & 0 \\ 1 & 1+e \end{pmatrix}}_{\tilde{\mathbf{A}}_{BB}} + \underbrace{\begin{pmatrix} -(1+e) \frac{A\omega^2}{g} \sqrt{1 - \left(\frac{(1-e)\pi n g}{(1+e)A\omega^2} \right)^2} \\ 0 \end{pmatrix}}_{\tilde{\mathbf{B}}_{BB}} \right) \underbrace{\begin{pmatrix} 1 & 1+e \end{pmatrix}}_{\tilde{\mathbf{C}}_{BB}} \begin{pmatrix} \frac{\omega}{g} \delta v[k] \\ \omega \delta t[k] \end{pmatrix} \end{aligned} \quad (4.1)$$

which is obviously equivalent to (3.16), with $\mathbf{A}_{BB} = \tilde{\mathbf{A}}_{BB} + \tilde{\mathbf{B}}_{BB} \tilde{\mathbf{C}}_{BB}$.

Eq. (4.1) admits the following state-space representation:

$$\mathbf{x}[k+1] = \tilde{\mathbf{A}}_{BB} \mathbf{x}[k] + \tilde{\mathbf{B}}_{BB}(A, \omega) u[k] \quad (4.2)$$

$$\begin{aligned} y[k] &= t[k+1] \\ &= \tilde{\mathbf{C}}_{BB} \mathbf{x}[k] \end{aligned} \quad (4.3)$$

where the ‘‘controller input’’ equals the ‘‘output’’: $u[k] = y[k]$, i.e. the impact time $t[k+1]$. This state-space representation is both controllable and observable.

In this representation, the sinusoidal actuation is interpreted as a proportional feedback of the output $t[k+1]$, emphasizing the importance of estimating the *next impact time* for stabilization. The first element of the matrix $\tilde{\mathbf{B}}_{BB}$ is a feedback gain equal to the actuator acceleration at impact. Negative feedback thus requires a negative acceleration.

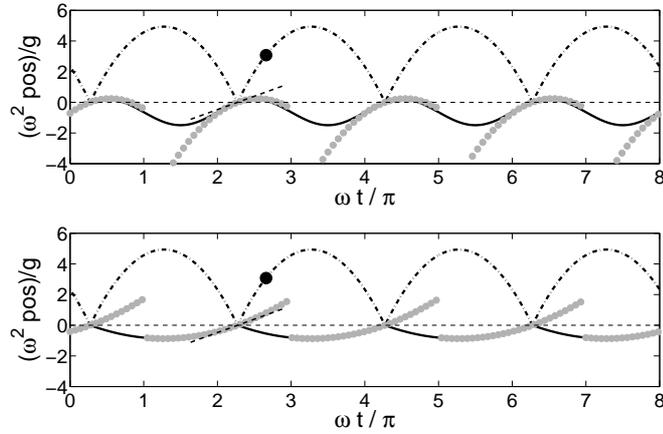


Figure 4.1: Comparison of the sinusoidal trajectory (top) and the mirror law (bottom) to stabilize the period-one pattern. The actuator (respectively the ball) non-dimensional position is depicted with solid (respectively dash-dotted) lines over non-dimensional time. The gray dots denote the piecewise quadratic trajectory that match the actuator position, velocity and acceleration at impact, see Section 4.2.3.

4.2.3 State feedback control

This section describes a more general class of controllers aiming at tracking a *time-varying* referenced trajectory, impact after impact. Given the flight time equation (3.6), exact matching between the real and the referenced next impact position, i.e.

$$s[k+1] = s_\rho[k+1] \quad (4.4)$$

is provided if the next impact occurs at time

$$t_u[k+1] = t[k] + \frac{v[k] + \sqrt{v[k]^2 - 2g(s_\rho[k+1] - s[k])}}{g}. \quad (4.5)$$

Given (3.5) and (4.5), the impactor velocity at impact must be equal to

$$\dot{s}_u[k+1] = \frac{v_\rho[k+1] - e\sqrt{v[k]^2 - 2g(s_\rho[k+1] - s[k])}}{1+e} \quad (4.6)$$

in order to provide $v[k+1] = v_\rho[k+1]$.

As illustrated in Fig. 4.1, the two control laws described in Section 4.2.1 achieve the same stabilization objective, but differ in the impact acceleration. In order to illustrate how the actuator acceleration at impact influences the robustness of the feedback system, we consider at first a mathematically convenient family of control laws in the form of a piecewise quadratic function of time, re-initialized after each impact:

$$s(t) = s_u[k+1] + \dot{s}_u[k+1](t - t_u[k+1]) + \frac{\gamma}{2}(t - t_u[k+1])^2 \quad (4.7)$$

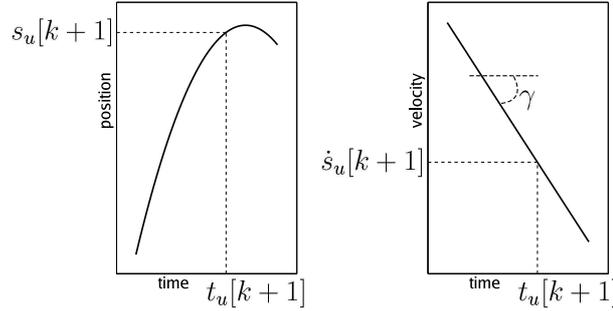


Figure 4.2: Local position (left) and velocity (right) profiles of the actuator trajectory. The position profile is given by (4.7) while the velocity is its first derivative. γ is the acceleration.

for $t[k] < t \leq t[k+1]$, γ denoting the actuator *acceleration*. At time $t = t_u[k+1]$, the actuator position (velocity, resp.) is equal to $s_u[k+1]$ ($\dot{s}_u[k+1]$, resp.), see Fig. 4.2. The gray dots on Fig. 4.1 illustrate the steady-state behavior of this control law when the parameters are tuned to match the ball position, velocity and acceleration at impact with the two control laws presented in Section 4.2.1.

The quadratic parameterization of the control (4.7) is convenient to obtain an explicit expression of the impact times: $t[k+1]$ is indeed the solution of (3.4) and (4.7) at time $t = t[k+1]$, which defines a second order polynomial in t . The controller inputs are $t_u[k+1]$ (4.5), $s_u[k+1] = s_\rho[k+1]$ and $\dot{s}_u[k+1]$ (4.6) and are functions of the state $(s[k], v[k])$ and impact time $t[k]$. Deadbeat convergence of this tracking control law, that is, convergence of the solution to the reference trajectory after a finite number of time steps, is established in the next section.

4.2.4 Deadbeat convergence

Substituting (4.4), (4.5) and (4.6) into (4.7) at time $t = t[k+1]$, we find:

$$\begin{aligned}
 s[k+1] &= s_\rho[k+1] \\
 &+ \frac{v_\rho[k+1] - e\sqrt{v[k]^2 - 2g(s_\rho[k+1] - s[k])}}{1+e} \left(t[k+1] - t[k] - \frac{v[k] + \sqrt{v[k]^2 - 2g(s_\rho[k+1] - s[k])}}{g} \right) \\
 &+ \frac{\gamma}{2} \left(t[k+1] - t[k] - \frac{v[k] + \sqrt{v[k]^2 - 2g(s_\rho[k+1] - s[k])}}{g} \right)^2. \tag{4.8}
 \end{aligned}$$

Substituting (3.6) into (4.8), we obtain a second order polynomial in $t[k+1]$, whose positive root is:

$$t[k+1] = t[k] + \frac{v[k] + \sqrt{v[k]^2 - 2g(s_\rho[k+1] - s[k])}}{g}. \tag{4.9}$$

Comparing this with (3.6) yields:

$$s[k+1] = s_\rho[k+1] \quad (4.10)$$

and therefore also:

$$v[k+1] = v_\rho[k+1] \quad (4.11)$$

$$t[k+1] = t_u[k+1] \quad (4.12)$$

reflecting then that the positive solution of (3.4), (3.5) and (4.7) is the exact matching between the reference and the real impact state.

To summarize, the continuous-time control law determined by the quadratic expression (4.7) and the discrete control (4.4), (4.5) and (4.6), ensures deadbeat convergence of the impact state after one time step. So far, the acceleration γ in (4.7) is a *free* parameter and does not influence the convergence. The control law is a *tracking* controller, that is, the reference trajectory $(s_\rho[\bullet], v_\rho[\bullet])$ is arbitrary, as long as it corresponds to a solution of the dynamical system (3.4) and (3.5).

4.2.5 Output feedback deadbeat control

The piecewise quadratic control proposed in Section 4.2.3 uses the full state $(s[k], v[k])$ of the system (3.4) and (3.5). To reduce the sensing requirements of the controller, we now assume that only the continuous-time actuator motion $s(t)$ is measured, together with the *impact times* $t[k]$: they provide the impact position $s[k]$ and velocity $\dot{s}[k]$. In this section, we derive a deadbeat *observer* that reconstructs the post-impact ball velocity $v[k]$ from the measured impact times $t[k]$ and actuator motion $s(t)$.

Post-impact velocity is estimated by an observer that is a copy of the velocity dynamics (3.5):

$$\hat{v}[k] = -e v[k-1] + e g (t[k] - t[k-1]) + (1+e)\dot{s}[k] \quad (4.13)$$

while $v[k-1]$ is obtained from (3.4):

$$\begin{aligned} \hat{v}[k] &= -e \left(\frac{s[k]-s[k-1]}{t[k]-t[k-1]} + \frac{g}{2}(t[k] - t[k-1]) \right) + e g (t[k] - t[k-1]) + (1+e)\dot{s}[k] \\ &= e \frac{g}{2} (t[k] - t[k-1]) - e \frac{s[k] - s[k-1]}{t[k] - t[k-1]} + (1+e)\dot{s}[k]. \end{aligned} \quad (4.14)$$

Equation (4.14) defines a deadbeat velocity observer using the *impact times* as sole input in addition to the actuator motion. Deadbeat convergence is ensured in one time-step since $\hat{v}[k] = v[k], \forall k > 1$.

The output feedback controller, whose the only measured signals are the *impact times*, is then obtained by replacing the actual state variable $v[k]$ by the estimated variable $\hat{v}[k]$ in (4.4), (4.5) and (4.6). Its deadbeat convergence is established in the following proposition:

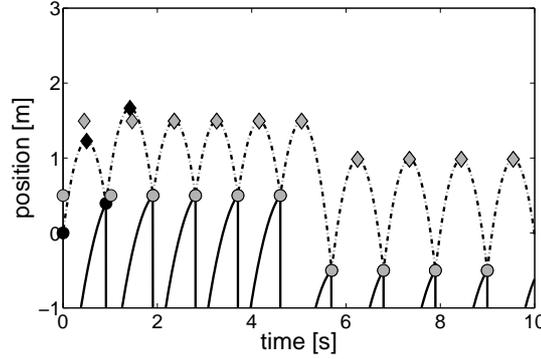


Figure 4.3: Output feedback control of the bouncing ball with the control law (4.15). The actuator (respectively the ball) position is depicted with solid (respectively dash-dotted) lines over time. Actual impacts position (apex position, resp.) are represented with black circles (black diamonds, resp.). Reference positions are accordingly represented with gray markers. $e = 0.7$, $\gamma < 0$.

Proposition 1 (Deadbeat convergence of the piecewise quadratic output feedback controller) *Consider the bouncing ball dynamics (3.4) and (3.5) and a reference trajectory $(s_\rho[k], v_\rho[k])$, $k \geq 0$. The output feedback control*

$$\begin{aligned}
 s(t) &= s_\rho[k+1] + \dot{s}_u[k+1](t - t_u[k+1]) + \frac{\gamma}{2}(t - t_u[k+1])^2 \\
 &\quad \text{for } t[k] < t \leq t[k+1] \\
 t_u[k+1] &= t[k] + \frac{\hat{v}[k] + \sqrt{\hat{v}[k]^2 - 2g(s_\rho[k+1] - s[k])}}{g} \\
 \dot{s}_u[k+1] &= \frac{v_\rho[k+1] - e\sqrt{\hat{v}[k]^2 - 2g(s_\rho[k+1] - s[k])}}{1+e} \\
 \hat{v}[k] &= e\frac{g}{2}(t[k] - t[k-1]) - e\frac{s[k] - s[k-1]}{t[k] - t[k-1]} + (1+e)\dot{s}[k] \quad (4.15)
 \end{aligned}$$

ensures deadbeat convergence of the 1D bouncing ball state toward the reference $(s_\rho[k], v_\rho[k])$ after two impacts (i.e. $k > 2$).

Proof The first impact is required to ensure convergence of the observer, since $\hat{v}[k] = v[k]$, $\forall k > 1$. As soon as $\hat{v}[k] = v[k]$, a second impact is required to achieve deadbeat convergence of the controller, as shown in Section 4.2.4. \square

The deadbeat convergence is illustrated in Fig. 4.3. At the third impact, both the reference impact position and the reference apex (which is an image of the reference post-impact velocity) are reached. The same figure illustrates the proper tracking of time-varying references since both the position and velocity references change at the seventh impact.

4.2.6 A blind mirror law based on output feedback

Assuming a ballistic flight between two impacts, the mirror law (3.18) can also be adapted to an *output* feedback control, with the impact times as measured output:

$$s(t) = \left(\frac{-(1-e)}{1+e} - \kappa_1 \left(E_\rho^* - \hat{E}(t) \right) \right) \hat{\beta}(t) \quad (4.16)$$

since both the ball position and energy can be estimated from impact state:

$$\begin{aligned} \hat{\beta}(t) &= s[k] + \hat{v}[k](t - t[k]) - \frac{g}{2}(t - t[k])^2 \\ \hat{E}(t) &= \hat{E}[k] \\ &= gs[k] + \frac{1}{2}\hat{v}[k]^2 \end{aligned}$$

for $t[k] \leq t < t[k+1]$. The tracking mirror law (3.18) requires a permanent tracking of the ball as sensory input. In contrast, the blind mirror law (4.16) based on output feedback only uses impact times and reconstructs the post-impact velocity $v[k]$ via the observer described in Section 4.2.5.

Since both the piecewise quadratic law (4.15) and the blind mirror law (4.16) require the same sensing capabilities, their robustness will be compared in the next section.

4.3 Robustness to Model Uncertainties

The *acceleration* parameter γ appearing in (4.15) played no role in the stability and convergence analysis. This section stresses the importance of this parameter for robustness purposes. We show that particular negative accelerations efficiently optimize either static or dynamical perturbations induced by a poor estimate of the coefficient of restitution e .

4.3.1 Uncertainty of the impact model

Among the several sources of uncertainty of the model (3.4) and (3.5), the *impact* model (3.3) is probably central. While the Newton impact law models e as a constant, this parameter is varying in experimental conditions. Furthermore, the linear relationship between the pre- and post-impact velocities (3.3) is certainly not respected outside a narrow range of impact velocities.

We model the uncertainty on the coefficient of restitution e by considering the following perturbed impact rule:

$$v[k+1] - \dot{s}[k+1] = -(e + \Delta e[k+1])(v^-[k+1] - \dot{s}[k+1]) \quad (4.17)$$

where $\Delta e[k+1]$ models the variation of the coefficient of restitution at time $t[k+1]$, w.r.t. the estimated value e . For the sake of simplicity, we study the robustness of the linearized feedback system, as derived in Appendix 4.A of this chapter.

From the linear state-space representation (4.28), we find the following closed-loop input-to-state transfer functions:

$$S(z) = S_\rho(z) + \frac{2(1-e)}{1+e} \frac{1}{z} E(z) \quad (4.18)$$

$$V(z) = V_\rho(z) + \frac{2}{1+e} \frac{z + \left(e^2 + \frac{\gamma}{g}(1+e)^2\right)}{z} E(z) \quad (4.19)$$

where $S(z)$, $V(z)$, $S_\rho(z)$, $V_\rho(z)$ and $E(z)$ refer to the z -transforms of $g\delta s[k]/(v_\rho^*)^2$, $\delta v[k]/v_\rho^*$, $g\delta s_\rho[k]/(v_\rho^*)^2$, $\delta v_\rho[k]/v_\rho^*$ and $\Delta e[k]$, respectively. In (4.18) and (4.19), the absence of dynamics in the transfer from references to states is due to the dead-beat convergence established in Section 4.2.5. The next section details the role of acceleration to reject the perturbations due to $\Delta e[k]$.

4.3.2 Robustness requires negative acceleration

From (4.19), we see that the acceleration γ can be designed to *place the zero* of the transfer function from $E(z)$ to $V(z)$. That design parameter will be discussed to optimize either the static or the dynamic performance.

Static error

To let the post-impact velocity converge toward the reference $v_\rho[k]$, assuming a *constant* perturbation $\Delta e[k] = \Delta e$, one has to cancel the static gain of the transfer function from $E(z)$ to $V(z)$. This amounts to place the zero of (4.19) at $z = -1$, which requires the following acceleration:

$$\gamma_{\text{stat.}} = -\frac{1+e^2}{(1+e)^2} g. \quad (4.20)$$

Interestingly, this optimal acceleration depends only on e , that is, the *estimated* coefficient of restitution. Fig. 4.4 illustrates the behavior of the feedback system when the coefficient of restitution is estimated at $e = 0.7$ while the real one is only $e + \Delta e = 0.5$. The desired post-impact velocity is reached because the difference between impact and apex positions is the same for the reference as for the actual trajectory. It should be noticed, however, that a static error persists on the reference position: the static gain from $E(z)$ to $S(z)$ is indeed independent of γ (see (4.18)). This static error does not appear to be detrimental to the robustness of the feedback system.

The optimal acceleration for static performance, as identified in (4.20), has also been derived from the original nonlinear equations, see Ronsse and Sepulchre (2006).

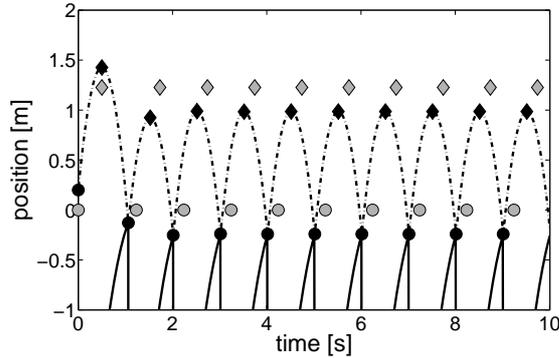


Figure 4.4: Output feedback control of the bouncing ball with piecewise quadratic trajectory. The actuator (ball, resp.) position is depicted with solid (dash-dotted, resp.) lines over time. Actual impacts position (apex position, resp.) are denoted with black circles (black diamonds, resp.). Reference positions are accordingly denoted with gray markers. $e + \Delta e = 0.5$, $e = 0.7$, γ is given by (4.20).

It is of interest to relate this particular acceleration to the sinusoidal control discussed in Section 4.2.1: the optimal acceleration (4.20) is exactly the middle point of the acceleration range where the period-one motion is stable (3.19).

Dynamic performance

Robustness to a static error on e is not the primary issue in real experiments because the average value of e is easy to determine. In contrast, robustness is required against the sustained variability of e . By placing the zero of (4.19) at $z = 0$, the dynamics from $E(z)$ to $V(z)$ are exactly canceled :

$$\gamma_{\text{dyn.}} = -\frac{e^2}{(1+e)^2}g \quad (4.21)$$

resulting in a static transfer function. Interestingly, this optimal value closely matches the value that minimizes the numerically computed variability of the sinusoidally actuated bouncing ball in Sternad et al. (2001a,b).

We summarize the robustness analysis of the piecewise quadratic law (4.15) with the following proposition:

Proposition 2 *Consider the bouncing ball dynamics (3.4) and (3.5) with a time-varying coefficient of restitution $e + \Delta e[k + 1]$. Using the output feedback control (4.15), the transfer function from $\Delta e[k]$ to $\delta v[k]/v_p^*$ is given by*

$$\frac{V(z)}{E(z)} = \frac{2}{1+e} \frac{z + \left(e^2 + \frac{\gamma}{g}(1+e)^2 \right)}{z}. \quad (4.22)$$

The choice $\gamma_{\text{stat.}} = -(1 + e^2)/(1 + e)^2 g$ (see (4.20)) ensures zero steady state error while the choice $\gamma_{\text{dyn.}} = -e^2/(1 + e)^2 g$ (see (4.21)) cancels the transfer function dynamics. Both choices result in negative acceleration at impact, with

$$\gamma_{\text{stat.}} < \gamma_{\text{dyn.}} \leq 0. \quad (4.23)$$

Simulation results

To illustrate the role of the impact acceleration for robustness in the nonlinear dynamics, we now compare the output piecewise quadratic controller (with the optimal negative accelerations previously identified) with the output controller mirror law (4.16). We have tested that $\kappa_1 \approx 0.025$ achieves the best trade-off between performance (rate of convergence) and robustness (noise sensitivity) in that blind mirror law.

Both the parabolic flight assumption and the Newton impact rule are perturbed by noise in an experimental setup. We simulated these perturbations by adding some noise to the reconstructed velocity $\hat{v}[k]$, whose dynamics use both the flight map and the impact rule. Equation (4.14) is thus replaced by:

$$\hat{v}[k] = \left(e \frac{g}{2} (t[k] - t[k-1]) - e \frac{s[k] - s[k-1]}{t[k] - t[k-1]} + (1 + e) \dot{s}[k] \right) (1 + \nu[k] \epsilon_{nl}) \quad (4.24)$$

where $\nu[k]$ is a random number between -1 and 1 , and ϵ_{nl} is the noise level. For the sake of illustration, we focus on stabilization of a period-one motion, characterized by $(s_\rho^*, v_\rho^*) = (0, g/2)$, i.e. one impact per second.

Fig. 4.5 depicts the standard deviation of the normalized post-impact velocity $v[k]/v^*$ over 100 impacts, for increasing noise level. The standard deviation of the impact *position* does not vary significantly depending on the control law, as suggested by (4.18). However, the piecewise quadratic law with the acceleration tuned to cancel the dynamics in (4.22) (γ defined by (4.21)) achieves quasi-zero variance in *post-impact velocity*, see the dashed line. For the tested noise levels, the standard deviation is intermediate with $\gamma = 0$. In contrast, both the blind mirror law and the piecewise quadratic law with γ defined by (4.20) generate twice as much variability. The excellent noise rejection obtained with $\gamma_{\text{dyn.}}$ (4.21), even considering the nonlinear dynamics, illustrates the robustness of the piecewise quadratic control law with a suitable negative acceleration. The range of negative acceleration that produces good noise rejection is obviously limited, since the more negative value $\gamma_{\text{stat.}}$ already results in poor dynamical performance.

4.3.3 Integral feedback control

The control objective considered in the previous sections was to track the position and velocity references (s_ρ^*, v_ρ^*) . However, as illustrated in Section 4.3.2, it is tedious to exactly cancel the static error of the post-impact velocity in the presence of model

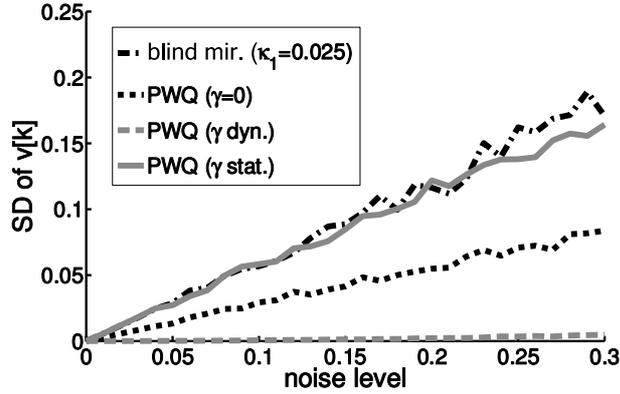


Figure 4.5: Simulations of the nonlinear noisy bouncing ball dynamics with different control laws: the blind mirror law (4.16) ($\kappa_1 = 0.025$, black dash-dotted), the zero-acceleration piecewise quadratic law (4.15) ($\gamma = 0$, black dotted), the piecewise quadratic law (4.15) with the optimal static acceleration (4.20) (gray plain) and the piecewise quadratic law (4.15) with the optimal dynamic acceleration (4.21) (gray dashed). The figure represents the standard deviation of the normalized post-impact velocity $v[k]/v^*$, calculated over 100 impacts. Noise level is defined in (4.24).

uncertainties. Given (3.8), the post-impact velocity static error will result in *flight time* static error (i.e. pattern frequency). This static error causes a linearly growing phase shift between the reference and the actual *impact times*, with detrimental consequences for pattern stabilization.

The velocity static error can be eliminated by integral feedback: the reference trajectory is now the position $s_\rho[\bullet]$ and the *impact time* $t_\rho[\bullet]$ (i.e. impact phase). The corresponding velocity input $v_\rho[k+1]$ in $\dot{s}_u[k+1]$ (4.15) is computed as the solution of (3.6):

$$v_\rho[k+1] = \frac{s_\rho[k+2] - s_\rho[k+1]}{t_\rho[k+2] - t_u[k+1]} + \frac{g}{2}(t_\rho[k+2] - t_u[k+1]). \quad (4.25)$$

The difference between the *reference* impact time $t_\rho[k+2]$ and the *estimated* impact time $t_u[k+1]$ must be then constant in steady-state:

$$(t_\rho[k+1] - t_u[k])^* = \frac{2v_\rho^*}{g}. \quad (4.26)$$

Since the difference between the estimated and the actual impact times is also a constant in steady-state, the delay between the reference and the actual impact times will remain constant through impacts, resulting in no static error between the desired and actual post-impact velocity, and therefore in the pattern frequency.

Deadbeat convergence in three impacts of the integral controller is straightforwardly established by adapting the derivations of Section 4.2.3. Note that this

controller requires reference signals *two* steps ahead since both $s_\rho[k+2]$ and $t_\rho[k+2]$ are used in (4.25).

4.4 Conclusion

This chapter presented the design, analysis and simulations of robust closed-loop control of the period-one motion in the bouncing ball. The control law only uses the impact times as feedback information, relaxing thereby the need for complex sensor design. The parameter γ , i.e. the acceleration at impact, plays no role in the stability analysis; in contrast, proper tuning of this parameter was shown to have a dramatic effect on robustness. The design is based on the simple Newton's impact model and we model the uncertainty by treating the variations of the coefficient of restitution as an external disturbance. Analyzing the transfer function from this disturbance to the impact velocity, we showed that the dynamics can be exactly canceled by a proper choice of the impact acceleration. This particular tuning requires *negative* impact acceleration, in accordance with sensorless control strategies (Holmes, 1982; Schaal et al., 1996; Ronsse et al., 2006) and with observed human strategies (Schaal et al., 1996; Sternad et al., 2001a,b). This contrasts with the mirror law algorithm proposed earlier in the literature (Buehler et al., 1988, 1990, 1994) and possibly explains why such control schemes — that have been shown to perform robustly in 1D, 2D and even 3D environments with a continuous-time sensing of the juggled object — may perform poorly in implementations with limited sensing capabilities (Gerard, 2005).

The chapter illustrated that measurement of impact times is both a cheap and relevant feedback source in juggling experiments. It may therefore supplement the continuous-time sensing required in more complicated juggling implementations, e.g. in 3D environments. The piecewise quadratic control introduced in the present chapter will be generalized to a smoother closed-loop trajectory in Chapter 7. This will permit to validate this strategy on a real experimental setup.

The results of the present chapter have been published in Ronsse et al. (2007a), Sections II to IV.

4.A Linearized Equations of the 1D Noisy Bouncing Ball with Piecewise Quadratic Control

The linearized dynamics of the perturbed 1D bouncing ball dynamics (3.4) and (4.17), and the piecewise quadratic law (4.15) are given by the following set of equations:

$$\begin{aligned}
\delta s[k+1] &= \delta s[k] + \frac{2\pi n}{\omega} \delta v[k] - \frac{\pi n g}{\omega} (\delta t[k+1] - \delta t[k]), \\
\delta v[k+1] &= -e \delta v[k] + e g (\delta t[k+1] - \delta t[k]) + (1+e) \delta \dot{s}[k+1] \\
&\quad + \frac{2}{1+e} \frac{\pi n g}{\omega} \delta e[k+1], \\
\delta s_\rho[k+1] &= \frac{2e}{1+e} \delta s_\rho[k+1] + \frac{1-e}{1+e} \delta s[k] \\
&\quad + \frac{1-e}{1+e} \frac{\pi n g}{\omega} (\delta t[k+1] - \delta t[k]) - \frac{1-e}{1+e} \frac{2\pi n}{\omega} \delta \hat{v}[k], \\
\delta \dot{s}[k+1] &= \frac{1}{1+e} \delta v_\rho[k+1] - \left(\frac{e}{1+e} + \frac{2\gamma}{g} \right) \delta \hat{v}[k] \\
&\quad + \left(\frac{e}{1+e} + \frac{\gamma}{g} \right) \frac{\omega}{\pi n} (\delta s_\rho[k+1] - \delta s[k]) + \gamma (\delta t[k+1] - \delta t[k]), \\
\delta \hat{v}[k+1] &= e \frac{g}{2} (\delta t[k+1] - \delta t[k]) - e \frac{\omega}{2\pi n} (\delta s[k+1] - \delta s[k]) + (1+e) \delta \dot{s}[k+1].
\end{aligned} \tag{4.27}$$

In (4.27), $\delta e[k+1]$ is the small perturbation on the coefficient of restitution and is considered as an additional input.

Using non-dimensional state variables, one obtains the following state-space model:

$$\begin{aligned}
\begin{pmatrix} \frac{g\delta s[k+1]}{(v_\rho^*)^2} \\ \frac{\delta v[k+1]}{v_\rho^*} \\ \frac{g\delta t[k+1]}{v_\rho^*} \\ \frac{\delta \hat{v}[k+1]}{v_\rho^*} \end{pmatrix} &= \underbrace{\begin{pmatrix} 0 & 1-e & 0 & -(1-e) \\ 0 & e^2 + \frac{\gamma}{g}(1+e)^2 & 0 & -\left(e^2 + \frac{\gamma}{g}(1+e)^2\right) \\ 1 & 1+e & 1 & 1-e \\ 0 & e^2 + \frac{\gamma}{g}(1+e)^2 & 0 & -\left(e^2 + \frac{\gamma}{g}(1+e)^2\right) \end{pmatrix}}_{\mathbf{A}'} \begin{pmatrix} \frac{g\delta s[k]}{(v_\rho^*)^2} \\ \frac{\delta v[k]}{v_\rho^*} \\ \frac{g\delta t[k]}{v_\rho^*} \\ \frac{\delta \hat{v}[k]}{v_\rho^*} \end{pmatrix} \\
&\quad + \underbrace{\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & \frac{2}{1+e} \\ -1 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}}_{\mathbf{B}'} \begin{pmatrix} \frac{g\delta s_\rho[k+1]}{(v_\rho^*)^2} \\ \frac{\delta v_\rho[k+1]}{v_\rho^*} \\ \delta e[k+1] \end{pmatrix}. \tag{4.28}
\end{aligned}$$

The state variables are small perturbations of the ball impact position $g\delta s[k+1]/(v_\rho^*)^2$ and velocity $\delta v[k+1]/v_\rho^*$; the impact time $g\delta t[k+1]/v_\rho^*$ and the observed velocity $\delta \hat{v}[k+1]/v_\rho^*$.

The matrix \mathbf{A}' is singular. This is a consequence of deadbeat convergence of the 1D bouncing ball, controlled with the piecewise quadratic law (4.15).

Chapter 5

Experimental Setup: the Wiper Robot

To be intelligent is to be
open-minded, active, memoried,
and persistently experimental.

Leopold Stein

5.1 Introduction

At the center of the thesis, this chapter describes the experimental setup we designed to conduct robotics and motor control experiments on impact juggling. One major contribution of the present work was to design this experimental setup, amenable to test both the theoretical predictions in robotics and to explore the human behavior. This setup is based on a simplified juggling paradigm, both capturing the main features of regular juggling, and amenable to simple mathematical modeling. This setup is described in Section 5.2, such as a crude model which is shown to be reducible to the bouncing ball model described in Chapter 3.

More technical details on the setup are given in Appendix B.

5.2 Wiper: an experimental setup amenable to mathematical modeling

One of the most popular juggling patterns is called the shower, and is depicted in Fig. 5.1(a): the balls follow a cyclic trajectory along two distinct parabolas produced by a low and a high toss. We study this steady-state pattern in an experiment that drastically simplifies the hardware: a *planar* motion of the puck between two impacting edges that idealize the juggler's arms. Fig. 5.1(b) depicts the so-called

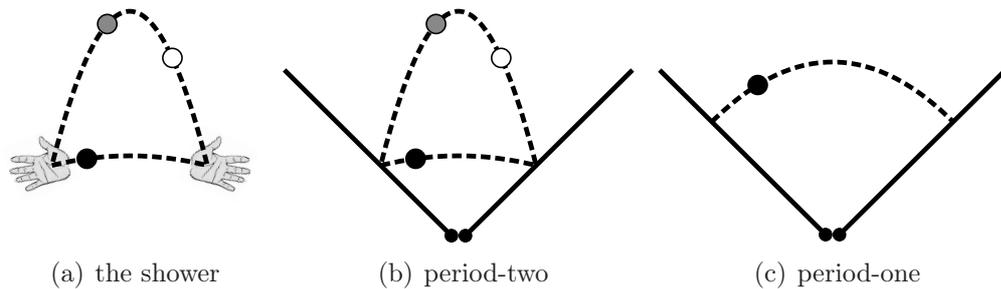


Figure 5.1: The *shower* pattern (Fig. 5.1(a)) is one of the simplest juggling pattern. It corresponds to a limit cycle of Wiper that is called the *period-two* (Fig. 5.1(b)), since the balls cycle between the edges along two parabolas. A degenerate (and simpler) case of the period-two, where both parabolas are similar, is called the *period-one* (Fig. 5.1(c)).

period-two orbit of this impact juggler, and the particular period-one orbit (both tosses are equal) is depicted in Fig. 5.1(c).

Laboratory implementation of this impact juggling experiment has been realized on the *Wiper* robot, pictured in Fig. 5.2. The motion plane is a tilted air-hockey table. Air-hockey is a popular game which is based on tight goal-shots of plastic pucks on an horizontal table. The puck trajectories are almost frictionless since the table is pierced with a lattice of little holes blowing air constantly. This frictionless table has been tilted with respect to the ground, such that gravity influences the puck motion, like in regular juggling. The gravity field \mathbf{g} can be adjusted by proper inclination of the table. The two metallic “arms” have a single (rotational) degree of freedom. The pucks have been manufactured from hertalon, a nylon derivative that is both light and elastic.

Wiper is tunable to different configurations (angle of inclination, sensors design, actuation level) and easy to instrument. The edges can indeed be directly actuated by two DC motors, as depicted in Fig. 5.2. The setup has been used to study the stabilization of simple periodic orbits in impact juggling, such as the period-one and the period-two depicted in Fig. 5.1.

Wiper can also be rapidly adapted to study human juggling. Replacing the motors by free rotational joints, human subjects can indeed actuate the edges through direct catching (see Fig. 5.3). The stabilization of the period-one juggling task (Fig. 5.1(c)) is easy and fast to learn. Wiper allows to test this stabilization task at different tempi, since the flight time between two impacts depends on the steady-state angle between the two edges. If the angle is large, the tempo will be slow, and vice-versa. In principle, this setup is also suitable to study more complex patterns, e.g. the period-two (Fig. 5.1(b)). This would permit to focus on learning issues and/or feedback selection issues, since the subjects would not be able to keep several pucks in visual tracking at the same time.



Figure 5.2: Picture of Wiper.

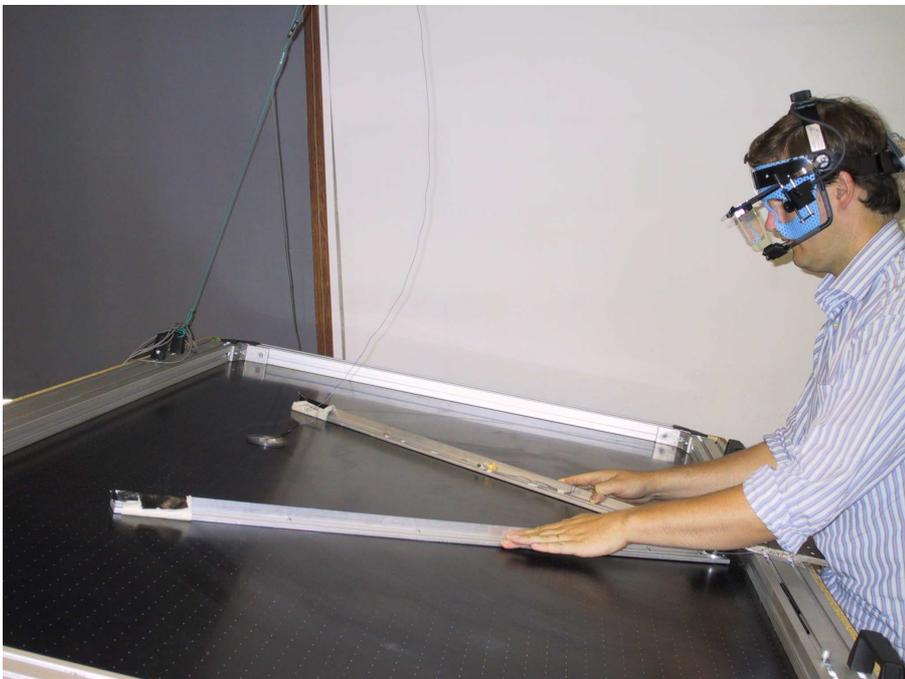


Figure 5.3: Wiper can be actuated by human subjects.

Wiper is also amenable to simple *mathematical modeling*. Similarly to the bouncing ball dynamics (see Section 3.2), the dynamical model consists of planar flight phases separated by impacts:

- (a) during flight phases, the juggled puck trajectories are ballistic flights along a parabola, solution of the Newton's equation $d^2\vec{p}/dt^2 = -\vec{g}$:

$$\begin{aligned} p_{\perp}(t) &= p_{\perp}[k] + v_{\perp}[k]t, \\ p_{//}(t) &= p_{//}[k] + v_{//}[k]t - 0.5gt^2 \end{aligned} \quad (5.1)$$

with $(p_{\perp}[k], p_{//}[k])$ and $(v_{\perp}[k], v_{//}[k])$ denoting the impact position and velocity at impact k , orthogonal and parallel to the gravity field, respectively;

- (b) at impact, there is a sharp discontinuity in the velocity profile. We model the impact with the simplest Newton's law:

$$\begin{aligned} v_n^+ - \dot{s}_n &= -e(v_n^- - \dot{s}_n), \\ v_t^+ - \dot{s}_t &= v_t^- - \dot{s}_t \end{aligned} \quad (5.2)$$

where (v_n^-, v_t^-) and (v_n^+, v_t^+) are the normal and tangential components of the velocity, with respect to the impacting surface, before and after the impact, respectively; and (\dot{s}_n, \dot{s}_t) are the impactor velocity at impact, in the normal and tangential directions w.r.t. the impacting surface. In (5.2), the normal equation is obviously exactly equivalent to (3.3).

The coefficient of restitution $0 \leq e \leq 1$ still models the dissipated energy at impact. The impact model is only a crude approximate of real impact dynamics, since for example it does not capture spin effects of the puck at impact (Spong, 2001). The complete dynamics of Wiper under these simplifying assumptions has been derived in previous papers (Sepulchre and Gerard, 2003; Gerard and Sepulchre, 2005; Ronsse et al., 2007a, 2006) and is derived in the next chapter of this thesis.

A further simplification of Wiper's dynamics is of interest to connect the model with the popular 1D bouncing ball model, studied in Section 3.2 (Holmes, 1982; Guckenheimer and Holmes, 1986; Bapat et al., 1986). Assuming an orthogonal wedge angle and parallel actuations of the edges (that is, the two edges are assumed to remain aligned with the two orthogonal axes of Fig. 5.4), the 2D motion of the juggled objects projects on each axis to a 1D motion that is unaffected by the bounces on the other axis (Sepulchre and Gerard, 2003; Ronsse et al., 2007a), see Fig. 5.4. In this special configuration, a period-two pattern in Wiper corresponds to two frequency-locked period-one bouncing ball patterns (i.e. constant bounce height, see Fig. 3.2) along the axes. The *phase* relationship between those two patterns determines biunivocally the shape of the periodic orbit. The period-one orbit corresponds to two balls bouncing exactly in anti-phase (i.e. one is at the apex when the other bounces). The period-two patterns correspond to any other phase relationship.

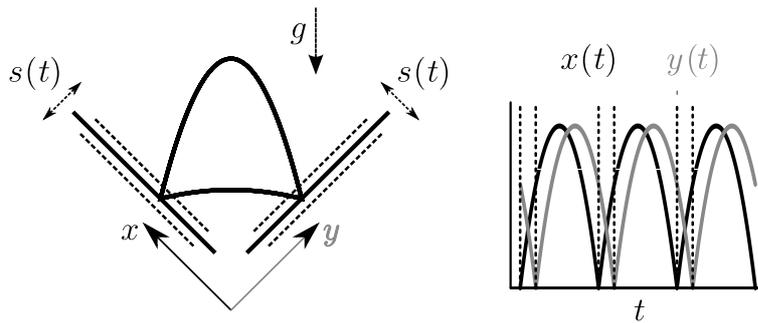


Figure 5.4: A special configuration of the wedge planar juggler that decouples the 2D ball motion (left) into two independent 1D bouncing ball motions (right). The right frame represents the ball trajectories over time along the x and y axes.

5.3 Conclusion

The *Wiper* has been designed from a simplified juggling experiment, aiming at mimicking one of the most popular juggling patterns — i.e. the shower — through its periodic orbits.

A further simplified model of *Wiper*'s dynamics reduces to the bouncing ball dynamics. Its stability properties can consequently be directly inherited from the bouncing ball ones, that are described in Section 3.2. The simplified model is useful to capture in a simple way the main properties of *Wiper*'s periodic orbits. It has been instrumental to develop modeling and design investigations discussed in the next chapters.

This chapter material has been partly published in Ronsse et al. (2007a, 2006, 2007b).

Chapter 6

Sensorless Stabilization of 2D Patterns of the Wiper Robot

Remember that there is nothing
stable in human affairs;
therefore avoid undue elation in
prosperity, or undue depression
in adversity.

Socrates

6.1 Introduction

The goal of the present chapter is to characterize a general model of Wiper's dynamics, their periodic orbits, and their *open-loop* stability properties. Indeed, since a crude approximation of Wiper's dynamics reduces to the bouncing ball model (see Section 5.2), it is sounded to investigate whether the mere sinusoidal actuation (3.7) also stabilizes some periodic pattern in our 2D juggling setup.

The control studied in this chapter is then a sinusoidal actuation of the edges, perhaps the simplest imitation of the fundamental cyclic motion of the hands of a juggler. Sinusoidal motion is also the simplest output of the oscillating circuits that have been used so far for modeling rhythmic movements (see Section 2.1.4). The vibration *frequency* of the edges is the key parameter since the stabilization mechanism rests on a synchronization (or frequency-locking) between the controller and the tempo of the juggling pattern. The vibration *amplitude* is a critical control parameter since the periodic orbits of Wiper are stable for a restricted amplitude range, as it was previously illustrated for the bouncing ball (3.17).

In this chapter, we describe the dynamics of a wedge-billiard, as a mathematical model of the Wiper planar juggler. This model is valid at any wedge angle, while the simplified version derived in Section 5.2 was restricted to the square configuration ($\theta = 45^\circ$). We analyze the existence and stability of periodic orbits of the

wedge-billiard. The chapter main result is to show that some exponentially unstable periodic orbits of the elastic model are stabilized in the non-elastic model with the sinusoidal control that uses no feedback measurement, hence the name *sensorless stabilization*, and to provide an experimental validation of this result. The result somewhat contradicts the intuition that the stabilization of unstable steady-states normally involves feedback. For the period-one orbit also discussed in this chapter, the result has been previously observed by Schaal and Atkeson (1993), at least numerically, who reported the planar juggler as one example of rhythmic system that can be stabilized without feedback.

Section 6.2 presents the wedge-billiard model and outlines some limitations of this model with respect to the expected real Wiper dynamics. In Section 6.3, the steady-state solutions of this model are derived. They correspond to fixed points of the iteration map, and therefore to periodic orbits of the wedge-billiard. In the unactuated, elastic ($e = 1$) wedge-billiard, none of these periodic orbits are stable.

These periodic orbits can nevertheless be stabilized through sinusoidal actuation. First this is demonstrated in Section 6.4 by considering parallel actuations of the edges. As mentioned in Section 5.2, this problem reduces to the analysis of bouncing ball dynamics in the square configuration ($\theta = 45^\circ$). Then the stability properties are straightforwardly inherited. In Section 6.5, the model is further generalized to account for the rotational actuation of the edges. The expected parametric stability region of Wiper are derived. Stability of the periodic orbits depends on the actuation amplitude A , the coefficient of restitution e , and — less intuitively — on the angle of impact θ . In particular, the periodic orbits that were all instable for $\theta > 45^\circ$ in the unactuated and elastic wedge-billiard, can be stabilized with the sensorless sinusoidal actuation.

Finally, Section 6.6 describes experimental validations of open-loop stabilization of Wiper's periodic orbits. The experimental parametric stability regions obtained with the setup described in Section 5.2 are compared with the theoretical parametric stability regions, derived in Section 6.5.

6.2 A bounce juggler model

The aim of the present section is to derive a crude mathematical model of the Wiper robot, presented in Chapter 5. This robot is an idealization of a human juggler. We consider a motion restricted to a plane under a constant gravitational field \mathbf{g} (with $|\mathbf{g}| = g$). The juggled ball undergoes collisions with two edges, which act as the juggler arms (see Fig. 6.1). In contrast with human juggler, the impacts between Wiper's edges and the ball are supposed to be instantaneous, hence the name *impact* (or *bounce*) juggling.

Our model is a control version of the model introduced by Lehtihet and Miller (1986) and was first presented by Sepulchre and Gerard (2003) in order to study several closed-loop *control* laws to stabilize impact juggling patterns. The four-

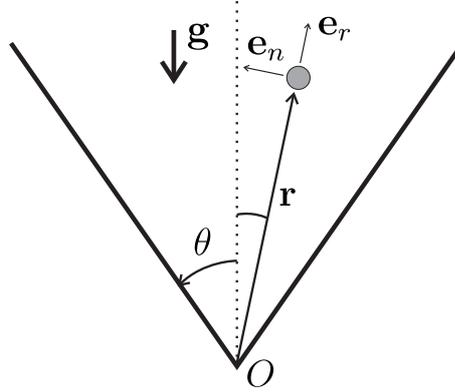


Figure 6.1: The wedge-billiard.

dimensional wedge-billiard dynamics are studied via the three-dimensional discrete Poincaré map relating the state *from one impact to the next one*, the ball motion between two impacts being parabolic (a ballistic flight in a constant gravitational field \mathbf{g}). Let $(\mathbf{e}_r, \mathbf{e}_n)$ be an orthonormal frame attached to the fixed point O with \mathbf{e}_r aligned with the ball position vector $\mathbf{r} = r\mathbf{e}_r$. The ball is assumed to be a unit mass point, let $\mathbf{v} = v_r\mathbf{e}_r + v_n\mathbf{e}_n$ denote its velocity. Therefore the discrete state vector denotes the state of the ball at impacts. This state being discontinuous at impacts, we choose the post-impact values to make up the state vector as a convention¹:

$$\mathbf{x}[k] \equiv \begin{pmatrix} V_r[k] \\ |V_n|[k] \\ R[k] \end{pmatrix} = \begin{pmatrix} V_r^+(t[k]) \\ |V_n^+(t[k])| \\ R^+(t[k]) \end{pmatrix}$$

where $V_r = v_r / \cos \theta$, $V_n = v_n / \sin \theta$ and $R = r / \cos \theta$ denote the state of the ball and the $\bullet^+[k]$ notations denote the post-impact values, evaluated at impact time $t[k]$. The corresponding pre-impact values are denoted $\bullet^-[k]$. We consider the absolute value of the normal velocity, the wedge-billiard being symmetric with respect to its bisecting line.

According to this state vector, the impact law of the wedge-billiard \mathcal{I} derives from (5.2):

$$\begin{aligned} V_r[k] &= V_r^-[k], \\ V_n[k] &= -eV_n^-[k] + (1+e)\dot{S}[k] \end{aligned} \tag{6.1}$$

where $\dot{S}[k] = \dot{S}(t[k]) = \dot{s}(t[k]) / \sin \theta$ and $\dot{s}(t[k])$ denotes the edge velocity at the impact time $t[k]$ and is consequently the system control input.

Assuming that the change in the edges position can be neglected to compute the ball flight map (this assumption rests basically on a *small amplitude* actuation

¹A similar choice has been made for the bouncing ball state vector, see Section 3.2.

of the edges, see the discussion in Section 6.5), the flight map is the velocity and position update of the ball integrated through a flight between two impacts on the *unactuated* wedge-billiard, i.e. (5.1). Two different flight maps must be considered whether these impacts occur on the same edge or not. These flight maps have been derived by Sepulchre and Gerard (2003). The rest of this thesis focusing only on solutions, and stability properties, where the ball hits the edges alternately, the wedge-billiard map \mathcal{B} is therefore the composition of the impact rule \mathcal{I} (6.1) and the second flight map (Lehtihet and Miller, 1986; Sepulchre and Gerard, 2003):

$$\begin{aligned} V_r^-[k+1] &= |V_n|[k] - V_r[k] - |V_n^-|[k+1], \\ |V_n^-|[k+1] &= \sqrt{\left(\frac{2V_r[k] + (\alpha^2 - 1)|V_n|[k]}{1 + \alpha^2}\right)^2 + \frac{4g}{1 + \alpha^2}R[k]} \end{aligned} \quad (6.2)$$

with

$$\alpha = \tan \theta.$$

One obtains the discrete billiard map \mathcal{B} :

$$\begin{aligned} V_r[k+1] &= |V_n|[k] - V_r[k] - |V_n^-|[k+1], \\ V_n[k+1] &= -e|V_n^-|[k+1] \operatorname{sign}(V_n[k]) + (1+e)\dot{S}[k+1], \\ R[k+1] &= R[k] - \frac{1}{2g}(V_r^2[k+1] - V_r^2[k]) - \frac{\alpha^2}{2g}(|V_n^-|^2[k+1] - V_n^2[k]). \end{aligned} \quad (6.3)$$

The position update of (6.3) derives from the *energy* expression:

$$E[k] = \frac{1}{1 + \alpha^2} \left(\frac{1}{2}V_r^2[k] + \frac{\alpha^2}{2}V_n^2[k] + gR[k] \right) \quad (6.4)$$

and the conservation of energy through the flight implies: $E^-[k+1] = E[k]$.

For later reference, one also notes the flight time, i.e. the time between two consecutive impacts, given by:

$$\begin{aligned} \Delta t[k] &\equiv t[k+1] - t[k] \\ &= \frac{1}{g} \left(|V_n^-|[k+1] + \frac{(\alpha^2 - 1)|V_n|[k] + 2V_r[k]}{1 + \alpha^2} \right). \end{aligned} \quad (6.5)$$

The wedge-billiard model \mathcal{B} (6.3) differs clearly from the actual Wiper dynamics. These differences are covered by the following mechanical assumptions:

1. The edges are not affected by the impacts, so that their velocity is continuous at impact times.
2. The impacts are localized in space around the fixed wedge (“small angle” assumption).
3. The contact is frictionless.

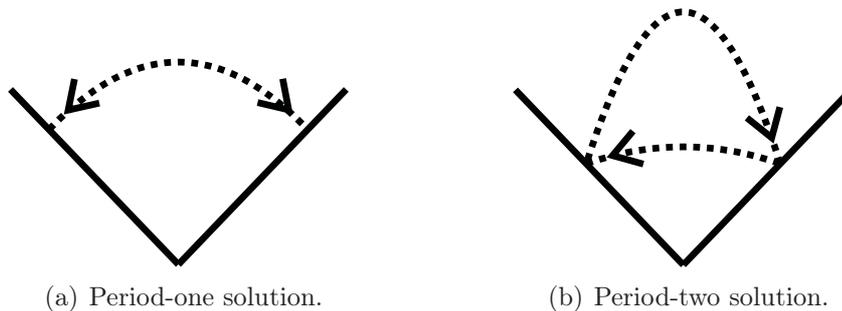


Figure 6.2: Two periodic orbits of the wedge-billiard. These periodic orbits are equivalent to Wiper’s periodic orbits, depicted in Fig. 5.1.

Assumption 1 relies on the fact that the edges are largely heavier than the puck. If the edges motors track a smooth reference signal, the perturbation caused at one impact time is rejected by the next impact time. Note that a finite mass ratio between the actuator and the object can be captured by a modified coefficient of restitution: Vincent and Mees (2000) studied a 1D bouncing ball model capturing this effect. The effect of assumption 2 is more tedious to analyse. Relaxing the “small angle” assumption leads to an implicit billiard map that is much more complicated to solve. Assumption 3 is reasonable. Friction at impacts induces the puck to *spin*. This leads to a more complicated model that will have topologically similar steady-state orbits. See the discussion in Section 6.6.2. The impact controllability of an air hockey puck has been studied by Spong (2001).

6.3 Energy balance of the steady-state solutions

One of the most common juggling patterns is called the *shower*. It involves a circle-shaped trip of several balls between the juggler hands (Fig. 5.1(a)). This section will describe and analyze a periodic solution of the model (6.3) which is very close to the shower pattern, i.e. the *period-two* orbit described in Chapter 5.

The first part of this section describes a degenerate case of the period-two orbits where both flight times are equal, i.e. the *period-one* orbit. The general period-two solution of the wedge-billiard will be investigated in Section 6.3.2. Stability properties of these solutions in the *unactuated elastic* wedge will be also investigated.

6.3.1 Period-one orbit

The period-one orbit is a round trip of the ball between the edges, both trajectories being exactly the same (Fig. 6.2(a)). Due to its symmetry, this periodic motion is characterized by a *unique* energy level E^* . The radial velocity V_r and the position R being conserved at impacts, the conservation of energy implies the conservation

of the square normal velocity:

$$(V_n^*)^2 = (V_n^{-*})^2. \quad (6.6)$$

Using (6.6), we obtain in the steady-state solution of (6.3):

$$V_r^* = |V_n^*| - V_r^* - |V_n^{-*}|, \quad (6.7)$$

$$|V_n^*| = e |V_n^{-*}| + (1 + e)|\dot{S}^*| \quad (6.8)$$

which implies $V_r^* = 0$. These results have a direct geometrical interpretation on Fig. 6.2(a): at the impacts, the radial velocity must be zero, and the normal velocity must be exactly reversed, for this steady-state motion. The steady-state edges velocity derives from (6.8):

$$|\dot{S}^*| = \frac{1 - e}{1 + e} |V_n^*|. \quad (6.9)$$

Using (6.4), (6.5) and (6.2), the fixed point of (6.3) is conveniently parametrized by the energy E^* :

$$|V_n^*| = \frac{1 + \alpha^2}{\alpha} \sqrt{\frac{2E^*}{3 + \alpha^2}}, \quad (6.10)$$

$$R^* = \frac{2(1 + \alpha^2)}{g(3 + \alpha^2)} E^*, \quad (6.11)$$

$$\Delta t^* = \frac{2\alpha}{g} \sqrt{\frac{2E^*}{3 + \alpha^2}}. \quad (6.12)$$

If the edges are elastic ($e = 1$), no energy supply is needed to sustain the period-one motion (see (6.9): $|\dot{S}^*| = 0$). This could be also simulated from a non-elastic wedge-billiard if the edges compensate for the energy dissipation at each impact: $|\dot{S}^*|[k + 1] = (1 - e)/(1 + e)|V_n^-|[k + 1]$.

Proposition 3 (see Lehtihet and Miller (1986)) *For every $\theta \in (0^\circ, 90^\circ)$ and for every energy level E^* , the wedge-billiard possesses a unique period-one orbit, determined by the fixed point (6.10), (6.11), (6.12) of the model (6.3). This orbit is marginally stable if $\theta < 45^\circ$ ($\alpha < 1$), unstable if $\theta = 45^\circ$ ($\alpha = 1$) and exponentially unstable if $\theta > 45^\circ$ ($\alpha > 1$).*

Proof Stability of the period-one orbits is investigated via the Jacobian linearization of (6.3) at the fixed point (6.10), (6.11), (6.12) which gives:

$$\begin{pmatrix} \delta V_r[k + 1] \\ \delta |V_n|[k + 1] \\ \frac{g}{|V_n^*|} \delta R[k + 1] \end{pmatrix} = \begin{pmatrix} \frac{1 - 4\alpha^2 - \alpha^4}{(1 + \alpha^2)^2} & \frac{4\alpha^2}{(1 + \alpha^2)^2} & \frac{-2}{1 + \alpha^2} \\ \frac{2(\alpha^2 - 1)}{(1 + \alpha^2)^2} & \frac{(\alpha^2 - 1)^2}{(1 + \alpha^2)^2} & \frac{2}{1 + \alpha^2} \\ \frac{2\alpha^2(1 - \alpha^2)}{(1 + \alpha^2)^2} & \frac{4\alpha^4}{(1 + \alpha^2)^2} & \frac{1 - \alpha^2}{1 + \alpha^2} \end{pmatrix} \begin{pmatrix} \delta V_r[k] \\ \delta |V_n|[k] \\ \frac{g}{|V_n^*|} \delta R[k] \end{pmatrix}. \quad (6.13)$$

The eigenvalues of this Jacobian matrix are:

$$\lambda_{1,2,3} = 1, \frac{1 - 4\alpha^2 - \alpha^4 \pm 2\alpha\sqrt{(\alpha^2 - 1)(\alpha^2 + 3)}}{(1 + \alpha^2)^2}. \quad (6.14)$$

One unitary eigenvalue ($\lambda_1 = 1$) is associated with the conservation of energy. The remaining two eigenvalues (λ_2 and λ_3) lie on the unitary circle for $\alpha < 1$, and it can be shown that the period-one solution is marginally stable in this case: in the state space plane of the elastic wedge-billiard, the period-one fixed point is surrounded by a continuum of closed orbits that correspond to *quasi-periodic* solutions (Lehtihet and Miller, 1986). When the wedge is a right angle ($\alpha = 1$), $\lambda_2 = \lambda_3 = -1$: the map \mathcal{B} has therefore an eigenvalue of algebraic multiplicity 2 on the unit circle. The unactuated wedge-billiard dynamics becoming *linear* in that case (see Sepulchre and Gerard, 2003), the period-one orbits are unstable. For $\alpha > 1$, the two eigenvalues are real, one of them being outside the unitary circle, the period-one solution is therefore exponentially unstable. \square

6.3.2 Period-two orbit

Period-two orbits model the shower juggling pattern presented in Fig. 5.1(a). The round trip of the ball between the edges is now characterized by two different trajectories depending on the direction of the ball. Fig. 6.2(b) is an example of a period-two solution. This solution will be characterized by two parameters E^{r^*} and E^{l^*} , associated to the ball energy on *each* trajectory, assuming arbitrarily that \bullet^{r^*} characterizes the right-edge impacts and \bullet^{l^*} the left-edge impacts. A positive edge velocity corresponds to a counterclockwise motion. The period-one orbit is a degenerate case of the period-two orbits for which $E^{r^*} = E^{l^*}$.

A period-two solution is characterized by two points in the state space, say \mathbf{x}^{r^*} and \mathbf{x}^{l^*} : these points ought to be fixed points of $\mathcal{B}^2 = \mathcal{B} \circ \mathcal{B}$. The position and the radial velocity being conserved at impact, the global energy balance of the period-two solution requires this time:

$$(V_n^{r^*})^2 - (V_n^{-r^*})^2 = - \left((V_n^{l^*})^2 - (V_n^{-l^*})^2 \right). \quad (6.15)$$

A possible loss of energy on one edge has to be compensated on the other one.

Eliminating the radial velocities $V_r^{r^*}$ and $V_r^{l^*}$ from their update equations:

$$V_r^{l^*} = |V_n^{r^*}| - V_r^{r^*} - |V_n^{-l^*}|, \quad (6.16)$$

$$V_r^{r^*} = |V_n^{l^*}| - V_r^{l^*} - |V_n^{-r^*}|, \quad (6.17)$$

we find another relation between the normal velocities:

$$|V_n^{r^*}| - |V_n^{-l^*}| = |V_n^{l^*}| - |V_n^{-r^*}|. \quad (6.18)$$

Equations (6.15) and (6.18) yield:

$$|V_n^{r^*}| + |V_n^{-l^*}| = -(|V_n^{l^*}| + |V_n^{-r^*}|). \quad (6.19)$$

Thanks to (6.18) and (6.19), one finds $|V_n^{-l^*}| = -|V_n^{l^*}|$ and $|V_n^{-r^*}| = -|V_n^{r^*}|$. These four variables being non-negative, (6.18) must be equal to zero to be satisfied, such as (6.15):

$$\begin{aligned} |V_n^{-l^*}| &= |V_n^{r^*}|, \\ |V_n^{-r^*}| &= |V_n^{l^*}|. \end{aligned} \quad (6.20)$$

The radial velocity and the impact position should then satisfy (see (6.16) and (6.3)):

$$V_r^{l^*} = -V_r^{r^*}, \quad (6.21)$$

$$R^{r^*} = R^{l^*} = R^*. \quad (6.22)$$

Geometrically, the period-two solutions are therefore characterized by two symmetrical parabolas: these parabolas reach their highest point (zenith) on the bisecting line of the wedge.

Introducing the normal velocity conservation (6.20) in (6.3), we obtain the steady-state relations:

$$\begin{aligned} |V_n^{l^*}| &= e|V_n^{r^*}| - (1+e)\dot{S}^{l^*}, \\ |V_n^{r^*}| &= e|V_n^{l^*}| + (1+e)\dot{S}^{r^*}. \end{aligned} \quad (6.23)$$

Equations (6.23) provide the steady-state edges velocities:

$$\begin{aligned} \dot{S}^{r^*} &= \frac{|V_n^{r^*}| - e|V_n^{l^*}|}{1+e}, \\ \dot{S}^{l^*} &= \frac{e|V_n^{r^*}| - |V_n^{l^*}|}{1+e}. \end{aligned} \quad (6.24)$$

Injecting (6.20) in the definitions of $|V_n^-|$, we find:

$$(V_n^{r^*})^2 = \left(\frac{2V_r^{r^*} + (\alpha^2 - 1)|V_n^{r^*}|}{1 + \alpha^2} \right)^2 + \frac{4g}{1 + \alpha^2} R^*, \quad (6.25)$$

$$(V_n^{l^*})^2 = \left(\frac{2V_r^{l^*} + (\alpha^2 - 1)|V_n^{l^*}|}{1 + \alpha^2} \right)^2 + \frac{4g}{1 + \alpha^2} R^* \quad (6.26)$$

which implies, taking (6.21) into account:

$$V_r^{r^*} = \frac{\alpha^2}{\alpha^2 - 1} (|V_n^{r^*}| - |V_n^{l^*}|) = -V_r^{l^*}. \quad (6.27)$$

The impact position and the energy levels are derived from (6.25), (6.27), (6.4) and (6.28):

$$R^* = \frac{\alpha^2 \left((1 + \alpha^2)^2 |V_n^{r*}| |V_n^{l*}| - \alpha^2 (|V_n^{r*}| + |V_n^{l*}|)^2 \right)}{g(1 + \alpha^2)(\alpha^2 - 1)^2}, \quad (6.28)$$

$$E^{r*} = \frac{\alpha^2 \left((\alpha^4 + \alpha^2 - 1)(V_n^{r*})^2 + \alpha^2 (V_n^{l*})^2 - 2|V_n^{r*}| |V_n^{l*}| \right)}{2(\alpha^2 - 1)(1 + \alpha^2)^2},$$

$$E^{l*} = \frac{\alpha^2 \left((\alpha^4 + \alpha^2 - 1)(V_n^{l*})^2 + \alpha^2 (V_n^{r*})^2 - 2|V_n^{r*}| |V_n^{l*}| \right)}{2(\alpha^2 - 1)(1 + \alpha^2)^2}. \quad (6.29)$$

If the energy levels denote the two parameters of the period-two solution, (6.29) must be inverted to find the normal velocities that can be replaced in the solution equations.

The flight times are derived from (6.5):

$$\Delta t^{r*} = \frac{2\alpha^2 (\alpha^2 |V_n^{r*}| - |V_n^{l*}|)}{g(\alpha^4 - 1)}, \quad (6.30)$$

$$\Delta t^{l*} = \frac{2\alpha^2 (\alpha^2 |V_n^{l*}| - |V_n^{r*}|)}{g(\alpha^4 - 1)}.$$

These relationships correspond to the period-one solution ((6.10), (6.11) and (6.12)) if $E^{r*} = E^{l*}$ ($|V_n^{r*}| = |V_n^{l*}|$).

An elastic wedge can be simulated from non-elastic edges by adding $(1 - e)/(1 + e)|V_n^-|[k + 1]$ to the edges velocity.

Proposition 4 *Period-two orbits exist in the unactuated elastic wedge-billiard only for $\theta = 45^\circ$. They are uniquely defined by their energy level E^* and their impact radial velocity V_r and are unstable. For every $\theta \in (0^\circ, 90^\circ)$ and for every pair (E^{r*}, E^{l*}) , the actuated elastic wedge-billiard possesses a unique period-two orbit determined by the fixed point of $\mathcal{B}^2 = \mathcal{B} \circ \mathcal{B}$ (6.3). If $\theta < 45^\circ$, this orbit is either marginally stable or exponentially unstable, depending on the energy difference $E^{r*} - E^{l*}$. It is unstable if $\theta = 45^\circ$ and exponentially unstable if $\theta > 45^\circ$.*

Proof First, we derive the period-two orbits in the *square* wedge-billiard ($\theta = 45^\circ$, $\alpha = 1$) because several equations previously derived become singular. In this case, (6.21), (6.25) and (6.26) force the normal velocities to be equal:

$$|V_n^{r*}| = |V_n^{l*}| = |V_n^*| = |V_n^{-r*}| = |V_n^{-l*}| \quad (6.31)$$

so that each trajectory has the same energy: $E^{r*} = E^{l*} = E^*$. The second parameter of the period-two orbit is the radial velocity V_r^* since (6.27) is undetermined when $\alpha = 1$.

The period-two solution of the square wedge-billiard is then derived from (6.4), (6.2) and (6.5) with $\alpha = 1$:

$$|V_n^*| = \sqrt{2E^*}, \quad (6.32)$$

$$R^* = \frac{1}{2g} (2E^* - (V_r^*)^2), \quad (6.33)$$

$$\Delta t^{r^*}, \Delta t^{l^*} = \frac{\sqrt{2E^*} \pm |V_r^*|}{g}. \quad (6.34)$$

From (6.24), we have:

$$\dot{S}^{r^*} = -\dot{S}^{l^*} = \frac{|V_n^{r^*}| - |V_n^{l^*}|}{2}. \quad (6.35)$$

The steady-state edges velocities \dot{S}^{r^*} and \dot{S}^{l^*} are therefore equal to zero only if $|V_n^{r^*}| = |V_n^{l^*}|$, i.e. in the square configuration. Period-two orbits exist in the *unactuated* elastic wedge-billiard only in that configuration. Period-two orbits exist with any other wedge angle but require an actuation of the edges, according to (6.35).

Stability of these orbits is studied via the linearization of $\mathcal{B}^2 = \mathcal{B} \circ \mathcal{B}$ (with \mathcal{B} given in(6.3)) around its fixed point. We find the following jacobian matrix:

$$\begin{aligned} \left(\begin{array}{c} \delta V_r[k+2] \\ \delta |V_n|[k+2] \\ \frac{g}{\sqrt{|V_n^{r^*}||V_n^{l^*}|}} \delta R[k+2] \end{array} \right) &= \mathcal{M}_2 \left(\begin{array}{c} \delta V_r[k+1] \\ \delta |V_n|[k+1] \\ \frac{g}{\sqrt{|V_n^{r^*}||V_n^{l^*}|}} \delta R[k+1] \end{array} \right) \\ &= \mathcal{M}_2 \mathcal{M}_1 \left(\begin{array}{c} \delta V_r[k] \\ \delta |V_n|[k] \\ \frac{g}{\sqrt{|V_n^{r^*}||V_n^{l^*}|}} \delta R[k] \end{array} \right) \end{aligned} \quad (6.36)$$

with

$$\mathcal{M}_i = \left(\begin{array}{ccc} \frac{1-4\alpha^2-\alpha^4+4\alpha^2\chi_i}{(1+\alpha^2)^2} & \frac{2\alpha^2(1+\tau_i)}{(1+\alpha^2)^2} & \frac{-2\sqrt{\tau_i}}{1+\alpha^2} \\ \frac{\alpha^4+1-2\alpha^2}{(\alpha^2-1)(1+\alpha^2)^2} & \frac{\alpha^4+1-2\alpha^2}{(1+\alpha^2)^2} & \frac{2}{(1+\alpha^2)\sqrt{\tau_i}} \\ \frac{2\alpha^2(1-\alpha^2+(1+\alpha^2)\chi_i+2\alpha^2\chi_i^2)}{(1+\alpha^2)^2\sqrt{\tau_i}} & \frac{2\alpha^4(1+\tau_i)(1+\chi_i)}{(1+\alpha^2)^2\sqrt{\tau_i}} & \frac{1-\alpha^2-2\alpha^2\chi_i}{1+\alpha^2} \end{array} \right)$$

and $\chi_i = (1 - \tau_i)/(1 - \alpha^2)$ for $i = 1, 2$. The parameter $\tau_1 = |V_n^{r^*}|/|V_n^{l^*}| = 1/\tau_2$ denotes the normal velocity ratio. These parameters capture the energy difference between both parabolas and have been defined as such only for computational convenience.

The eigenvalues of (6.36) were numerically computed, depending on α and τ_1 . Fig. 6.3 depicts the stability regions of the period-two solution in the parameter space (θ, τ_1) . As in the period-one stability analysis, one of the eigenvalues is always equal to 1, the energy being conserved over two impacts in the period-two orbit. The

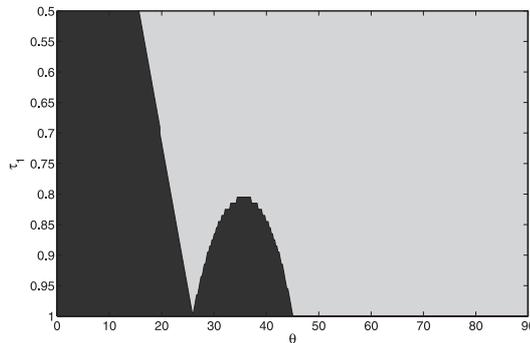


Figure 6.3: Stability regions of the period-two solution: marginal stability (black) and exponential instability (grey).

two remaining eigenvalues either are complex and lie on the unitary circle (black zone, marginal stability) or are real, at least one of them being outside the unitary circle (grey zone, instability).

Similarly to what was observed for the period-one orbit², Fig. 6.3 shows that all period-two orbits are exponentially unstable for $\theta > 45^\circ$. In contrast, marginally stable orbits coexist with unstable orbits for $\theta < 45^\circ$, depending on the energy difference $|E^{r*} - E^{l*}|$ via the ratio τ_1 . Finally, for $\theta = 45^\circ$, $|V_n^{r*}| = |V_n^{l*}|$ induces $\tau_1 = \tau_2 = 1$. The three eigenvalues of $\mathcal{M}_2\mathcal{M}_1$ are equal to 1. Because the dynamics become linear (see Sepulchre and Gerard, 2003), then the period-two orbits are unstable. \square

It is of interest to observe the unstable behavior of all period-two orbits for a particular value of θ about 26° . This value corresponds to $(1 - \alpha^2)/((1 + \alpha^2)^2) = 0.5$ (i.e. $\alpha = \sqrt{\sqrt{5} - 2}$), in which case two eigenvalues of the linearized system (6.13) are equal to $\pm i$. The linearized system (6.36) corresponds to a double iteration of (6.13) when $\tau_1 = 1$, resulting in two eigenvalues equal to $(\pm i)^2 = -1$. They split into two real eigenvalues, one of them outside the unit circle, when $\tau_1 \neq 1$. The same critical value of θ causes *uncontrollability* of a linearized model of the elastic wedge-billiard which is controlled with one edge. For this particular aperture, the dynamics of the normal velocity, the radial velocity and the energy decouple. At first order, the control input leaves each of these dynamics invariant (Gerard and Sepulchre, 2004).

²Note that the period-one eigenvalues (6.14) lie on a “slice” of Fig. 6.3, for $\tau_1 = \tau_2 = 1$ ($\chi_1 = \chi_2 = 0$). In that case, \mathcal{M}_1 and \mathcal{M}_2 are equal to (6.13).

6.4 Periodic orbits of the sinusoidally actuated wedge-billiard

None of the periodic orbits studied in Section 6.3 are attractors: either the linearized system is marginally stable or unstable. In this section, we introduce how a simple periodic actuation of the wedge can *isolate* and *stabilize* one of the periodic orbits, characterized by its energy level(s) (E^* for a period-one orbit or E^{r^*} and E^{l^*} for a period-two orbit).

The stabilization method presented in this section is completely similar to the one studied by Holmes (1982) in the bouncing ball dynamics (see Section 3.2.2), i.e. a sinusoidal actuation of the edges (3.7). Such an harmonic motion is the closest imitation of the fundamental cyclic motion of the juggler's hands. Moreover, we assume that this edges remain synchronized in anti-phase (in intrinsic space), such that their aperture remains constant throughout the cycles.

As previously mentioned, the special configuration of the square wedge-billiard turns out to bridge the bouncing ball dynamics studied in Section 3.2 and the wedge dynamics.

6.4.1 The square wedge-billiard

The steady-state velocity-energy relation of the square wedge-billiard (6.32) has exactly the same form as for a bouncing ball³, emphasizing the decoupling of the square wedge: the dynamics along each edge can be viewed as a 1 DOF independent bouncing ball motion that is unaffected by the bounces on the other axis (Sepulchre and Gerard, 2003), see Chapter 5 (Fig. 5.4). The parameter V_r^* determines the phase shift between the two bouncing balls.

The steady-state regime is characterized by two frequency-locking relations between the ball and the wedge (Ronsse et al., 2004):

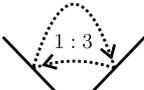
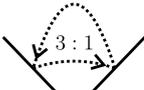
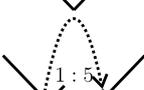
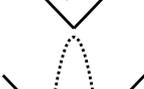
$$(t[k+2] - t[k])^* \equiv \Delta t^{r^*} + \Delta t^{l^*} = n \frac{2\pi}{\omega}, \quad (6.37)$$

$$(t[k+1] - t[k])^* \equiv \Delta t^{r^*} = (2m-1) \frac{\pi}{\omega} \quad (6.38)$$

where $(\bullet)^*$ denotes the steady-state solutions. These relations rest on the trivial assumption $m \leq n$, with m and $n \in \mathbb{N}$. Eq. (6.37) expresses that the ball period is a multiple of the edge vibration period: this is the frequency-locking relation of each dynamics similar to (3.10). Eq. (6.38) expresses that the phase difference between two successive impacts must be equal to an odd multiple of the vibration half-frequency. As a convention, m will be associated with the flight time between the right edge and the left edge: $\Delta t^{r^*} = (2m-1)\pi/\omega$, while the flight time between

³Consider simply that the bouncing ball dynamics refer to an exchange of a maximal *potential* energy E^* (at the top-point) and a maximal *kinetic* energy $\propto (V_n^*)^2$ (just before impacts), i.e. $|V_n^*| \propto \sqrt{E^*}$.

Table 6.1: Periodic orbits for the square wedge-billiard. $(x : y)$ denotes the ratio between both flight times where x (y) is associated to the flight from left to right (from right to left).

	m=1	m=2	m=3	m=4	...
n=1		X	X	X	
n=2			X	X	
n=3				X	...
n=4					
⋮			⋮		

the left edge and the right edge will be $\Delta t^{l*} = (2n - 2m + 1)\pi/\omega$ in the steady-state regime. Table 6.1 illustrates the first periodic orbits for the vibrating square wedge-billiard and the ratios between the low toss and the high toss flight times for each of these orbits. A sustained steady-state shower pattern will be characterized exactly by the same ratios, n denoting also the number of juggled balls. It is interesting to point out how the symmetry of the square wedge-billiard captures the symmetry of the juggler behavior. Beek and Lewbel (1995) wrote a very accessible paper on the “scientific aspects of juggling” where they present a compact notation for juggling patterns: *Site-swap* notation represents the order in which props are thrown and caught in each cycle of the juggle, assuming throws happen on beats that are *equally spaced in time*, being the case both in most of the common juggling patterns, and in the square wedge-billiard. The site-swap notation of the 3 balls shower is simply ‘51’, where the 5 refers to the duration of the high toss and the 1 to the time needed to pass the ball from one hand to the other on the lower part of the arc. Each $(x : y)$ orbit presented in Table 6.1 will then be “site-swap” noted ‘ xy ’.

Injecting the flight time solutions (6.34) in (6.37) and (6.38), we obtain the steady-state velocities of the periodic orbits of the square wedge-billiard, while the

energy E^* , the impact position R^* derive from (6.32) and (6.33):

$$E^* = \frac{(n\pi g)^2}{2\omega^2}, \quad (6.39)$$

$$|V_n^*| = \frac{n\pi g}{\omega}, \quad (6.40)$$

$$V_r^{r*} = -\frac{(n-2m+1)\pi g}{\omega} = -V_r^{l*}, \quad (6.41)$$

$$R^* = \frac{g}{2} \left(\frac{\pi}{\omega}\right)^2 (2m-1)(2n-2m+1). \quad (6.42)$$

The steady-state energy is independent of m . The radial velocity naturally corresponds to the period-one solution ($V_r^* = 0$) when $n = 2m - 1$ (see Table 6.1).

By considering the cartoon model introduced in Chapter 5 and a sinusoidal actuation along the axis perpendicular to the edges (see Fig. 5.4), the normal velocity is updated as: $V_n[k+2] = eV_n[k] + (1+e)\dot{S}[k+2]$, exhibiting the decoupling between the dynamics along each edge. The stability properties of the periodic orbits are immediately inherited from the decoupled dynamics, i.e. from the stability properties of the bouncing ball, assuming the small amplitude assumption⁴. Under this assumption, the parametric stability region is slightly different of (3.17) (Holmes, 1982):

$$n\pi \frac{1-e}{1+e} < \frac{\sqrt{2}A\omega^2}{g} < \sqrt{n^2\pi^2 \left(\frac{1-e}{1+e}\right)^2 + 1}. \quad (6.43)$$

Note that only $g/\sqrt{2}$ of the gravity field applies along each axis. The differences between the parametric stability regions characterized by (3.17) and (6.43) have been studied by Bapat et al. (1986). Significant differences between the exact and the approximated model are more likely if e is below about 0.8; nevertheless, the approximated model provides a good description of the qualitative dynamical behavior.

Consequently, each periodic orbit depicted in Table 6.1 is composed of two phase-locked period-one bouncing balls. They are stable if (6.43) is respected.

6.4.2 The general vibrating wedge-billiard

The dynamics of a ball in the general wedge-billiard do not decouple in two independent dynamics along each edge. The frequency-locking relation (6.37) between two successive impacts times on the same edge is still valid. In contrast, we introduce a scaling parameter ρ in (6.38) to take the energy dissymmetry into account:

$$\Delta t^{r*} = (2m-1)\rho \frac{\pi}{\omega}. \quad (6.44)$$

⁴The small amplitude assumption of the bouncing ball assumes that the edges motion is much smaller than the ball motion, the impacts therefore occurring at a constant position, i.e. $s[k]$ constant $\forall k$.

The parameter ρ must fulfill $0 < \rho < \frac{2n}{2m-1}$, tuning the flight time Δt^{r^*} between 0 and $(t[k+2] - t[k])^*$.

Injecting the flight time solutions (6.30) in (6.37), we obtain a relation about the mean of the normal velocities:

$$|V_n^{r^*}| + |V_n^{l^*}| = \frac{\pi n g}{\omega} \frac{1 + \alpha^2}{\alpha^2} \quad (6.45)$$

that must be combined with (6.44) to obtain the steady-state normal velocities of the periodic orbits, while the energy levels, the radial velocities and the impact position derive from (6.27), (6.28) and (6.29):

$$E^{r^*} = \frac{1}{2\alpha^2} \left(\frac{n\pi g}{\omega} \right)^2 \left(\left(\frac{(2m-1)\rho}{2n} \right)^2 (\alpha^2 - 1) + 1 \right), \quad (6.46)$$

$$E^{l^*} = \frac{1}{2\alpha^2} \left(\frac{n\pi g}{\omega} \right)^2 \left(\left(\frac{2n - (2m-1)\rho}{2n} \right)^2 (\alpha^2 - 1) + 1 \right), \quad (6.47)$$

$$|V_n^{r^*}| = \frac{n\pi g}{\omega} \frac{(\alpha^2 - 1) \frac{(2m-1)\rho}{2n} + 1}{\alpha^2}, \quad (6.48)$$

$$|V_n^{l^*}| = \frac{n\pi g}{\omega} \frac{(1 - \alpha^2) \frac{(2m-1)\rho}{2n} + \alpha^2}{\alpha^2}, \quad (6.49)$$

$$V_r^{r^*} = \frac{-\pi g}{\omega} (n - (2m-1)\rho) = -V_r^{l^*}, \quad (6.50)$$

$$R^* = \frac{g(1 + \alpha^2)}{4\alpha^2} \left(\frac{\pi}{\omega} \right)^2 (2m-1)\rho(2n - (2m-1)\rho). \quad (6.51)$$

The energy exchange between the ball and the edges at each impact is:

$$\pm(E^{l^*} - E^{r^*}) = \pm \frac{\alpha^2 - 1}{2\alpha^2} \left(\frac{n\pi g}{\omega} \right)^2 \left(\frac{n - (2m-1)\rho}{n} \right).$$

In the bouncing ball dynamics, the bifurcation parameter is a non-dimensional amplitude proportional to the vibration amplitude A and to the square of the vibration frequency ω^2 (Holmes, 1982; Schaal et al., 1996). Therefore both these parameters can be used to generate the cascade of bifurcations. This non-dimensional parameter depends on ρ . This rests on a trigonometric relation between the steady-state edges phases at impact deriving from $\sin^2 \phi^{r^*} + \cos^2 \phi^{r^*} = 1$:

$$\cos^2 \phi^{r^*} + \cos^2 \phi^{l^*} - 2 \cos \phi^{r^*} \cos \phi^{l^*} \cos((2m-1)\rho\pi) = \sin^2((2m-1)\rho\pi) \quad (6.52)$$

which is not invertible.

Summarizing the derivations in this section, we obtain the following proposition:

Proposition 5 *The wedge-billiard model (6.3) with harmonic actuation $s(t) = A \sin(\omega t)$ isolates particular periodic orbits among those derived in Proposition 4.*

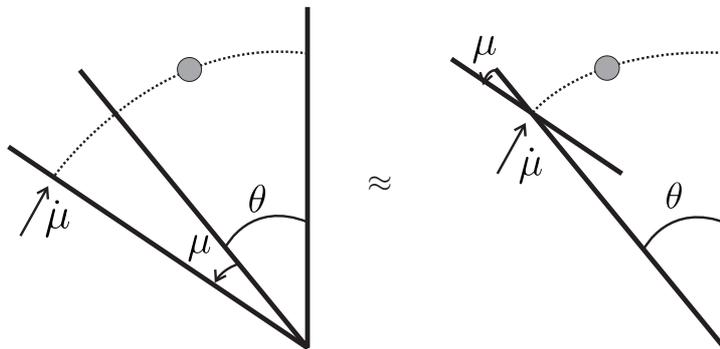


Figure 6.4: The controlled rotational wedge (left), and the simplified model when μ is small (right).

These orbits are characterized by (6.46) to (6.51) via n and m and satisfy frequency-locking relations between the ball and the edges ((6.37) and (6.44)). For the particular square wedge configuration ($\theta = 45^\circ$), (6.44) degenerates to (6.38). Considering a cartoon model for a sinusoidal actuation (the parallel actuation depicted in Fig. 5.4), the dynamics decouple into two 1 DOF dynamics along each edge. Stability of the isolated periodic orbits follows immediately from the bouncing ball stability (6.43).

6.5 Stabilization of periodic orbits by rotational actuation of the edges

The Wiper robot, presented in Section 6.6, uses rotational actuation of the edges instead of linear actuation as depicted in Fig. 5.4. A rotational implementation is indeed much simpler to design, since the arms can be directly actuated by DC motors. We examine in this section how rotational actuation modifies the model studied in the previous sections, both for period-one and period-two orbits.

With rotational actuation, the angle θ of each edge with the vertical is no longer constant, which significantly complicates the derivation of the flight map. To avoid the complication of computing that new flight map, a “small amplitude” assumption is introduced: we neglect the variation of θ in the derivation of the flight map but only take it into account in the derivation of the impact map. As illustrated in Fig. 6.4, this simplification amounts to assume that the impacts always occur at angle θ but that the angular actuation μ rotates the normal and *tangential* directions of the impacted edge by an angle μ (Fig. 6.4, right). This simplification neglects the displacement of the impact point and is more likely if $|\mu| \ll \theta$.

The Poincaré map of the rotational wedge is derived in Appendix 6.A of this chapter.

The general steady-state equations derived in Section 6.4 still hold in that case

(Eqs. (6.46) to (6.51)). The edges velocity being equal to $\dot{\mu}(t) = A\omega \cos(\omega t)$, the steady-state edges phases derive from (6.48), (6.49) and (6.60):

$$\begin{aligned}\phi^{r^*} &= \arccos \left(\frac{1-e}{1+e} \frac{n\pi g}{AR^*\omega^2} \frac{(\alpha^2-1)^{\frac{(1+e)\frac{(2m-1)\rho-e}{1-e}}{2n}} + 1}{\alpha} \right), \\ \phi^{l^*} &= \arccos \left(\frac{1-e}{1+e} \frac{n\pi g}{AR^*\omega^2} \frac{(\alpha^2-1)^{\frac{(1+e)\frac{(2m-1)\rho-1}{1-e}}{2n}} - 1}{\alpha} \right) - \pi.\end{aligned}\quad (6.53)$$

Because the steady-state impact position R^* is proportional to $1/\omega^2$ (6.51), the non-dimensional bifurcation parameter is now proportional to the vibration amplitude A only. The vibration frequency acts as a *temporal scaling factor* and does not influence the stability properties.

To analyze how the modified model affects the stability of periodic orbits, we numerically computed the eigenvalues of the linearized Poincaré map in two simple cases:

- The period-one orbits ($n = 2m - 1$, $\rho = 1$).
- The period-one and period-two orbits in the *square* rotational wedge ($\theta = 45^\circ$, $\rho = 1$).

6.5.1 Stability of period-one orbits

The simplest period-one orbit ($n = m = 1$) has been first mentioned as an open-loop stable orbit in a wedge-billiard with a rotational sinusoidal actuation by Schaal and Atkeson (1993). The stability region presented in this paper is in agreement with our results. This stability region has been obtained from a linearized model which does not rest on a “small amplitude” assumption. Accordance between stability regions obtained from both models is viewed as a validation of our “small amplitude” model.

The period-one orbits are particular cases of Equations (6.46) to (6.51) and Equation (6.53) with $n = 2m - 1$ and $\rho = 1$, which is required for period-one. We find:

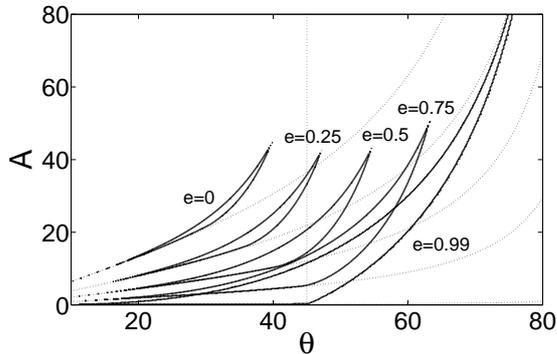
$$E^* = \frac{3 + \alpha^2}{8\alpha^2} \left(\frac{(2m-1)\pi g}{\omega} \right)^2, \quad (6.54)$$

$$|V_n^*| = \frac{1 + \alpha^2}{\alpha^2} \frac{(2m-1)\pi g}{2\omega}, \quad (6.55)$$

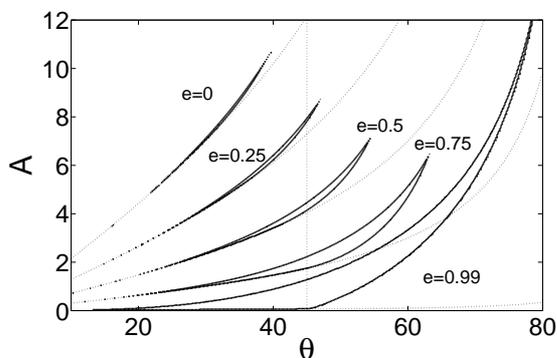
$$V_r^* = 0, \quad (6.56)$$

$$R^* = \frac{1 + \alpha^2}{4\alpha^2} \left(\frac{(2m-1)\pi}{\omega} \right)^2 g, \quad (6.57)$$

$$\phi^* = \arccos \left(\frac{1-e}{1+e} \frac{2\alpha}{A} \frac{1}{(2m-1)\pi} \right). \quad (6.58)$$



(a) $n = 1, m = 1$



(b) $n = 3, m = 2$

Figure 6.5: Parametric stability region of two period-one orbits in the general wedge (solid lines). The dotted lines denote the physical minimum value for the amplitude A (the arccos argument in (6.58) must be ≤ 1).

The stability is studied via the linearized Poincaré map of $\tilde{\mathcal{B}}$ (6.60) and of the flight time (6.5) around the period-one solution just derived. We find the Jacobian matrix (6.63) derived in Appendix 6.B of this chapter. Its eigenvalues were numerically computed for several values of e and α . Fig. 6.5 depicts the stability region for the first two period-one orbits ($n = 1, m = 1$ (a) and $n = 3, m = 2$ (b)). The superposed curves stand for different values of e .

A decreasing coefficient of restitution reduces the stability regions and shifts it in a zone corresponding to closer angles between the edges. We can see that even for $e = 0$, the first period-one orbit is still theoretically stabilizable with a sinusoidal vibration of the edges if the impacts occur with $\theta \in [15^\circ, 40^\circ]$.

We conclude that sinusoidal actuation of edges stabilizes period-one orbits, for any coefficient of restitution $e < 1$ and for a broad domain of wedge geometry. For $\theta > 45^\circ$ this exponential stability is in sharp contrast with the instability of the same periodic orbit in the fixed elastic wedge.

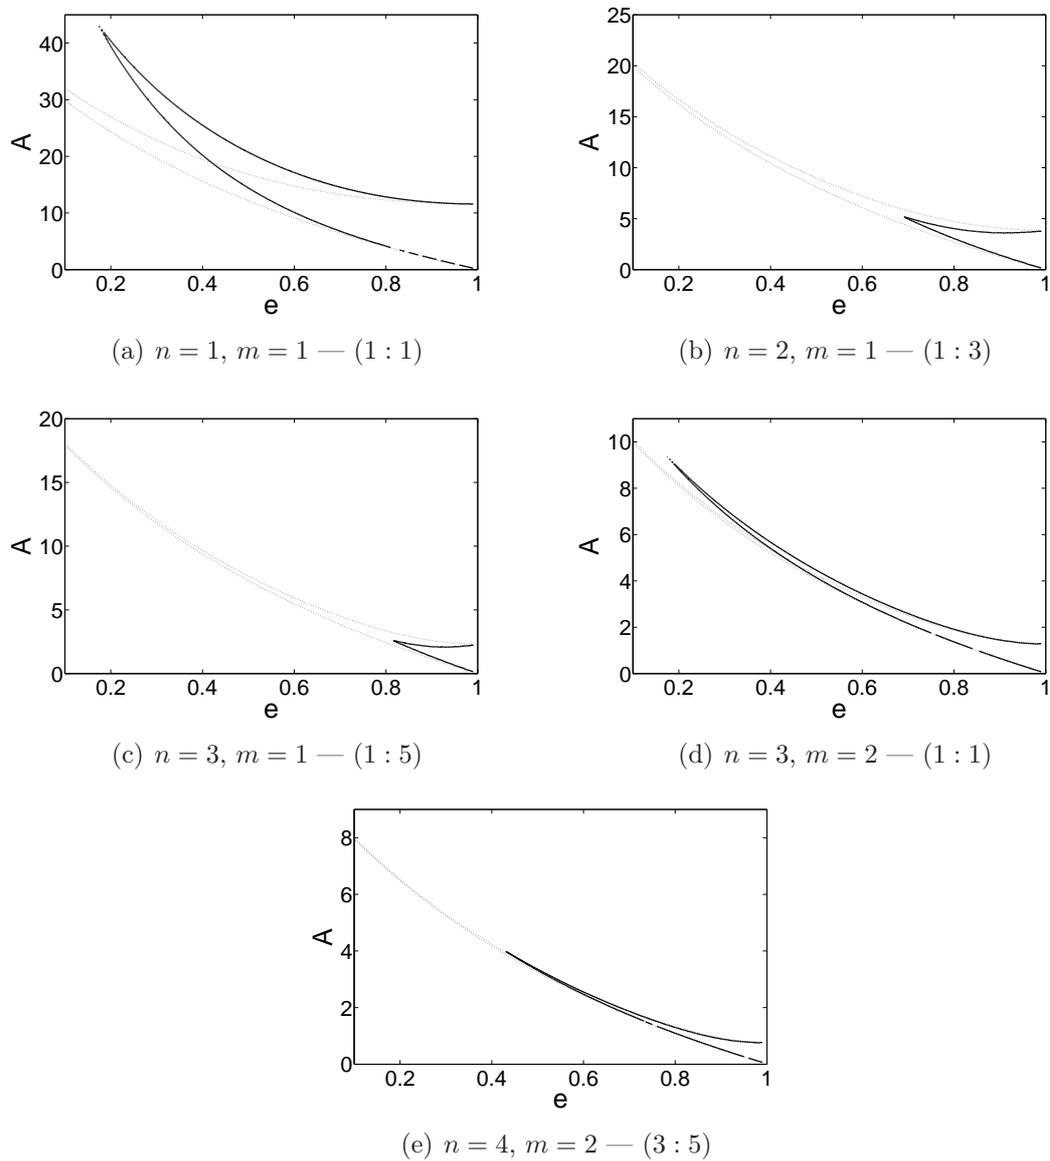


Figure 6.6: Parametric stability region of periodic orbits in the square wedge with linear actuation (dotted) and rotational actuation (solid).

6.5.2 Stability of period-two orbits in the square wedge

With *linear* actuation of the edges, we have shown in Section 6.4.1 that stability of period-two orbits in the square wedge follows from the bouncing ball dynamics analysis, yielding exponentially stable orbits in the parameter range (6.43). We now show that this stabilization result also holds with *rotational* actuation of the edges, by computing the eigenvalues of the Jacobian linearization of the map $\tilde{\mathcal{B}} \circ \tilde{\mathcal{B}} = \tilde{\mathcal{B}}^2$ (6.60) and of (6.5) around the period-two solution. We find the matrix $\tilde{\mathcal{M}}_2 \tilde{\mathcal{M}}_1$ derived in Appendix 6.B of this chapter (6.66).

Despite the new dynamics of the rotational wedge, there still exists a region in the parameters space where the absolute values of the four eigenvalues of this matrix are less than 1. This region is depicted on Fig. 6.6 (solid lines) for five of the periodic solutions emphasized: $(n = 1, m = 1)$ and $(n = 3, m = 2)$ are two $(1 : 1)$ period-one solutions depicted in 6.6(a) and 6.6(d), respectively; while $(n = 2, m = 1 - (1 : 3))$, $(n = 3, m = 1 - (1 : 5))$ and $(n = 4, m = 2 - (3 : 5))$ are three period-two solutions depicted in 6.6(b), 6.6(c) and 6.6(e), respectively. Each of those stability regions is compared with the corresponding stability region of the double bouncing balls system (dotted lines), where $\sqrt{2}A$ has been replaced by AR^* in (6.43). These stability regions clearly overlap more accurately for a high coefficient of restitution e . This makes a physical sense: the largest the coefficient of restitution, the smallest the vibration, and thus the more acceptable the small angle assumption.

Note that each of the curves depicted in Fig. 6.5 crosses the $\theta = 45^\circ$ line (light dotted) with intervals corresponding to those drawn in Figs. 6.6(a) and 6.6(d), for the corresponding values of e .

6.5.3 Stability of period-two orbits in the actuated non-square wedge

The eigenvalues of the Jacobian matrix of the period-two orbits in the non-square wedge have not been computed. By analogy with the period-one orbits, we expect the conclusions obtained for the square wedge to persist in a range of values around $\theta = 45^\circ$.

The results of this section are summarized in the following proposition:

Proposition 6 *For a broad range of parameters (θ, e, A) , a sinusoidal actuation of the edges around a common fixed point (i.e. rotational actuation) achieves exponential stability of isolated periodic orbits. In particular, Fig. 6.5 illustrates the parametric stability region of two period-one orbits; and Fig. 6.6 illustrates the parametric stability region of five periodic orbits in the square configuration.*

6.6 Experimental results

This section describes an experimental validation of the stability results derived in the previous section for period-one and period-two orbits⁵.

The experimental validations have been realized with the Wiper robot, described in Chapter 5 and pictured in Fig. 5.2. The goal of the crude model we derived in Sections 6.2 to 6.5 is to capture the main effect of sensorless actuation of Wiper. The good matching between our experimental data and the theoretical predictions (see the rest of this section) is the best argument we can provide to validate the model, at least for the period-one motion.

6.6.1 Stabilization of period-one orbits

Our first experimental result is the stabilization of the first period-one orbit ($n = m = 1$) with rotational sinusoidal actuation of the edges. According to Fig. 6.5(a), two parameters are supposed to be crucial in the determination of the theoretical stability region: α , which is a geometrical parameter of the wedge and is therefore derived from the edges position, and the coefficient of restitution e , which on the contrary is difficult to estimate.

For a set of initial wedge aperture, we isolated experimentally the amplitude domains where the first period-one motion is stable. The experiments were conducted with a table inclined at an angle of 15° and a vibration frequency tuned in order to have a steady-state impact position R^* close to 0.7m (Equation (6.57)). Both parameters serve as tempo scaling factors and do not influence the stability regions.

Fig. 6.7 exhibits the experimental results obtained: all the tested conditions are marked by a point. The white and light gray zones are simply crude “contour” interpolations of the observed stability regions: all experimentally stable conditions belong to the white set. The light gray set contains the conditions which have been sorted as fragile but stable in the sense that either the period-one was stable, but only for a limited number of impacts (basically more than 10 but less than 30), or the puck described a complex motion confined around the period-one solution. The dark gray zone covers the experimental instability region, and the black zone excludes the experimental conditions leading to a risk of collision between the edges or with the table frame.

In order to compare the experimental results with the theoretical predictions, we superposed the theoretical stability region (solid lines) on Fig. 6.7, for $e = 0.5$. This region fits very well the experimental results and therefore validates the model and the mathematical analysis presented in the previous sections.

We note that Fig. 6.7 and Fig. 6.5(a) should be compared with caution. In Fig. 6.7, the parameter along the horizontal axis is the mean edge angle which is easily determined experimentally. However, the mean edge angle differs from the impact

⁵The interested reader will find two movies illustrating the stabilization of these periodic orbits on <http://ieeexplore.ieee.org> (Ronsse et al., 2006) or on the first author’s homepage.

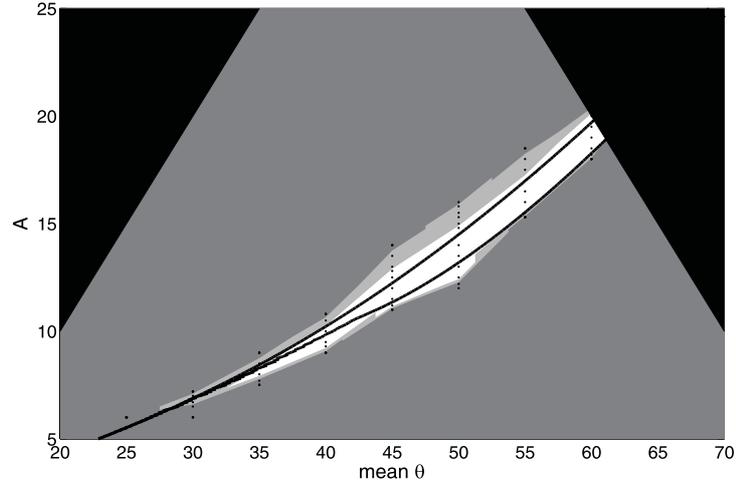


Figure 6.7: Experimental results on period-one stabilization: the points correspond to the tested conditions, the white zone surrounds the stability region, the light gray zone surrounds the “weak stability” region (see text), the dark gray zone depicts the instability region and the black zone excludes the non-secure zone. The results are compared with the theoretical prediction (solid lines) for $e = 0.5$.

angle, used as a parameter in Fig. 6.5. This explains why the theoretical parameter stability regions in Fig. 6.7 appears to be shifted and elongated compared to the corresponding region in Fig. 6.5(a).

6.6.2 Stabilization of period-two orbits

Experimental stabilization of period-two orbits turned out to be much more challenging. This is partly explained by the reduced stability regions of period-two orbits (see Fig. 6.6), but we also point out two additional reasons supporting that fact:

- The *basins of attraction* of the period-two orbits are much smaller than the period-one ones: a period-two orbit will be stabilized with a faster vibration frequency than a period-one orbit with an equivalent steady-state impact position (compare the ω obtained from Equations (6.51) and (6.57) for similar values of R^*). The period-two orbits are therefore more sensitive to the *phase* initial condition.
- The nonzero tangential impact velocity in period-two orbits causes the puck to *spin*. This phenomenon, not captured by the model, likely affects the stability properties and requires further investigations. See Spong (2001) for a study on the control of the spin dynamics.

We focused on the stabilization of the first period-two orbit ($n = 2, m = 1$). Fig. 6.6(b) shows that this periodic solution is not stable in the square configuration for

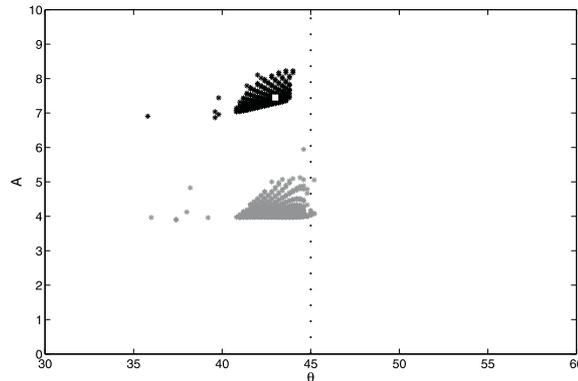


Figure 6.8: Stability region of the $(n = 2, m = 1)$ period-two orbit in the general Wiper for $e = 0.5$ (black points) and $e = 0.7$ (gray points).

the expected coefficient of restitution ($e \approx 0.5$). The Jacobian matrix of the period-two orbits in the general Wiper being hard to calculate, we preferred to simulate numerically the effect of small perturbations in our dynamical model, around the period-two limit cycle.

The points lying on Fig. 6.8 correspond to parameters values where a small perturbation of the period-two solution decreases through the impacts. These points therefore give an idea of the parametric stability region of that period-two orbit. Fig. 6.8 is consistent with Fig. 6.6(b)⁶: the parametric stability region has no intersection with the 45° black dotted line for e smaller than 0.7.

We successfully stabilized the period-two orbit in the lab during more than 50 impacts with a vibration amplitude A equal to 7.45° . The impact angular position has been visually estimated around 43° . This point is marked by a white square in Fig. 6.8 and clearly belongs to the numerically predicted stability region of the $(n = 2, m = 1)$ period-two orbit. This result opens the door to the stabilization of juggling patterns implying several pucks (two in this case), requiring only a temporal separation of the pucks equal to $2\pi/\omega$. In this way, Wiper would mimic the popular *shower* juggling pattern.

6.7 Conclusion

This chapter has established the experimental validation of the sensorless stabilization of bounce juggling patterns in the Wiper robot, described in Chapter 5. This result is also supported by a mathematical analysis of the modified stability prop-

⁶Let us recall that Figs 6.6(b) and 6.8 slice differently the stability region in the 3D parameter space (e, θ, A) : Fig. 6.6(b) slices the parameter space for $\theta = 45^\circ$ while Fig. 6.8 slices the parameter space for $e = 0.5$ and $e = 0.7$. These figures match at their intersections: $(\theta = 45^\circ, e = 0.5)$ and $(\theta = 45^\circ, e = 0.7)$.

erties: from periodic orbits which have been proved to be unstable, we derived an actuation law which stabilizes the same patterns.

Aiming at mimicking the popular *shower* juggling pattern, in Section 6.3 two particular periodic solutions of the wedge-billiard were studied: the period-two orbit, and its degenerate case, i.e. the period-one orbit. These periodic solutions have been proved to be either marginally stable or unstable for an uncontrolled elastic wedge-billiard. In Section 6.4, we derived the steady-state periodic orbits of a periodically actuated wedge-billiard. For the particular *square* wedge-billiard configuration, this generalizes the results on the bouncing ball dynamics (Holmes, 1982) to stabilize the wedge-billiard periodic orbits. A practical implementation of this stabilization requires a different configuration which has been modeled in Section 6.5. Within a single robot, several juggling patterns can therefore be stabilized through sensorless actuation, at least theoretically.

Finally, these theoretical predictions have been validated by several experimental results on *Wiper* robot, as explained in Section 6.6. Fig. 6.7 emphasized that our model under-estimated the stability region for the first period-one orbit. The dynamical properties that are not captured by our model seem to have a *stabilizing* effect on the period-one. Conversely, the period-two orbits were never stabilized during more than 30 seconds in our actual setup, while the model predicted reasonable parametric stability regions. We suppose that the *spin* effect, not captured in the map \mathcal{B} (6.60) while permanent in period-two trajectories, is highly destructive for stability; and that the sensorless control is not *robust* enough to cope with these unmodeled dynamics.

This chapter content has been published in Ronsse et al. (2006).

6.A Iteration map of the rotationally actuated wedge-billiard

Under the "small amplitude" assumption (Fig. 6.4), the impact rule \mathcal{I} captures the angular rotation. It is now given by

$$M(\mu[k]) \begin{pmatrix} V_r[k] \\ V_n[k] \end{pmatrix} = \begin{pmatrix} 1 & 0 \\ 0 & -e \end{pmatrix} M(\mu[k]) \begin{pmatrix} V_r^-[k] \\ V_n^-[k] \end{pmatrix} + \begin{pmatrix} 0 \\ \frac{1+e}{\alpha} R[k] \end{pmatrix} \dot{\mu}[k] \quad (6.59)$$

with $M(\mu)$ denoting the rotation matrix of the edge:

$$M(\mu) = \begin{pmatrix} \cos \mu & \alpha \sin \mu \\ -\frac{\sin \mu}{\alpha} & \cos \mu \end{pmatrix}$$

and $\mu[k] = \mu(t[k])$ ($\dot{\mu}[k] = \dot{\mu}(t[k])$) denoting the edge position (velocity) at impact time $t[k]$. The presence of $R[k]$ in (6.59) is due to a second important feature introduced by the rotational actuation: the energy exchange with the edges depends

now on the impact position $R[k]$. The map $\tilde{\mathcal{B}}$ of the rotational wedge is given by:

$$\begin{aligned} \begin{pmatrix} V_r[k+1] \\ V_n[k+1] \end{pmatrix} &= J(\mu[k+1]) \begin{pmatrix} |V_n|[k] - V_r[k] - |V_n^-|[k+1] \\ |V_n^-|[k+1] \cdot \text{sign}(V_n[k]) \end{pmatrix} \\ &\quad + \frac{1+e}{\alpha} \begin{pmatrix} -\alpha \sin \mu[k+1] R[k+1] \\ \cos \mu[k+1] R[k+1] \end{pmatrix} \dot{\mu}[k+1], \\ R[k+1] &= R[k] + \frac{1}{2g} V_r^2[k] + \frac{\alpha^2}{2g} V_n^2[k] \\ &\quad - \frac{1}{2g} (|V_n|[k] - V_r[k] - |V_n^-|[k+1])^2 - \frac{\alpha^2}{2g} |V_n^-|^2[k+1] \end{aligned} \quad (6.60)$$

with

$$J(\mu) = M(-\mu) \begin{pmatrix} 1 & 0 \\ 0 & -e \end{pmatrix} M(\mu) = \begin{pmatrix} \cos^2 \mu - e \sin^2 \mu & \frac{\alpha(1+e)}{2} \sin 2\mu \\ \frac{1+e}{2\alpha} \sin 2\mu & \sin^2 \mu - e \cos^2 \mu \end{pmatrix}. \quad (6.61)$$

The flight time is still given by (6.5).

The ‘‘small amplitude’’ assumption is done around the steady-state periodic orbits. They are therefore unchanged with respect to those which have been derived in Section 6.4. The actual actuation law is then:

$$\mu(t) = A (\sin \omega t - \sin (\{\phi^{r*}, \phi^{l*}\})), \quad (6.62)$$

providing $\mu^* = 0$.

6.B Jacobian matrices

The linearized Poincaré map of the period-one orbits of the rotational wedge-billiard is the matrix $\tilde{\mathcal{M}}$:

$$\begin{pmatrix} \frac{\omega}{g} \delta V_r[k+1] \\ \frac{\omega}{g} \delta |V_n|[k+1] \\ \frac{\omega^2}{\pi g} \delta R[k+1] \\ \omega t[k+1] \end{pmatrix} = \tilde{\mathcal{M}} \begin{pmatrix} \frac{\omega}{g} \delta V_r[k] \\ \frac{\omega}{g} \delta |V_n|[k] \\ \frac{\omega^2}{\pi g} \delta R[k] \\ \omega t[k] \end{pmatrix} \quad (6.63)$$

that is given by:

$$\tilde{\mathcal{M}} = \begin{pmatrix} \frac{1-4\alpha^2-\alpha^4}{(1+\alpha^2)^2} - 2\mu_{11} & \frac{4\alpha^2}{(1+\alpha^2)^2} - (\alpha^2-1)\mu_{11} & -\frac{4\alpha^2}{(2m-1)(1+\alpha^2)^2} - \frac{2}{(2m-1)}\mu_{11} & \frac{-\frac{1-e}{1+e}2(1+\alpha^2)}{4\alpha^3} \\ \frac{2e(\alpha^2-1)}{(1+\alpha^2)^2} + (2m-1)\mu_{21} & \frac{e(\alpha^2-1)^2}{(1+\alpha^2)^2} + \mu_{22} & \frac{4e\alpha^2}{(2m-1)(1+\alpha^2)^2} + \mu_{21} & \frac{-(1-e)(1+\alpha^2)(2m-1)^2\Gamma'}{4\alpha^3} \\ \frac{(2m-1)(1-\alpha^2)}{1+\alpha^2} & \frac{2(2m-1)\alpha^2}{1+\alpha^2} & \frac{1-\alpha^2}{1+\alpha^2} & 0 \\ \frac{4\alpha^2}{(1+\alpha^2)^2} & \frac{2\alpha^2(\alpha^2-1)}{(1+\alpha^2)^2} & \frac{4\alpha^2}{(2m-1)(1+\alpha^2)^2} & 1 \end{pmatrix} \quad (6.64)$$

with

$$\begin{aligned}\mu_{11} &= \frac{1-e}{1+e} \frac{4\alpha^2}{1+\alpha^2}, \\ \mu_{21} &= (1-e) \left(\frac{-(2m-1)\Gamma'}{\alpha(1+\alpha^2)} + \frac{2(1-\alpha^2)}{(2m-1)(1+\alpha^2)} \right), \\ \mu_{22} &= (1-e) \left(\frac{(1-\alpha^2)(2m-1)^2\Gamma'}{2\alpha(1+\alpha^2)} + \frac{4\alpha^2}{1+\alpha^2} \right), \\ \Gamma' &= \sqrt{\left(\left(\frac{1+e}{1-e} \right) A\pi^2 \right)^2 - \left(\frac{2\pi\alpha}{2m-1} \right)^2}.\end{aligned}$$

Note that the determinant of $\tilde{\mathcal{M}}$ is equal to

$$|\tilde{\mathcal{M}}| = \frac{e((1-e)(3\alpha^2-1) + 2(1+\alpha^2))}{(1+e)(1+\alpha^2)} \quad (6.65)$$

which is equal to 1 in the elastic case ($e = 1$), $\forall \alpha$. This illustrates that the sinusoidal input does not achieve exponential stability of the period-one orbit in the *elastic* wedge, because all the eigenvalues of this matrix cannot be < 1 in that case.

The linearized Poincaré map of the periodic orbits of the rotational *square* wedgebilliard is the matrix $\tilde{\mathcal{M}}_2\tilde{\mathcal{M}}_1$:

$$\begin{aligned}\begin{pmatrix} \frac{\omega}{g}\delta V_r[k+2] \\ \frac{\omega}{g}\delta|V_n|[k+2] \\ \frac{\omega^2}{\pi g}\delta R[k+2] \\ \omega t[k+2] \end{pmatrix} &= \tilde{\mathcal{M}}_2 \begin{pmatrix} \frac{\omega}{g}\delta V_r[k+1] \\ \frac{\omega}{g}\delta|V_n|[k+1] \\ \frac{\omega^2}{\pi g}\delta R[k+1] \\ \omega t[k+1] \end{pmatrix} \\ &= \tilde{\mathcal{M}}_2\tilde{\mathcal{M}}_1 \begin{pmatrix} \frac{\omega}{g}\delta V_r[k] \\ \frac{\omega}{g}\delta|V_n|[k] \\ \frac{\omega^2}{\pi g}\delta R[k] \\ \omega t[k] \end{pmatrix},\end{aligned} \quad (6.66)$$

$\tilde{\mathcal{M}}_1$ being given by

$$\tilde{\mathcal{M}}_1 = \begin{pmatrix} -\frac{2\eta}{\eta+\nu} - \frac{1-e}{1+e} \frac{2(\eta+\nu)}{\nu} & 1 & -\frac{2}{\eta+\nu} - \frac{1-e}{1+e} \frac{2(\eta+\nu)}{\eta\nu} & -\frac{1-e}{1+e} \frac{(\eta+\nu)^2}{\eta\nu} \\ \frac{-e(\nu-\eta)}{\eta+\nu} + (1-e)\left(\frac{2\eta\beta}{\eta+\nu} + \frac{\nu-\eta}{\nu}\right) & (1-e)\frac{\eta+\nu}{\nu} & \frac{2e}{\eta+\nu} + (1-e)\left(\frac{2\beta}{\eta+\nu} + \frac{\nu-\eta}{\nu}\right) & (1-e)\beta \\ \frac{(\nu-\eta)\eta}{\eta+\nu} & \eta & \frac{\nu-\eta}{\eta+\nu} & 0 \\ \frac{2\eta}{\eta+\nu} & 0 & \frac{2}{\eta+\nu} & 1 \end{pmatrix} \quad (6.67)$$

with

$$\beta = \frac{(\nu+\eta)(\eta-\nu)}{\nu\eta} - \frac{\nu\eta}{2} \sqrt{\left(\left(\frac{1+e}{1-e} \right) A\pi^2 \right)^2 - \left(\frac{2n\pi}{(2m-1)(2n-2m+1)} \right)^2}.$$

In (6.67), $\eta = 2m - 1$ and $\nu = 2n - 2m + 1$ depends on the steady-state flight times. $\tilde{\mathcal{M}}_2$ has exactly the same structure as $\tilde{\mathcal{M}}_1$ with an exchange between ν and η .

Chapter 7

Rhythmic Feedback Control of Wiper in Blindness

Music and rhythm find their way
into the secret places of the soul.

Plato

7.1 Introduction

In Chapter 6, we provided experimental validation of a *sensorless* control law (a purely sinusoidal actuation of the juggler arms) that robustly stabilizes the period-one orbit (Fig. 6.2(a)), but failed to robustly stabilize the period-two orbit (Fig. 6.2(b)) for more than a few seconds. The aim of the present chapter is consequently to investigate whether adding feedback to the loop may help to robustly stabilize the period-two. The control law proposed in the present chapter can be interpreted as a phase and amplitude modulation of the sinusoidal control law, based on the feedback information provided by the impact times. The technical material of this hybrid control strategy (discrete feedback for continuous-time actuation) has been presented for the bouncing ball in Chapter 4.

In Section 7.2, we briefly generalize the bouncing ball strategy to the special configuration of Wiper that reduces to two coupled bouncing balls (see Section 5.2). Along each decoupled dynamics — i.e. each edge — the control algorithm is exactly similar to (4.15). Zero static error between the frequencies of these dynamics¹ is ensured through the integral control described in Section 4.3.3.

The more accurate model of Wiper, based on the wedge-billiard described in Section 6.2 and derived in Section 6.A, is considered in Section 7.3. Discrete feedback control based on impact times is realized in a similar manner to the bouncing ball control: a model-based *observer* is synchronized with the actual dynamics through

¹Such that the phase-lag between them is fixed and invariant.

the impact times. Sustained experimental stabilization of the period-two orbit is also provided. Once again, let us emphasize that our experimental validation suffers from large discrepancies between the wedge-billiard model and the real setup dynamics. For example, the Newton impact model does not capture the *spin* effect of the puck, clearly visible in the period-two pattern. However, the focus of our approach is precisely to validate the *robustness* of a model-based control scheme in a real setup in spite of the many discrepancies between the model and the experimental setup. The robustness has been quantified in Chapter 4 via the *acceleration* of the actuator at impact. Obviously, we tuned our controller to match with the optimal acceleration (4.21) locally around impacts, such that the closed-loop robustness there obtained is large enough to compensate for the model discrepancies.

7.2 Feedback Control of the Simplified Wedge-Billiard

First, let us generalize the robust feedback control strategy discussed for the bouncing ball (Chapter 4) to the special configuration of the wedge-billiard, presented in Section 5.2 and pictured in Fig. 5.4. The two edges are then assumed to remain aligned with the two orthogonal axes of a fixed reference frame (this requires parallel actuation of the edges, in contrast with the rotational actuation of the Wiper robot, see Chapter 5). In this special configuration, the resulting dynamics nicely decouple: the dynamics along each axis are the dynamics of a one-dimensional bouncing ball (see Section 5.2 and Sepulchre and Gerard, 2003).

A period-two orbit of the planar juggler projects onto a period-one orbit along each axis (see Fig. 5.4). The phase shift Φ^* (not represented in the figure) between the x -orbit and the y -orbit controls the shape (that is the apex difference between the low toss and the high toss) of the period-two planar orbit, since the impact times obey:

$$\begin{aligned} t[k_x] &= \frac{2\pi k_x}{\omega}, \\ t[k_y] &= \frac{2\pi k_y + \Phi^*}{\omega} \end{aligned} \tag{7.1}$$

where k_x and k_y denote the k^{th} impact along the x and y axes, respectively; and ω still denotes the cycle frequency.

Based upon this analogy, the stabilization of the period-two planar orbit of the juggler reduces to the stabilization of a period-one orbit in two independent 1D bouncing balls. As a consequence, the blind feedback control — i.e. based only on impact *times* measurement — introduced in Chapter 4 directly applies for the simplified wedge billiard. Any period-one or period-two orbit can be stabilized through proper control of the *phase* between the two decoupled dynamics. Phase

control with no static error can be achieved with the integral control described in Section 4.3.3.

7.3 Experimental Validation with the Wiper Robot

7.3.1 Experimental setup

In the *Wiper* robot, both edges are actuated around their fixed point, like in the general wedge-billiard model (see Sections 6.2 and 6.A, and Fig. 6.1). Fig. 5.2 is a picture of *Wiper*.

The actuated metallic edges are controlled by two independent motors, which are mounted close to each other (about 10cm) since the model assumes that both edges rotate around the same point. Both motors are controlled with a real-time computer running with XPCTARGET (THE MATHWORKS[®]). Impact times are detected by two accelerometers mounted at the top of the edges to record the high-frequency vibrations generated by the impacts. More technical details can be found in Appendix B.1.

7.3.2 From the bouncing ball model to the wiper model

The double bouncing ball model presented in Section 7.2 is obviously a crude model of the Wiper robot: it neglects the coupling resulting from a rotational actuation of the edges, and a steady-state aperture possibly different from $2\theta = 90^\circ$. However, a more accurate model of Wiper was derived in Chapter 6², under the small angle assumption (Fig. 6.4). Despite this assumption, the model predicted a parametric stability region of the period-one orbit which is in excellent agreement with experiments (Fig. 6.7).

The Wiper model (see 6.60) has three state variables — V_r , V_n and R — and two input variables — $\mu[k]$ and $\dot{\mu}[k]$, denoting the impacted edge position w.r.t. θ (left) or $-\theta$ (right), and velocity at impact k — see Figure 6.1. Injecting the flight time (6.5) into the state-space model (6.60) gives the following state-space model, depending explicitly on the time difference:

$$\begin{aligned} \begin{pmatrix} V_r[k+1] \\ V_n[k+1] \end{pmatrix} &= J(\mu[k+1]) \begin{pmatrix} \frac{1-\alpha^2}{1+\alpha^2} V_r[k] + \frac{2\alpha^2}{1+\alpha^2} |V_n[k] - g(t[k+1] - t[k])| \\ \left(\frac{-2}{1+\alpha^2} V_r[k] + \frac{1-\alpha^2}{1+\alpha^2} |V_n[k] + g(t[k+1] - t[k])| \right) \text{sign}(V_n[k]) \end{pmatrix} \\ &\quad + \frac{1+e}{\alpha} \begin{pmatrix} -\alpha \sin \mu[k+1] R[k+1] \\ \cos \mu[k+1] R[k+1] \end{pmatrix} \dot{\mu}[k+1], \quad (7.2) \\ R[k+1] &= \frac{1+\alpha^2}{4g} (V_n^2[k] - (|V_n[k] - g(t[k+1] - t[k])|)^2) \end{aligned}$$

with the rotation matrix $J(\mu)$ given by (6.61).

²See also Lehtihet and Miller (1986); Sepulchre and Gerard (2003); Gerard and Sepulchre (2004, 2005); Chang et al. (2005); Ronsse et al. (2005, 2006)

The output control law (4.15) derived in Chapter 4 for the bouncing ball is adapted for Wiper as follows:

- (i) **State observer:** The deadbeat velocity observer (4.14) is replaced by a copy of Wiper's dynamics (7.2), in which we substitute estimated variables \hat{V}_r , \hat{V}_n , \hat{R} to the state variables V_r , V_n , R ; and where the measured flight time is injected from the measurements (see (6.5)) to synchronize the observer with the actual state variables:

$$\begin{aligned} \begin{pmatrix} \hat{V}_r[k] \\ \hat{V}_n[k] \end{pmatrix} &= J(\mu[k]) \begin{pmatrix} \frac{1-\alpha^2}{1+\alpha^2} \hat{V}_r[k-1] + \frac{2\alpha^2}{1+\alpha^2} |\hat{V}_n|[k-1] - g(t[k] - t[k-1]) \\ \left(\frac{-2}{1+\alpha^2} \hat{V}_r[k-1] + \frac{1-\alpha^2}{1+\alpha^2} |\hat{V}_n|[k-1] + g(t[k] - t[k-1]) \right) \text{sign}(\hat{V}_n[k-1]) \end{pmatrix} \\ &\quad + \frac{1+e}{\alpha} \begin{pmatrix} -\alpha \sin \mu[k] \\ \cos \mu[k] \end{pmatrix} \hat{R}[k] \dot{\mu}[k], \quad (7.3) \\ \hat{R}[k] &= \frac{1+\alpha^2}{4g} \left(\hat{V}_n^2[k-1] - \left(|\hat{V}_n|[k-1] - g(t[k] - t[k-1]) \right)^2 \right). \end{aligned}$$

We suppose that the inputs $\mu[k]$ and $\dot{\mu}[k]$ are directly available from the measured trajectory of the edges $\mu(t)$ and the impact times. This observer displays excellent convergence properties in simulations.

- (ii) **Tracking controller:** The tracking controller (4.4), (4.5), (4.6) is adapted to the Wiper model. The desired next impact time is estimated on the basis of (6.5):

$$t_u[k+1] = t[k] + \frac{1}{g} \left(|\hat{V}_n^-|[k+1] + \frac{(\alpha^2 - 1)|\hat{V}_n|[k] + 2\hat{V}_r[k]}{1 + \alpha^2} \right). \quad (7.4)$$

The position reference is simply $\mu_\rho^* = 0$ since the objective is to stabilize periodic orbits of Wiper whose impacts occur at $\pm\theta$. The ‘‘local’’ velocity of the edge, at the impact point, depends obviously both on the velocity input $\dot{\mu}_\rho[k+1]$ and the impact radial position $R[k+1]$. Defining $\dot{S} = \dot{s}/\sin\theta$, where \dot{s} is equal to this local velocity, one has therefore $\dot{\mu}[k+1] = \alpha\dot{S}[k+1]/R[k+1]$. This ‘‘local’’ edge velocity derives from the impact rule:

$$\dot{S}_u[k+1] = \frac{|V_n|_\rho[k+1] - e|\hat{V}_n^-|[k+1]}{1+e} \quad (7.5)$$

where the reference post-impact velocity $|V_n|_\rho[k+1]$ is equal to:

$$|V_n|_\rho[k+1] = \frac{g}{2\alpha^2} (t_\rho[k+3] - t_u[k+1]) - \frac{g(1-\alpha^2)}{2\alpha^2} (t_\rho[k+2] - t_u[k+1]) \quad (7.6)$$

and is computed from the impact times reference. In the square configuration ($\alpha = 1$), this equation becomes similar to (4.25). In that case, the wedgebilliard dynamics decouple into two bouncing ball dynamics, one along each edge, as mentioned in Sections 6.4 and 7.2.

Finally, the estimate of the radial position is obtained from the estimated state variables (see (6.60)):

$$\begin{aligned}
 R_u[k+1] &= \hat{R}[k] + \frac{1}{2g}\hat{V}_r^2[k] + \frac{\alpha^2}{2g}\hat{V}_n^2[k] \\
 &\quad - \frac{1}{2g}(|\hat{V}_n|[k] - \hat{V}_r[k] - |\hat{V}_n^-|[k+1])^2 - \frac{\alpha^2}{2g}(\hat{V}_n^-)^2[k+1].
 \end{aligned} \tag{7.7}$$

- (iii) **Continuous-time actuation:** The piecewise quadratic control law (4.7) generates obviously sharp position transients at impact, even in steady-state. In order to smoothen the actuation, we chose to design a closed-loop control that reaches the desired position and velocity at impact, but that is smooth in steady-state. This modification helped to prevent false impact detection and motor damages in real experiments. This was realized by adapting the open-loop sinusoidal law (3.7) to take feedback into account, for amplitude and phase tuning:

$$\begin{aligned}
 \mu(t) &= A \left(\kappa_{FB} \frac{\dot{S}_u[k+1]}{\dot{S}^*} \frac{R^*}{R_u[k+1]} + (1 - \kappa_{FB}) \right) \\
 &\quad (\sin(\omega(t - t_u[k+1]) + \phi^*) - \sin \phi^*) \text{sign}(\bullet)
 \end{aligned} \tag{7.8}$$

where $\text{sign}(\bullet) = 1$ for the left arm and -1 for the right one. The steady-state phase ϕ^* is given by (6.58), with $m = 1$. At time $t = t_u[k+1]$, we obtain the desired impact position ($\mu(t_u[k+1]) = 0$) and the impact velocity:

$$\dot{\mu}(t_u[k+1]) = \kappa_{FB}\dot{\mu}_u[k+1] + (1 - \kappa_{FB})\dot{\mu}^*. \tag{7.9}$$

The amplitude A is tuned to match, in steady-state, the local impact acceleration $\dot{S}(t_u[k+1])$ with the optimal value defined by (4.21). This acceleration is equal to $\ddot{\mu}(t)$, with μ defined in (7.8), and is obviously negative in steady-state since $0^\circ < \phi^* < 90^\circ$ (see (6.58)).

The feedback gain $0 \leq \kappa_{FB} \leq 1$ is tuned to achieve the best possible trade-off between the stabilizing performance of the observer-based output feedback controller ($\kappa_{FB} = 1$) and its sensitivity to the model uncertainty.

7.3.3 Results

The experimental challenge was to stabilize period-one and, more importantly, period-two orbits of Wiper. This section describes the results we obtained for an aperture of $\theta = 40^\circ$. The parameters of the actuation law (7.8) were tuned to $A = 9^\circ$ and $\omega = 1.1\pi\text{rad/s}$, while the feedback gain has been empirically tuned to $\kappa_{FB} = 0.4$. The coefficient of restitution was estimated about $e = 0.7^3$. The

³The puck used in that experiment was then more elastic than the puck used in the experiments reported in Chapter 6.

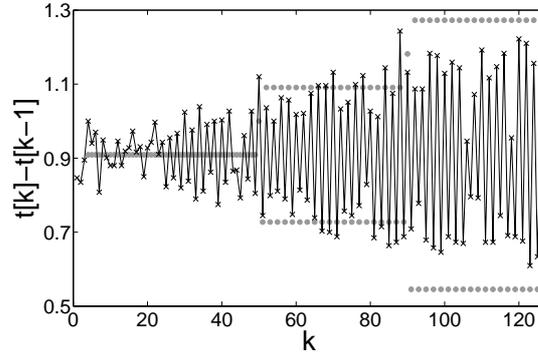


Figure 7.1: The flight times between two successive impacts (detected by the accelerometers) are depicted with the black crosses. The gray circles denote the reference flight times, as defined by $t_\rho[k] - t_\rho[k-1]$. The first part of the trajectory (about $0 < k < 50$) displays the closed-loop stabilization of a period-one motion: $\Delta t^* = \pi/\omega \simeq 0.91\text{s}$. The second part of the trajectory (about $50 < k < 90$) displays the closed-loop stabilization of a “small” period-two motion, i.e. when the reference flight time alternates between $0.9\Delta t^*$ and $1.1\Delta t^*$. The last part of the trajectory (about $k > 90$) displays the closed-loop stabilization of a larger period-two motion: the flight time reference alternates between $0.8\Delta t^*$ and $1.2\Delta t^*$.

experiment initialization has been realized by proper throwing of the puck while the edges were sinusoidally actuated. Initialization then exploited the good basin of attraction of the period-one orbit in open-loop (see Chapter 6). Closed-loop control was switched on after convergence of the observer. Then the reference was switched from period-one to a period-two with small aperture, and finally to a period-two with large aperture. A movie is available⁴ to illustrate this experiment.

The flight times between two successive impacts, as detected by the accelerometers, are depicted in Figure 7.1, w.r.t. the reference. The first part of the trajectory (till $k \simeq 50$) displays the closed-loop stabilization of the period-one motion, i.e. when the target flight time is always equal to $\Delta t^* = \pi/\omega \simeq 0.91\text{s}$. The second part of the trajectory (about $50 < k < 90$) displays the closed-loop stabilization of a “small” period-two motion, i.e. when the flight time alternates between $0.9\Delta t^*$ and $1.1\Delta t^*$. The last part of the trajectory (about $k > 90$) displays the closed-loop stabilization of a larger period-two motion: the flight time reference alternates between $0.8\Delta t^*$ and $1.2\Delta t^*$. The mismatch between the reference and the actual trajectory increases for the “large” period-two orbit. This can be explained as follows: the steady-state velocity \dot{S}^* and impact position R^* depend on the reference pattern (see (6.35), (6.48), (6.49) and (6.51)). However, for the sake of simplicity, we let them equal to the steady-state values of the period-one motion ((6.55) and (6.57))

⁴<http://ieeexplore.ieee.org> (Ronsse et al., 2007a) or on the first author’s homepage. The material is 15.9 MB in size.

in the control law (7.8), regardless of the reference. This introduces a steady-state error w.r.t. the two reference flight times when they significantly differ from each other, when $\alpha \neq 1$.

7.4 Conclusion

The successful experimental validation of period-two orbits in Wiper, with robust control of the trajectory, contrasts with previous results obtained with a sensorless control (Chapter 6 and Ronsse et al., 2006). In that case, indeed, the period-two orbits collapsed after a few seconds. This was due to both the small model-predicted basins of attraction, and the lack of robustness with respect to the real dynamics variability. This illustrates that feedback is required to stabilize the shower pattern with Wiper. A first “naive” attempt to add feedback in Wiper’s dynamics would be to generalize the mirror law algorithms developed by Buehler, Koditschek and Kindlmann (1988, 1990, 1994) for the bouncing ball. However, this law would have sharply contrasted with the encouraging results of the sensorless strategy discussed in Chapter 6, which had already satisfying behavior for the simplest periodic orbit, i.e. the period-one. Indeed, the mirror law requires permanent tracking of the puck in order to measure its continuous-time position and energy (see (3.18)). Such permanent tracking would necessitate a complex sensor design. Preliminary results of the mirror law implementation with a video-based tracker in Wiper can be found in Manuel Gerard’s PhD thesis (Gerard, 2005).

In contrast, we discussed in Chapter 4 an alternative feedback source, based on limited sensing demand: the *impact times*. Indeed, the times of impact can be cheaply measured by accelerometers on the edges (as we did), or microphones, which require limited signal processing to extract the relevant information (see Appendix B.1). This limited sensing led to *robust* control, which compensated for the discrepancies between the model-based observer and the actual dynamics. In agreement with the robustness analysis of Section 4.3, we illustrated that sustained stabilization in the lab required a proper (negative) tuning of the impact acceleration and could never be achieved with positive acceleration.

The present chapter material has been published in Ronsse et al. (2007a).

Chapter 8

Human Control of Bimanual Impact Juggling: Trade-off between Performance and Robustness

He who controls the present,
controls the past. He who
controls the past, controls the
future.

George Orwell

8.1 Introduction

As summarized in Section 1.1, the computational approach is widely used in the motor control literature, providing a unified framework to study motor planning, control, estimation, prediction and learning (see also Jordan and Wolpert, 1999; Wolpert and Ghahramani, 2000). Such concepts are inherited from well established engineering principles from system-theory. In particular, it is now accepted that the planning of a desired limb trajectory can be successfully quantified with the help of *optimal control* theory (Bryson and Ho, 1969), by measuring to what extent the planned trajectory maximizes a given cost function (e.g. the non-smoothness, the energy expenditure, the movement duration, etc. . . , see Todorov and Jordan, 2002; Scott, 2004; Todorov, 2004, 2006). Through the present chapter, we aim at shedding light on a particular computational aspect that has not been studied previously in the motor control literature, while ubiquitous in control theory (see Section A.3): the trade-off between *performance* and *robustness*. Performance quantifies how well a closed-loop system performs (cost function); and robustness measures the ability to maintain the closed-loop stability despite uncertainties and perturbations. For example, robustness can be quantified by the noise level and the computational

delay which are tolerated before losing stability. Importantly, performance and robustness compete against each other, in the sense that they cannot be arbitrarily improved simultaneously. Many textbooks in control theory (see e.g. Franklin et al., 2005; Astrom and Murray, 2005) provide caveat examples of closed-loop designs performing excellently, but fragile to a slight change in any parameter. In this chapter, we discuss how this trade-off may be influential for trajectory planning in the control of rhythmic movements by human subjects.

The task studied in this chapter is the period-one of Wiper (Fig. 5.1(c)) and requires the stabilization of the puck through periodic impacts with the metallic arms which are actuated by the subject. The control strategy must be therefore robust in the stabilization of both the coordinated bimanual movement and the puck trajectory.

Since both the setup and the task are similar, interesting analogies can be established with the robotics experiments described in Chapters 6 and 7. The experiment is thus suited to address questions in robotics and motor control in parallel (Schaal and Schweighofer, 2005).

We investigate how the trade-off between performance and robustness is influenced in the juggling task by the suppression of a major source of sensory inflows, i.e. visual feedback, and by the time available to plan and execute the movement, i.e. the task tempo. Both the absence of visual feedback and the increase in the tempo should increase the uncertainties in the task. In this study, we investigate to what extent the maintenance of the robustness at a satisfactory level, despite uncertainties, will affect the performance of the subjects. Direct analogies are established with the robotics experiments, in which we investigated mainly the *minimum* sensory feedback necessary for robust stabilization of the juggled puck.

8.2 Methods

8.2.1 Task and experimental setup description

Nine healthy human subjects (4 female, 5 male, 23-28 years old) volunteered to perform the task. Two subjects were left-handed, seven were right-handed. One subject is the author of the thesis while the others were naive regarding the goals of the experiment. They provided informed written consent, and reported no history of neurological or musculoskeletal disorder. All had normal vision, either natural or corrected. All the procedures conducted were approved by the local ethics committee, in compliance with the Helsinki declaration.

A general picture of the experimental setup has been described in Fig. 5.3: the subjects stood in front of the air-hockey table (tilted 12.5° w.r.t. the ground), and actuated two independent metallic arms that were free to rotate on the table about their lower extremity. The two points of rotation were mounted close to each other, on the table frame. The subjects had to rotate the metallic arms in order to impact

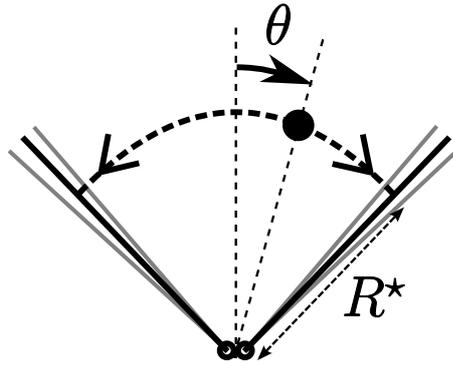


Figure 8.1: Sketch of the experimental setup: the two solid black lines represent the metallic arms, rotating around their respective lower extremity. The puck is depicted by a black disk. The steady-state period-one parabola is the dashed line. θ denotes the angular position of the puck in the table system of coordinate, with respect to the vertical. R^* denotes the radial position of the puck at impact in the same system of coordinate.

a plastic puck (72mm diameter) back and forth. The goal was to maintain the so-called “period-one” pattern (see Figs. 5.1(c) and 8.1).

The task tempo was assigned by a metronome, that beeped either at $\Delta t^* = 400\text{ms}$, $\Delta t^* = 600\text{ms}$, $\Delta t^* = 800\text{ms}$ or $\Delta t^* = 1000\text{ms}$ intervals. The subjects were instructed to synchronize the impacts with the metronome beeps, such that the tempo Δt^* represents in fact the half-period of the periodic pattern. At each tempo, the subjects were asked to perform the juggling task for 2 minutes in normal visual condition, followed by 2 minutes with closed eyes. These 2×2 minutes blocks were repeated 6 times per session: one at $\Delta t^* = 400\text{ms}$, one at $\Delta t^* = 600\text{ms}$, two at $\Delta t^* = 800\text{ms}$, and two at $\Delta t^* = 1000\text{ms}$ ¹. The succession of tempi was randomly generated before each session. Each of the nine subjects completed 4 sessions. They were authorized to acquaint themselves with the task during 15 minutes before the first session.

The coefficient of restitution of the arms was estimated at $e = 0.48$. It means that the puck lost roughly $(1 - 0.48^2) \times 100 = 77\%$ of its energy at each impact. During the task completion, this energy was obviously restored by impacting the puck with a non-zero arm velocity.

The eye movements (horizontal and vertical) were recorded at 200Hz using a Chronos eye tracker (CHRONOS VISION GmbH, Berlin, Germany), which is based on high-frame rate CMOS sensors (Clarke et al., 2002). The eye tracker has a

¹This generated $120\text{sec}/0.4\text{sec} = 300$ beeps at $\Delta t^* = 400\text{ms}$; $120\text{sec}/0.6\text{sec} = 200$ beeps at $\Delta t^* = 600\text{ms}$; $2 \times 120\text{sec}/0.8\text{sec} = 300$ beeps at $\Delta t^* = 800\text{ms}$; and $2 \times 120\text{sec}/1\text{sec} = 240$ beeps at $\Delta t^* = 1000\text{ms}$ per subject \times visual feedback condition \times session, and then potentially the same number of impacts. Obviously, the subjects could not maintain steady-state juggling cycles during the whole blocks.

resolution superior to 0.1° in a measurement range of $\pm 40^\circ$. The eye tracker was also used to monitor the correct closing of the eyes when instructed. The 3D positions of infrared light-emitting diodes (IREDs) on the Chronos helmet, at the top of the arms and at the center of the puck were measured using an OptoTrak 3020 system (Northern Digital, Ontario, Canada). The OptoTrak was mounted on the ceiling about 3m in front of the subject. The positions of the IREDs were rotated to be expressed in a coordinate system with two axes parallel to the air-hockey table frame, and the third one pointing upward, and centered between the rotation points of the arms. The position of each IRED was sampled at 200Hz with a resolution of about 0.1mm within the working environment. The Chronos eye tracker and the OptoTrak were synchronized by an external computer (see Section B.2). Each block was executed over a period of 120s, and its data recorded in separate files. A complete description of the method used to compute the gaze orientation in 3D space, from the head IREDs and the Chronos movies, is described in Appendix C. The calibration protocol was executed at the beginning, middle and end of each session.

8.2.2 Data analysis

Digital processing of the raw data was performed with MATLAB (the MathWorks Inc., Natick, MA, USA). The eye and head IREDs position signals were filtered at 48Hz by a zero-phased digital filter (autoregressive, forward and backward). The arms IREDs were filtered at 10Hz by a similar filter. The puck trajectories were not filtered because the puck velocity was discontinuous at impacts: these discontinuity points identified the impact times. Eyes, head, arms and puck velocities and accelerations were computed from position signals by means of a central difference algorithm.

This study does not focus on transient or initialization phases. We kept therefore only the steady-state cycles in the database. Abnormal flights and impacts were also excluded from our analysis. Furthermore, we kept only the trials containing at least four successive impacts, and we removed the two first and the last one. When the subjects were instructed to close their eyes during steady-state cycles, we kept only the sequences during which the eyes were detected to be closed for more than 80% of the flight time. Other impacts were not analyzed. After this complete selection, the database contained 31,997 pairs of flights and impacts.

To assess whether the arms maintained a pose interval between two successive impacting movements during the task execution, we computed the activity period of the impacting arm as the percentage of the movement cycle (i.e. two consecutive flight times) during which the arm was moving, according to a mixed velocity ($v_{arm} > 5^\circ/s$) OR acceleration ($a_{arm} > 20^\circ/s^2$) criterion. For instance, if this ratio equals 90% for a cycle, it means that the arm velocity and acceleration were simultaneously below $5^\circ/s$ and $20^\circ/s^2$ respectively, — i.e. the arm was at rest — during 10% of the flight time. We also computed a synchronization index of the arms during flight

time as the averaged phase relationship between both arms, via the non-normalized correlation coefficient of the two velocity signals v_l (left arm) and v_r (right arm):

$$\cos \hat{\phi} = \frac{\widehat{v_r \cdot v_l}}{\sqrt{\widehat{|v_r|^2} \cdot \widehat{|v_l|^2}}} = \frac{\int_{\Delta t} v_r(t) v_l(t) dt}{\sqrt{\int_{\Delta t} v_r^2(t) dt \cdot \int_{\Delta t} v_l^2(t) dt}}$$

Synchronized in-phase or anti-phase movements corresponded to $\cos \hat{\phi} = -1$ and $\cos \hat{\phi} = 1$, respectively; while decoupled movements, i.e. when the arms moved one after the other with 90° phase-lag, corresponded to $\cos \hat{\phi} = 0$.

8.2.3 Statistics

Some analyzed variables were not normally distributed, either because they were sharply bounded, or because they had a bimodal distribution (normality was tested using the Lilliefors test). The analysis of variance of these variables was performed through Kruskal-Wallis ANOVAs, a dedicated tool for non-normal populations. The tests were made for the two visual feedback conditions, with the tempo Δt^* as factor; and for the four tempi Δt^* , with the visual feedback condition as a factor (all significant p -levels at 0.05).

For normal populations, the dependent variables means were analyzed using classical factorial analysis of variance (ANOVAs). The first within-subject factor was the tempo Δt^* and the second factor was the visual feedback condition. All ANOVAs were evaluated as significant for p -levels of 0.05. Main and interaction effects were further analyzed using Tukey's HSD post hoc tests ($p < 0.05$).

Linear regressions were processed using a standard regression algorithm (least square minimization). The slopes of the regressions were compared using the "homogeneity of slopes" module. This module is an extension of the analysis of covariance (ANCOVA) module, that is used when the continuous and categorical predictors interact in influencing the responses.

All statistical analysis was completed using Statistica (StatSoft Inc., OK, USA).

8.3 Results

8.3.1 Tempo dictated the angular position at impact

On average, the subjects performed the task very well at a tempo close to the reference Δt^* , both with eyes open or closed (see Table 8.1).

The steady-state juggled pattern is a parabola, with radial impact position R^* and angular aperture $2\theta^*$ (see Fig. 8.1). In order to increase (or decrease) the tempo, the subjects had the choice to reduce (or augment) either the radial position or the angular aperture (see (6.11) and (6.12)). In the data, we reported no significant variation of radial impact position. In contrast, the *angular* position of impact

Table 8.1: Averaged across all subjects, the table reports the mean flight time between two successive impacts (Δt) and the angular position at impact (θ_{impact}), depending on the tempo and visual feedback (vFB) condition (w.out=without). Factorial ANOVAs revealed that both the puck flight time between two successive impacts and the impact position depended on the tempo. Tuckey’s HSD post hoc analysis revealed that all the pairwise comparisons were always significant, for a given visual feedback condition.

tempo Δt^*	400ms		600ms		800ms		1000ms	
vFB	with	w.out	with	w.out	with	w.out	with	w.out
Δt [ms]	412	406	600	599	806	805	997	977
θ_{impact} [deg]	21.6	22.1	31.9	32.2	42.5	42.3	51.5	49.9

strongly depended on the tempo (see Table 8.1), both with eyes open or closed. Thus the subjects only adapted the aperture between the arms to juggle the different tempi: the faster the tempo, the smaller the aperture (the range of θ in Fig. 8.1).

8.3.2 Intermittent or rhythmic behavior

In normal visual conditions, two stereotyped actuation profiles were adopted by the subjects. At slow tempi (e.g. $\Delta t^* = 1000\text{ms}$, see Fig. 8.2, left), the arms were actuated intermittently, one after the other, to impact the puck. In this case, the actuation was a train of discrete arm movements. In contrast, faster tempi (e.g. $\Delta t^* = 400\text{ms}$, see Fig. 8.2, right) were characterized by rhythmic and sustained actuation of both arms in synchrony (see velocity profiles in Fig. 8.2, right).

As illustrated in these typical examples, the synchronization between the arms strongly depended on the task tempo. Fig. 8.3 represents the histogram distributions of the synchronization index $\cos \hat{\phi}$, for the four tested tempi. At the fastest tempo ($\Delta t^* = 400\text{ms}$), strongly synchronized movements were executed (sync index close to 1). As the tempo decreased, a second mode emerged, characterized by completely decoupled actuation (sync index close to 0). This transition did not occur at the same tempo for all the subjects: Table 8.2 represents the number of subjects that relied either on the synchronized behavior (sync index close to 1), on the decoupled behavior (sync index close to 0), or on an intermediate behavior (sync index around 0.5), for each tempo. The switching tempo varied across subjects: the nine subjects relied on the synchronized behavior at the fastest tempo ($\Delta t^* = 400\text{ms}$), six of them relying on the decoupled behavior at the slowest tempo ($\Delta t^* = 1000\text{ms}$). One subject was still at the intermediate level at $\Delta t^* = 1000\text{ms}$, while two subjects kept their arms in synchrony from $\Delta t^* = 400\text{ms}$ to $\Delta t^* = 1000\text{ms}$, suggesting that their switching tempo was beyond those tested. The synchronization index distribution is bimodal, since in most cases the subjects adopted one of the two actuation

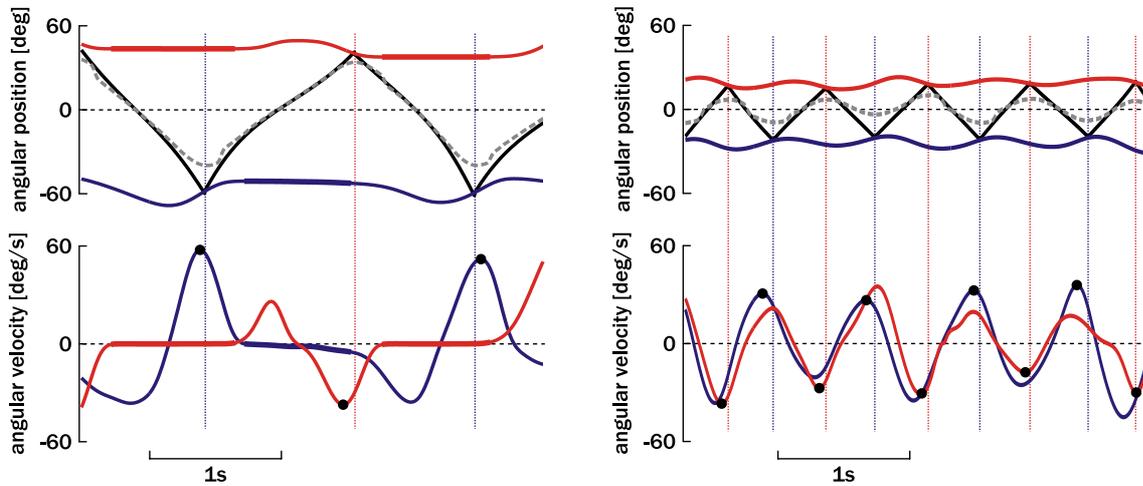


Figure 8.2: Typical behavior at slow tempo ($\Delta t^* = 1000\text{ms}$, left) and fast tempo ($\Delta t^* = 400\text{ms}$, right). Top panels, the right and left arm angular position (red and blue, respectively) and the angular position (θ) of the puck (black) are represented. For clarity, the radial position of the puck is not represented in this graph. The thicker portions (left) emphasize the periods when the arms were at rest. The vertical lines denote the impact times. The gray curve denotes the intersection between the gaze and the air-hockey table in the same system of reference (see Appendix C). Bottom panels, the right and left arm velocity (red and blue, respectively) are represented. The black dots denote the point of maximum velocity (velocity peak) around impacts.

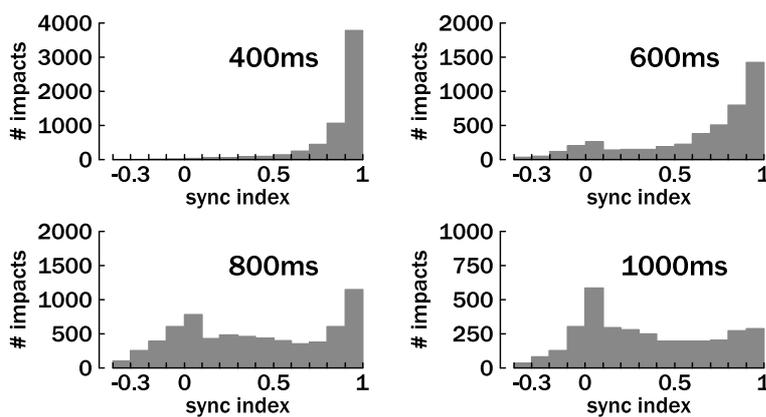


Figure 8.3: Distributions of the synchronization index for each tempo Δt^* . Only the data with visual feedback are displayed in these graphs.

Table 8.2: Inter-subjects variability of the transition from synchronized (sync index close to 1) to decoupled (sync index close to 0) behavior. For each tempo Δt^* , the corresponding column displays the number of subjects that relied mainly on each behavior (sync index around 0.5 for intermediate).

tempo Δt^*	400ms	600ms	800ms	1000ms
synchronized	9	6	2	2
intermediate	0	2	2	1
decoupled	0	1	5	6

modes. Importantly, the average synchronization index statistically decreased with the tempo for all nine subjects (ANOVAs).

The transition from sustained and synchronized movements (fast tempo) to intermittent and decoupled movements (slow tempo) is further illustrated in Fig. 8.4. This figure shows that the faster the tempo, the larger the synchronization index and the longer the activity period of the arms. At the fastest tempo ($\Delta t^* = 400\text{ms}$), both arms were almost always moving (activity period close to 100%) in synchrony (sync index close to 1). As the tempo decreased, the arms were actuated more and more intermittently, resulting in a loss in synchrony and a decreased activity period.

Fig. 8.4 shows that the same trend applies regardless of the visual feedback condition. However, it further reveals that the subjects relied more on the sustained and sync actuation mode when “deprived” from visual feedback. This effect is significant at each tempo. The smallest difference is reported at $\Delta t^* = 400\text{ms}$ since, in this case, the actuation was sustained and sync regardless of the visual feedback condition.

8.3.3 Performance vs. robustness of the sensorimotor loop

Gaze tracking

In normal visual conditions, the subjects performed a dynamical tracking of the puck (see Fig. 8.2). The good accuracy of the gaze tracking resulted in both a small *position* error of the gaze w.r.t. the puck, and also a high *velocity* gain of the pursuit during the flying phases, centered around the middle of the puck trajectory. This is similar to gaze strategies in the three-balls cascade juggling. Indeed, Huys and Beek (2002) reported that the point-of gaze was confined close to the zenith of the balls trajectory, while their dynamics were frequency-locked, suggesting both small position and velocity errors around zenith.

The position error has been quantified by measuring the angle between the gaze and the target (i.e. the vector from the eye to the puck, see (C.16)). This angle was smaller than 5° — such that the subjects kept the puck in their central vision —

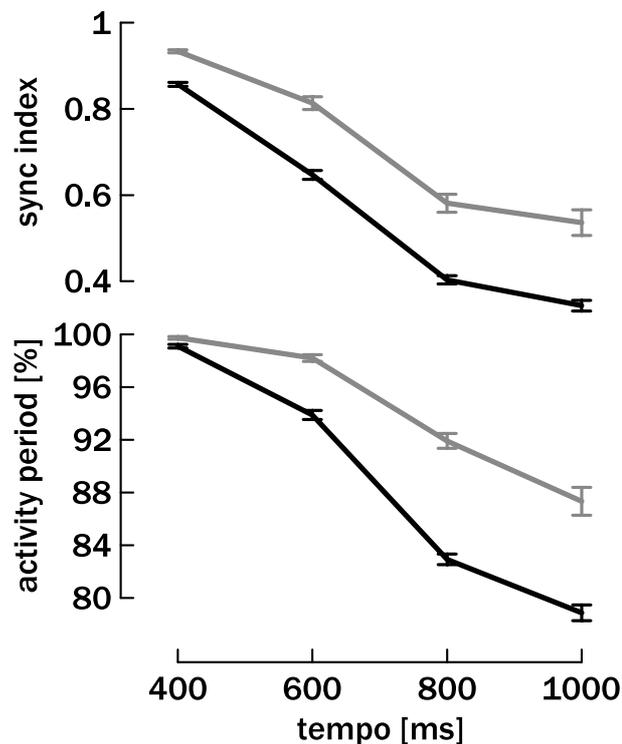


Figure 8.4: Synchronization index (top) and activity period (bottom) of the arms, as a function of the tempo Δt^* and the visual feedback condition: with visual feedback (black) and without visual feedback (gray). Mean \pm 0.95 Conf. Interval. Both parameters depended on the tempo for the two visual feedback conditions; and on the visual feedback condition for the four tempi, as suggested by the confidence intervals (Kruskal-Wallis ANOVAs).

during more than 63% of the flight time (on average), regardless of the tempo.

Likewise, the measured pursuit gain, i.e. the gain between the gaze and the puck velocity (in the subject's frame of reference), was on average larger than 0.8 at mid-flight regardless of the tempo.

An important consequence of this high accuracy is that subjects could use visual feedback in order to estimate the puck energy (at mid-flight), a relevant state information depending on the position and the squared velocity of the puck.

Closed-loop control of the puck energy

The regulation of the puck energy by means of the impact velocity is a central control loop of the juggling task. The arm velocity at impact indeed determined the energy *restored* to the puck. Deviation from the steady-state energy must be compensated for by corrections of the arm velocity at impact. We studied this negative feedback

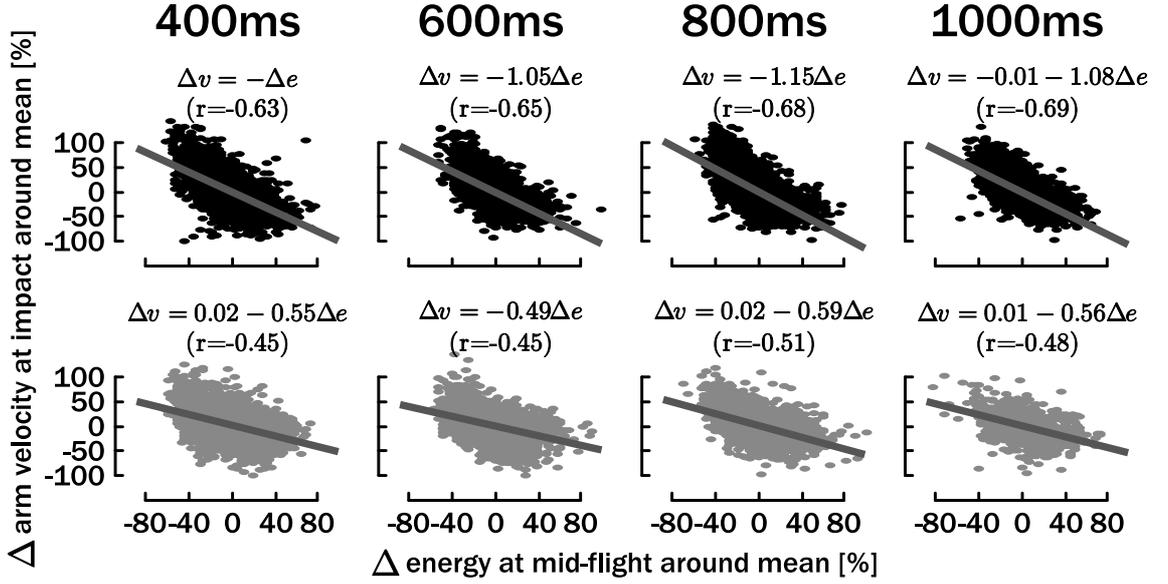


Figure 8.5: The relative (Δ) energy of the puck at mid-flight as a function of the relative (Δ) velocity of the arm at the following impact, around means, for each tempo and visual feedback conditions. The black points (top) depict the impacts with visual feedback and the gray points (bottom) those without visual feedback. The straight lines represent the linear regressions of these data, the regression slopes and the correlation coefficients being indicated above each graph ($p < 0.0001$ in all conditions). The “homogeneity of slopes” analysis revealed that the slope coefficients significantly depended on the visual feedback condition ($p < 0.0001$).

control loop by computing the correlation between the puck energy at mid-flight² and the angular velocity of the arm at the forthcoming impact. Both the puck energy and the arm velocity were normalized w.r.t. the average values, in order to compare the correlations across different conditions. For a given parameter x , whose instances are x_i and average \bar{x} , the normalized instances are:

$$\Delta x_i[\%] = \frac{x_i - \bar{x}}{\bar{x}} \times 100\%$$

Fig. 8.5 reveals that the negative feedback control loop was indeed implemented by the subjects at the four tempi, with or without visual feedback. The slopes were consistent across tempi *for a given visual feedback condition*: they ranged between -1 and -1.15 in the presence of visual feedback, while they were about twice smaller (between -0.49 and -0.59) in the absence of visual feedback. It is further observed that the data variance was higher without visual feedback ($r \simeq -0.47$) than with visual feedback ($r \simeq -0.66$).

²Since the puck flights were almost frictionless, the energy is roughly constant during the whole flight.

Table 8.3: Averaged across all subjects, the table reports the mean and standard deviation (SD) of the puck energy (per unit of mass) at mid-flight [$10^6\text{m}^2/\text{s}^2$], depending on the tempo and visual feedback (vFB) condition. The table displays also the F-ratio (F-r.), that is the ratio of the variance without visual feedback and the variance with visual feedback, for each tempo. The F-ratios are significantly different from 1 (all $p < 0.01$).

tempo Δt^*	400ms		600ms		800ms		1000ms	
	mean \pm SD	F-r.	mean \pm SD	F-r.	mean \pm SD	F-r.	mean \pm SD	F-r.
with vFB	2.01 ± 0.4	1.67	1.9 ± 0.37	1.86	1.92 ± 0.38	1.66	1.93 ± 0.38	1.56
without vFB	1.97 ± 0.56		1.85 ± 0.5		1.89 ± 0.49		1.94 ± 0.47	

Thus the puck energy was regulated through a negative feedback loop both with and without visual feedback: the subjects estimated the puck energy during the flying phases, and adapted the arm trajectory to impact at the required velocity. As reported in Section 8.3.3, the puck energy was potentially estimated by means of gaze tracking when the visual feedback was available. The present data suggest that, without visual feedback, the subjects were still able to estimate the puck energy, though with less accuracy. Importantly, the feedback gain was both larger (larger slope) and less variable (larger r -value) with visual feedback than without.

A larger gain permitted to regulate the puck energy more accurately. Table 8.3 shows the mean and standard deviation of the puck energy at mid-flight, as a function of the tempo and visual feedback condition. While there is no clear tendency for a dependence of the *mean* on the tempo or on the visual feedback, the *variance* of the data was much larger without than with visual feedback, as figured out by the standard deviations. This is further confirmed by the F-ratios, i.e. the ratios of the variances, that were significantly larger than 1 in favor of the ‘without visual feedback’ condition for each tempo.

In summary, by tuning a larger gain in the energy feedback loop, the subjects maintained the puck energy more efficiently around its steady-state value with visual feedback. Because the energy estimation was more uncertain, this large gain was not maintained without visual feedback, resulting in an increase in the puck energy variability.

8.3.4 Performance vs. robustness in movement amplitude

For a fixed visual feedback condition, the sensorimotor gain computed in the previous section was constant across tempi. Nevertheless, clear differences were observed in the actuation mode (see Section 8.3.2). The purpose of this section is to investigate how the two different control strategies (sustained/sync or intermittent/decoupled) balance the trade-off between performance and robustness at the level of trajectory

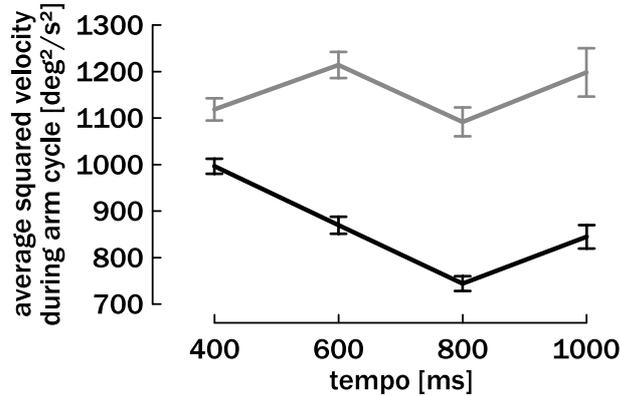


Figure 8.6: Averaged squared velocity of the impacting arm across a complete arm cycle, as a function of the tempo Δt^* and the visual feedback condition: with visual feedback (black) and without visual feedback (gray). Mean \pm 0.95 Conf. Interval. The influence of the tempo, the visual feedback condition, and their interaction were significant (factorial ANOVAs). As suggested by the confidence intervals, Tuckey’s HSD post hoc analysis revealed that the pairwise differences between the data with and without visual feedback were always significant, for any given tempo Δt^* ($p < 0.0001$).

planning. Indeed, one may hypothesize that the strategy adopted without visual feedback (i.e. when the sensory capability was severely degraded) must compensate for the degradation of the sensing capabilities by increasing the robustness. Increased robustness should however result in a decrease in performance.

The performance has been quantified by the energy expended to produce the movement. We estimated this quantity indirectly through the movement amplitude, by averaging the squared velocity of the impacting arm across a complete cycle (i.e. between two successive impacts on the same arm). Fig. 8.6 reveals that the absence of visual feedback increased the energy expense, at the four tested tempi. Moreover, the energy expense also increased at the fastest tempo ($\Delta t^* = 400\text{ms}$) in the presence of visual feedback. In contrast, there was no clear influence of the tempo without visual feedback. In summary, the sync and sustained actuation strategies were more energy expensive. Accordingly, the smallest difference between visual feedback conditions is reported at the fastest tempo ($\Delta t^* = 400\text{ms}$), since the actuation modes with and without visual feedback were sustained and sync in both cases.

We must now clarify why the augmented muscular effort allows for an increase in robustness. As reported in Chapter 4, on the basis of previous studies in motor control (Schaal et al., 1996; Sternad et al., 2001a,b), a crucial parameter has been emphasized to quantify the control robustness in impact dynamics: the arm *acceleration* at impact. These studies showed that *negative* acceleration at impact is a necessary condition for sensorless control, and we extended this requirement for

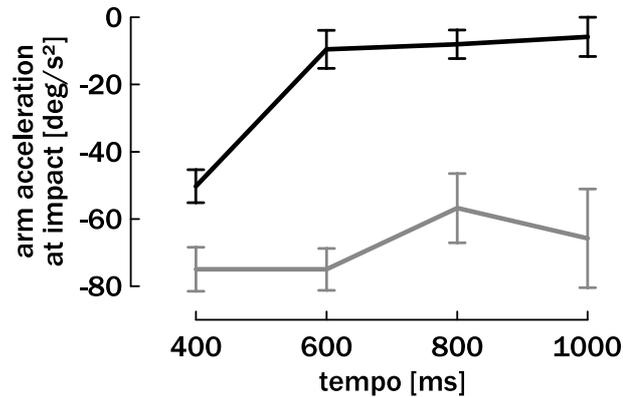


Figure 8.7: Acceleration of the arm at impact as a function of the tempo Δt^* and the visual feedback condition: with visual feedback (black) and without visual feedback (gray). Mean \pm 0.95 Conf. Interval. The influence of the tempo, the visual feedback condition, and their interaction were significant (factorial ANOVAs). As suggested by the confidence intervals, Tuckey’s HSD post hoc analysis revealed that the pairwise differences between the data with and without visual feedback were always significant, for any given tempo Δt^* ($p < 0.0001$).

robust closed-loop control (see Chapter 4).

For a fixed tempo, neither the impact position nor the velocity of the arms varied across visual feedback conditions: their steady-state values are fixed by the reference tempo and the coefficient of restitution ((6.12) and (6.55)). Conversely, the arm acceleration at impact strongly depended both on the tempo and *on the visual feedback condition*. Fig. 8.7 shows that the subjects adopted more negative acceleration at impact in the rhythmic actuation mode (fastest tempo and/or without visual feedback). In contrast, with visual feedback AND at slower tempi, i.e. at $\Delta t^* = 600\text{ms}$, 800ms and 1000ms , the acceleration at impact was close to zero. Apparently, the increased energy expense is the price to pay for a more negative acceleration at impact, resulting in more robustness.

Near zero acceleration means that the impact occurs near the velocity peak of the arm (see the black dots on Fig. 8.2). This strategy is optimal in term of energy expenditure. Conversely, negative acceleration at impact requires a larger effort (a larger velocity peak) *before* the impact to eventually reach the same requested impact velocity when decelerating. This further confirms that the sustained actuation mode is more robust but more expensive than the intermittent mode, and vice-versa. Interestingly, the robust mode of control was always preferred when the sensorimotor loop was degraded: either because of the visual feedback suppression and/or because of the increased bandwidth requirement at faster tempi. In consequence, the subjects balanced the loop degradation by adopting a robust mode of control, sacrificing some performance in trajectory planning.

8.4 Discussion

The present chapter investigates an original rhythmic task in motor control, requiring the stabilization of a juggled puck through an appropriate bimanual actuation pattern. The impact task is a 2D bimanual extension of the classical movement of a ball bouncing on a racket (see Chapter 3 and Schaal et al., 1996; Sternad, 1999; Sternad et al., 2001a,b; Katsumata et al., 2003; de Rugy et al., 2003; Dijkstra et al., 2004). The task was performed at four different tempi, and in two distinct visual feedback conditions. We studied how the subjects balance the performance and the robustness of the task across these different experimental conditions.

A primary role of any feedback loop is to reduce the sensitivity to uncertainties (see Appendix A). A higher feedback gain normally improves the closed-loop performance but deteriorates its robustness, increasing the proneness to instability in the presence of delays or noise measurement (see e.g. Franklin et al., 2005; Astrom and Murray, 2005). Indeed, the feedback gain directly multiplies the sensed outputs (including noise) to compute the input commands: the larger the gain, the larger the noise propagation. The feedback gain tuning is therefore a first crude manifestation of the fundamental trade-off between performance and robustness of feedback systems.

We studied this trade-off in the feedback control from the arm impact velocity to the puck energy. The gain of this feedback loop was twice larger in the presence of visual feedback, resulting in less variability of the puck energy around the steady-state. In contrast, the gain of the feedback was *not* affected by the *tempo* in a given visual condition. This result is somehow counter-intuitive since, when the tempo increases, the processing of sensory information (constant delay) becomes more and more significant with respect to the duration of movement (see e.g. Desmurget and Grafton, 2000). Consequently, the feedback signal starts to influence the ongoing motor command later in the flying phase, or even after the impact. Faced with delays becoming more and more significant with respect to the cycle period, the feedback gain must be decreased to maintain the closed-loop stability. However, our subjects were not affected by this requirement, since they maintained roughly the same gain even when they had to produce an impact each 400ms.

Two questions remain unsolved at this point: How was the negative gain generated in the absence of visual feedback, and how was the negative gain maintained even at the fastest tempi? First, the negative gain is also due to the presence of other sensory inflows than vision. Importantly, we further reported that the subjects changed their control strategies in these cases, i.e. to compensate for the uncertainties due to both the loss of visual feedback and the prominent delay in the loop. Indeed, the subjects reduced the pose intervals between strokes as the tempo increased, and eventually adopted a sustained behavior. The increase in the activity period was correlated with an increase in the synchronization index, since the arms cannot be independently actuated due to their intrinsic coordination rules (Swinnen, 2002; Swinnen and Wenderoth, 2004). The subjects also adapted the motion plan-

ning across visual feedback conditions, for a fixed tempo: without visual feedback, they favored the sustained and sync bimanual actuation.

The sustained and sync actuation mode was clearly adopted in order to benefit from a robustness mechanism emerging from this particular strategy. Indeed, the average arm acceleration at impact was negative when the sustained mode of control was adopted. In Section 4.3, we showed that *negative* acceleration at impact is a necessary condition for *robust* closed-loop control in impact dynamics with limited sensing capabilities. This is due to the sensorless dynamical properties of the task (i.e. control without feedback) which can be *stable* only if the acceleration at impact is negative (see (3.19) for the bouncing ball). Robust control with negative acceleration at impact has been moreover confirmed for experiments with Wiper within Chapter 7. In sum, robust control was then favored (1) without visual feedback, i.e. with degraded sensing capabilities, and (2) at the fastest tempo, i.e. when the delays were prominent with respect to the tempo.

In contrast, a majority of the subjects preferred to rest on the intermittent mode of control when the tempo was slow enough to plan and execute each movement, and when the visual feedback was available. This strategy does not benefit from the robustness due to negative acceleration (it resulted in nearly-zero arm acceleration at impact). However, these conditions were less subject to uncertainties, and the need for robustness can consequently be traded for performance in trajectory planning: the energy expenditure was indeed smaller when the intermittent mode was favored: with visual feedback at the three slowest tempi.

The sustained and synchronized actuation mode is then robust — due to negative acceleration of the arm at impact — but more energy expensive than the intermittent and decoupled mode. The robust mode of control was adopted when the most important source of sensory inflow was not available, or at the fastest tempo (even with visual feedback). Therefore, we hypothesize that the balance between the sustained and the intermittent mode was not only achieved to maintain a sufficient feedback gain across tempi — dictated by the sensing capabilities —, but also to compromise performance and robustness in planning of the movement trajectory.

The reported negative gain without visual feedback was also a consequence of residual closed-loop mechanisms, and not only a by-product of the strategy with negative acceleration. Other sensory inflows are indeed still available to estimate the puck energy. For instance, we demonstrated through the derivation of the state observer (7.3) that the sole detection of impact times was sufficient for closed-loop stabilization of periodic orbits (see Chapter 7). The actual implementation required however negative acceleration at impact, i.e. to supplement a robust control strategy to the limited sensing capability. The detection of impact times is available for human subjects through auditory and/or haptic feedback.

A general framework has been elaborated by Wolpert and colleagues to model the ability to generate accurate and appropriate motor behavior under many different and often uncertain environmental and contextual conditions (the MOSAIC model, see Wolpert and Kawato, 1998; Jordan and Wolpert, 1999; Wolpert and Ghahramani,

2000; Haruno et al., 2001). This architecture has been proposed for motor control, on the basis of multiple pairs of forward (predictor) and inverse (controller) models stored in the brain. The present chapter refers to a rhythmic task to which this framework extends. Indeed, the subjects adopted either the intermittent mode to maximize the performance (energy expenditure), or the sustained mode to maximize the robustness to uncertainties, depending on the task context. These controllers are modulated by the same “inputs” (e.g. the puck energy) but their outputs are thought to be from different levels of complexity. The intermittent controller has to plan (and execute) a complete arm trajectory, to reach the desired velocity at the desired time (Land and McLeod, 2000; Ronsse et al., 2007a). The sustained controller modulates a baseline rhythmic movement, strongly synchronized between the arms. Consequently, we postulate that the internal models in this mode encompass not only the puck and arm dynamics, but also the steady-state oscillation of the arms. The feedback is used to modulate this oscillation, for example through its amplitude (Buschges, 2005).

Many motor control investigations have focused on performance criteria, using for example the theory of optimal control to illustrate how human subjects perform a broad set of tasks by optimizing a given performance index (Jordan and Wolpert, 1999; Wolpert and Ghahramani, 2000; Todorov and Jordan, 2002; Scott, 2004; Todorov, 2004, 2006). However these performance criteria have never been paralleled with *robustness* criteria. Nevertheless, it is well established in control theory that performance and robustness are in competition, since they cannot be arbitrarily improved simultaneously (see e.g. Section A.3 and Franklin et al., 2005; Astrom and Murray, 2005). For that reason, a closed-loop design which focuses only on performance indices may lead to disastrous behavior in the face of uncertainties, because it is not robust. The present chapter illustrates two compromises between performance and robustness at the behavioral level, for the control of a rhythmic task across various contexts. First, the task performance was quantified by the feedback gain of the energy control loop. The feedback gain was reduced when the sensing capabilities were degraded, improving the robustness to uncertainties. Secondly, the performance of the planning (energy cost) was traded for robustness (negative acceleration). The robust control mode was preferred when the sensorimotor loop was degraded.

8.5 Conclusion

This chapter main objective was to demonstrate the importance of an essential concept of control theory which is often disregarded in motor control studies: the robustness of the actuation strategy. Through an informative example, we illustrated that the control behavior of a given task may either be dictated by *performance* criteria or by *robustness* criteria, while the task objective remains unchanged. Moreover, we confirmed that these two concepts cannot be arbitrarily improved simultaneously, since they result from a trade-off in motor planning. This chapter is submitted for publication (Ronsse et al., 2007c).

Chapter 9

General Discussion and Perspectives

Nature’s music is never over; her silences are pauses, not conclusions.

Mary Webb

Thanks to a common experimental setup — the Wiper robot described in Chapter 5 —, fundamental questions in robotics and human motor control have been studied in parallel. These questions focused on the systems that are *rhythmic* (in the sense that the input command must be periodic) and *underactuated* (i.e. the actuator degrees of freedom are fewer than the system degrees of freedom). This makes the dynamics *hybrid*, since continuous-time dynamics (flight) interact with discrete-time dynamics, i.e. at impacts. Chapters 4, 6 and 7 contributed to the robotics part. They provided control design solutions for underactuated systems in which the input command influences the system state only through instantaneous impacts. Moreover, these results have been successfully validated in the lab. Chapter 8 analyzed the human behavior when controlling the *same* task, and reported different strategies depending on the context.

The aim of this general discussion is to gather the most important conclusions of the thesis, and to discuss them within the perspective of an integrated approach, between robotics and neuroscience. The discussion is divided into three sections:

- Section 9.1 is dedicated to impact juggling. We emphasize why this thesis aims at reconciling two “historically” distinct approaches for the control of impact juggling: open-loop and closed-loop. Of prime interest is the discussion of the trade-off between open-loop and closed-loop that is also present in human behaviors.
- Section 9.2 stresses the main arguments calling for an integrated approach between neuroscience and robotics. Mainly, we recall why Wiper is a unique and very appropriate setup for further investigations and experiments.

- Section 9.3 proposes a model capturing the human control strategies in bouncing experiments. The keystone of this model is to produce either discrete and individually planned movements, or rhythmic movements depending on the context.

In Section 9.4, we wrap up the central message of the thesis as a general conclusion, while in Section 9.5, we mention future perspectives and open questions in the continuity of the present manuscript.

9.1 Impact juggling control

One-dimensional impact juggling of a bouncing ball with a racket has been recognized as a benchmark in robotics and neuroscience (see Chapter 3). Despite an apparent simplicity, this task is complex, since a large number of kinematic profiles can be adopted by the racket to reach the required position and velocity at impact. From a control perspective, two approaches have been investigated for the control of this task. The *mirror law algorithms* developed by Buehler, Koditschek and Kindlmann (1988, 1990, 1994) rest on an intuitive strategy to compute the racket trajectory: the racket has to mirror the ball — up to a fixed coefficient — which is consequently permanently tracked. In 2D and 3D, further feedback loops are added to confine the ball(s) motion to vertical bouncing. The mirror law introduces sharp discontinuities in the racket trajectory at impact, due to the mirroring strategy¹. In real-time implementations, the actual trajectory is obviously smoothed due to the n -order dynamics of the actuator. Nevertheless, the model remains valid assuming that the actuator re-tracks the “ideal” trajectory significantly before the next impact.

Sharply contrasting with the permanent tracking required by the mirror law, sensorless sinusoidal actuation achieves the stabilization of several periodic orbits in open-loop (see Holmes, 1982; Guckenheimer and Holmes, 1986; Bapat et al., 1986; Tufillaro et al., 1992; Schaal et al., 1996). Nevertheless, the parametric stability region is limited into finite ranges of both the amplitude and frequency. Schaal, Atkeson and Sternad (1996) noticed an important feature of the stabilization of the period-one (the simplest periodic orbit): the racket acceleration at impact must be *negative*, within the range defined by (3.19) according to the model. In contrast, the mirror law is closed-loop controlled and impacts the ball with positive acceleration (the mirror of the gravity).

In Chapter 7, we proposed an approach to reconcile those different strategies, i.e. to produce smooth trajectories which are based on feedback sensing. The principle of the proposed strategy is to modulate the amplitude of the sinusoidal vibration depending on the estimated velocity of the ball. The negative acceleration

¹Indeed, with the mirror law, the position of the racket is continuous at impact, but the velocity obeys a discontinuous bound which is the mirror of the impact law (3.3).

at impact is then no longer required as a necessary condition for open-loop stability. However, we proved in Chapter 4 through a simplified actuation trajectory that it is a necessary condition for *robust* control. The need for robustness has been particularly emphasized in a complex impact juggling experiment, requiring the stabilization of a puck in a 2D *plane* and under “bimanual” actuation.

Such hybrid strategies (continuous actuation based on discrete feedback) have already been studied in the literature for controlling the bouncing ball, on the basis of a modulation of the sinusoidal actuation (Vincent and Mees, 2000; de Rugy et al., 2003). With respect to these earlier studies, our contributions are threefold:

1. We quantified the *robustness* of our design, and we showed that negative acceleration at impact is necessary for robust control.
2. We studied a feedback control which is only based on the detection of impact times and not on sensing the ball state, reducing the need for sensors to mere impact detectors.
3. We validated our control strategy in the lab with an actual setup that required to extend the design to a two-dimensional space.

Another contribution of the thesis was to disentangle how humans performing the same task also face the trade-off between open- and closed-loop control. We have seen that the strategy is not sensorless in normal conditions, since the suppression of visual feedback significantly altered the gain of the puck energy control loop. However, the subjects executed the task with smooth movements (either a train of discrete strokes or a sustained rhythmic actuation), sharply contrasting with the mirror law profile. Suppression of visual feedback or fast tempo favored the sustained actuation mode. This mode was sinusoidally-shaped and corresponded to negative arm acceleration at impact, highlighting the need for robust control when the sensorimotor loop was degraded.

We have not studied in this thesis to what extent the negative feedback of the energy control in the absence of visual feedback was due to (1) only the negative impact acceleration (sensorless); (2) remaining closed-loop mechanisms (due to the other sources of sensory inflow); or (3) a mixture of both of them. Based on preliminary simulation results — which are not reported in the present manuscript — we favor the third hypothesis: the sensorless sinusoidal actuation (with negative acceleration at impact) indeed results in a negative gain in the energy control loop, but much smaller than those reported in Fig. 8.5. A significant amount of this negative gain must then result from additional closed-loop mechanisms. For instance, a state observer based on impact times could be implemented on the basis of a structure similar to (7.3). Indeed, the impact times are still measurable in the absence of visual feedback both through the auditory and haptic feedback.

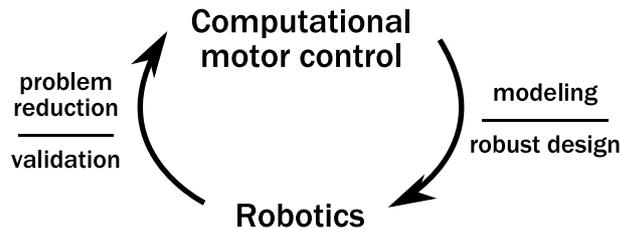


Figure 9.1: Experimental results in computational motor control and robotics cross-fertilize.

9.2 Neuroscience and robotics

The neural and behavioral study of motor control parallel with questions in robotics (mainly with humanoid robots) in many points (Schaal and Schweighofer, 2005). We tried to capture these interactions within Fig. 9.1. On the first hand, computational models can provide useful guidance in the interpretation of complex, high dimensional biological data sets. More particularly, the mathematical modeling of the task may help at reducing the problem complexity by emphasizing a limited number of relevant parameters, being useful to analyze the behavior and discriminate between different control strategies. On the other hand, models and algorithms from robotics research provide useful inspiration, baseline performance, and sometimes direct analogs for motor control.

Our study took advantage of the interaction between computational motor control and robotics. It provided a direct comparison between the theoretical predictions, the robotics implementations, and the human behavior. The data reported in Chapter 8 suggested that humans use mainly two distinct primitives to control rhythmic tasks, i.e. an intermittent train of discrete strokes or a sustained rhythmic movement. This might inspire a flexible computational model to generate complex movements, based on simple interactions between fundamental primitives.

The major contribution of the thesis in that respect was to design an original experimental setup which is easily configured for both kinds of experiments. Using the same setup and the same task, the comparison of both data sets is straightforward. For example, given the robot performance with the sensorless strategy and its lack of robustness in real experimental conditions, we anticipated that the human behavior could not be sensorless, even in the absence of visual feedback. Indeed, even in control strategies that resembled a sinusoidal actuation (sustained and sync), we have observed an active closed-loop control of the puck energy. Moreover, the negative acceleration reported with that mode was shown to be an effective measure of robustness, as illustrated in the closed-loop robotics experiment.

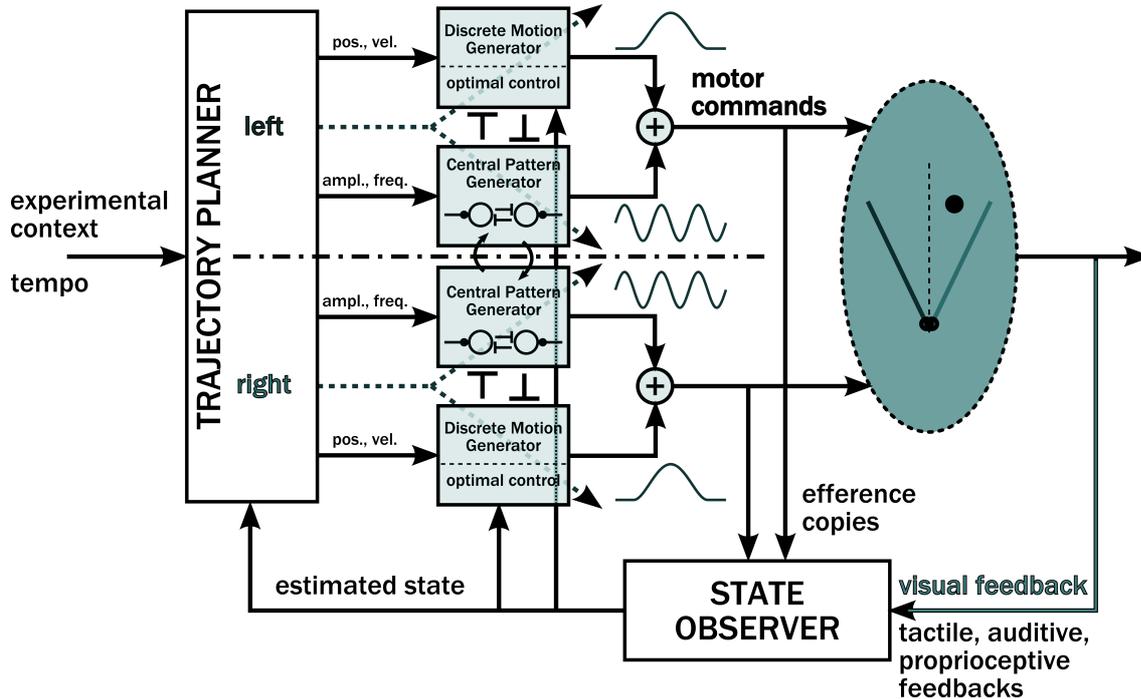


Figure 9.2: Sketch of a black-box model for the control of rhythmic tasks.

9.3 Towards a model for the control of rhythmic movements

A model capturing the observed human behavior in the control of complex impact dynamics has not been described in this manuscript, but a first sketch could be established on the basis of the conclusions of Chapter 8: the intermittent and sustained control modes are based on different fundamental primitives of movements, and consequently rest on different control architectures.

A plausible model architecture is proposed in Fig. 9.2. Its central black-box is the so-called trajectory planner (Section 1.1), which has to decide which movements will be executed depending on the experimental context (e.g. the tempo) and the available sensory feedback. The trajectory planner generates movements on the basis of two fundamental primitives:

- A discrete movement generator is used to produce individually planned discrete strokes. In this mode of control, one may hypothesize that the planning and the execution of the movement do not overlap. In impact tasks, the goal of the planned movement is to reach the desired position at the desired time, and with a desired velocity (Land and McLeod, 2000; Ronsse et al., 2007a). Optimization theory provides a computational framework which is ideal for a selection process such as motor planning (Jordan and Wolpert, 1999; Wolpert

and Ghahramani, 2000; Todorov and Jordan, 2002; Scott, 2004; Todorov, 2004, 2006).

- A central pattern generator (CPG) produces the rhythmic movement (see e.g. Cohen et al., 1988; Duysens and Van de Crommert, 1998; Swinnen, 2002; Zehr et al., 2004; White et al., 2007). In this case, we hypothesize that the trajectory planner does not plan each movement cycle separately, but rather modulates the baseline rhythmic movement (amplitude, frequency and phase). Potentially, trajectory planning and execution may overlap in this case.

The left-arm and right-arm central pattern generators are coupled with each other, in order to produce rhythmic movements which are in agreement with the well established bimanual coordination principles (Swinnen, 2002; Swinnen and Wenderoth, 2004).

The quality of the feedback path obviously depends on the available sensory inflows. In Fig. 9.2, we highlighted the visual feedback, since this particular source of information has been manipulated in the experiment reported in Chapter 8. The presence or absence of visual feedback is also the fundamental difference between the mirror law and the new family of closed-loop control strategies for bouncing tasks presented in Chapters 4 and 7. Both for our robot or for the human subjects, the quality of the state estimation (through the state observer) was significantly impaired without visual feedback, claiming for the need of a trajectory planning which is robust to these uncertainties. Consequently, the trajectory planner will favor the rhythmic central pattern generator when the state estimation is impaired due to the absence of visual feedback; or when separate planning of discrete movements is not possible given the delays in the loop. In the other cases, the trajectory planner favors the planning of intermittent discrete movements, in order to minimize the energy cost.

The model described above remains conceptual and requires further refinements.

9.4 Conclusion of the thesis

The key message of this thesis is that rhythmicity favors robustness. This has been illustrated within an integrated approach in robotics and neuroscience (motor control). In robotics experiments, we have shown that a slight modulation of a passive rhythmic actuation on the basis of limited sensing capabilities can dramatically enlarge the robustness of the control. This paradigm provides a useful guideline for the design of robots in rhythmic environments: the reduction of the sensors cost is guided by a suitable exploitation of the sensorless dynamics. The robustness of rhythmic trajectories has been confirmed by the experiments with human subjects, who sacrificed a performance index (the energy cost of their movements) to adopt a rhythmic control mode when robustness was imposed by the experimental context.

9.5 Future perspectives

At the time of completing this manuscript, a lot of exciting challenges are left unsolved. The most obvious of them is the stabilization of the n -pucks (with $n > 1!$) shower with Wiper. The most critical step toward this challenge has been solved since we succeeded to stabilize the period-two with one puck (see Section 7.3). However, two technical issues still remain before the extension to n -pucks:

The initialization: Another initialization mode has to be found, since we used the sensorless control of the period-one to initialize the movement and to let the observer converge. However, it is not possible to juggle the period-one with more than one puck.

The impact detection: The accelerometers we used would not be able to distinguish which puck impacted the arms. However, an individual processing of the feedback generated by each puck has to be computed, and the state observer of each puck could only be implemented if the setup identified which puck caused which impact.

Other research perspectives are also pointed in the continuity of the present project.

Design of continuous feedback controllers for impact tasks. The reduction of sensing capabilities is somehow pushed to the limit in the present thesis. It would be of interest to exploit the robustness analysis (which e.g. showed the importance of negative acceleration) in the design of control laws that can rely on additional sensing capabilities, e.g. a permanent tracking of the puck. It would be interesting to characterize the role of the acceleration at impact in the sensitivity to other types of perturbations, for example delays in the loop.

Vision dynamics in impact tasks. We have recorded the eye movements in the experiment reported in Chapter 8, however only static gains of vision have been analyzed in this chapter (position and velocity errors). Further analysis of the gaze dynamics would provide useful insights into the way the ocular and upper-limbs systems are coupled in juggling tasks, and to inspire humanoid robotics designs as well.

Disentangling the human control strategies. Chapter 8 focused on the human behavior to control the period-one of Wiper. However, the present analyzes capture only partially the actual control strategies. On the first hand, the differences between the discrete and the rhythmic primitives, and the subsequent trade-off between acceleration at impact and actuation energy, are reasonably clear. On the other hand, the closed-loop control of the puck energy has only been crudely understood so far. Indeed, the signature of negative feedback tuning is significant, but the low

correlation levels suggests that they only partially capture the closed-loop control. Thus these points need to be further investigated.

A model for the control of periodic movements at various tempi. In Section 9.3, we presented a first version of a computational model aiming at reproducing the human behavior. The central question of such a model is to propose a structure which is able to produce both discrete and rhythmic movements. In the literature, this kind of model has been investigated on the basis of separate primitives (see e.g. Sternad et al., 2000) or on a similar neural architecture (see e.g. Schoner, 1990; de Rugy and Sternad, 2003; Jirsa and Kelso, 2005). However, it has never been established how the planning of either discrete or rhythmic movements could be dictated by a performance vs. robustness trade-off. Thus future perspectives include an accurate elaboration of such a model.

The role of learning. This thesis did not focus on learning issues. Once again, this concept plays an important role, both in robotics and neuroscience: these fields may serve as a mutual source of inspiration (Schaal and Schweighofer, 2005). In behavioral experiments, learning refers to the changes in the control strategy which are only guided by the past experience and mastering in the task fulfillment. For example, it would be interesting to investigate whether the choice between the sustained or the intermittent control mode varies during the course of learning. In robotics, machine learning algorithms could also be considered to increase the controller performance on the basis of past data.

Appendices

Appendix A

Relevant Insights in Control Theory - a Benchmark Example

There is no teaching to compare
with example.

Sir Robert Baden-Powell

This thesis deals with several important concepts of control theory. Two of them are particularly emphasized throughout the manuscript: the balance and the complementarity between *feedback* and *feedforward* control, and the trade-off between *performance* and *robustness*. This appendix objective is to illustrate those two concepts on a simple example, in order to provide the reader with the major insights that are necessary for understanding how they are related to the rest of the thesis. The treated example belongs to the family of the *linear, time-invariant* (LTI) systems. This systems class is extremely convenient to deal with, since the mathematical tools for analysis and design are inherited from linear algebra. Nevertheless, this appendix stays also at the level of concepts, which are relevant for non-linear systems, such as the bounce jugglers presented in the other chapters.

A.1 A benchmark example

Let us consider a simple electrical DC motor as open-loop system. The motor is represented in Fig. A.1. The open-loop model can be split into an electrical and a mechanical part. The electrical part derives from the Kirchhoff's laws:

$$V_m(t) = R_m I_m(t) + L_m \frac{dI_m(t)}{dt} + k_m \omega_m(t). \quad (\text{A.1})$$

The last term denotes the back-electromotive force, induced by the current into the rotating motor, and therefore proportional to the motor velocity. In the Laplace

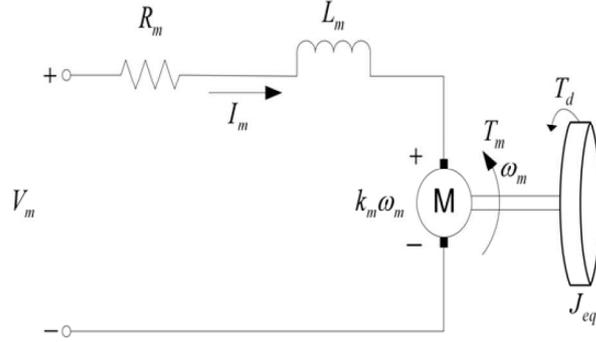


Figure A.1: Electrical DC motor.

domain, (A.1) gives:

$$V_m(s) = R_m I_m(s) + L_m s I_m(s) + k_m \Omega_m(s). \quad (\text{A.2})$$

Assuming that the electrical time constant (L_m/R_m) is very small, and can consequently be neglected w.r.t. the mechanical time constant, (A.2) simplifies to:

$$V_m(s) = R_m I_m(s) + k_m \Omega_m(s). \quad (\text{A.3})$$

The mechanical equation of motion of the motor derives from the rotational Newton law (the friction is neglected):

$$J_{eq} \frac{d\omega(t)}{dt} = k_m I_m(t) \quad (\text{A.4})$$

since the motor torque is proportional to the input current. J_{eq} denotes the moment of inertia of the motor and its load. In the Laplace domain, (A.4) becomes:

$$J_{eq} s \Omega_m(s) = k_m I_m(s). \quad (\text{A.5})$$

Combining (A.3) and (A.5), we obtain the following DC motor equation, between the voltage provided to the motor (input) and the motor velocity (output):

$$k_m V_m(s) = (R_m J_{eq} s + k_m^2) \Omega_m(s) \quad (\text{A.6})$$

which is equivalent to the following open-loop transfer function:

$$G_m(s) \equiv \frac{\Omega_m(s)}{V_m(s)} = \frac{k_m}{R_m J_{eq} s + k_m^2} = \frac{K}{\tau s + 1} \quad (\text{A.7})$$

where the open-loop gain and the open-loop time constant are equal to $K = 1/k_m$ and $\tau = R_m J_{eq}/k_m^2$. The (realistic) numerical values used in the rest of this section are $K = 20 \text{rad/Vs}$ and $\tau = 0.1 \text{s}$.

A.2 Feedback and feedforward control

The term *feedback* is used to refer to a situation in which two (or more) dynamical systems are connected together such that each system influences the other and their dynamics are thus strongly coupled (Astrom and Murray, 2005). In the treated example, one system is the motor, and the second is a *controller*. The coupling from the controller to the motor is obvious, since the controller assigns the input voltage of the motor. The motor state influences also the controller, since its output (i.e. the motor velocity) is *measured* to adapt the controller output (i.e. the motor voltage) consequently. The classical feedback loop is depicted in Fig. 1.2(b): the controller C_{FB} updates the system input $u(t)$ (the voltage in this case), depending on the difference between the reference signal $r(t)$ and the actual measured output $y(t)$ (the motor velocity in this case).

Usually, we may summarize the advantage of feedback control in three points:

- Feedback can be used to stabilize unstable open-loop dynamics. Nevertheless, this is not the case of the DC motor (A.7).
- Feedback can be used to speed up the dynamics, by tuning the closed-loop bandwidth at a larger value than the open-loop one.
- Feedback enhances the system *robustness* with respect to uncertainties or perturbations.

For example, imagine that the requirements of the loop design are: asymptotic tracking (zero static error), no overshoot, and doubling of the open-loop bandwidth in closed-loop, i.e. $\omega_0 = 1/0.05 = 20\text{rad/s}$. A first naive approach to design a controller respecting these requirements is to do *feedforward* control, according to an architecture relaxing the need of sensors (see Fig. 1.2(a)). Indeed, in such a structure, the output $y(t)$ has not to be measured to compute the input signal $u(t)$. Moreover, given the open-loop transfer function (A.7), a feedforward controller which respects all of the design requirements can be straightforwardly derived:

$$C_{FF}(s) = \frac{0.1s + 1}{20(0.05s + 1)} \quad (\text{A.8})$$

and the step response of the corresponding transfer function $C_{FF}(s)G_m(s)$ is given in Fig. A.2(a), black line. The step response has indeed no static error (the static gain equals 1); the settling time¹ is about 0.19s, i.e. twice faster than the open-loop; and there is no overshoot, since the resulting dynamics is of first order.

However, when studying the controller performance with respect to unexpected *perturbations*, Fig. A.2(a) reveals large discrepancies between the desired and the actual behavior. Indeed if the gain K varies of $\pm 20\%$ (red curves), the static error is now of the same order of variation. If the time constant τ varies of $\pm 20\%$ green

¹I.e. the time from which the step response enters forever into a “tube” of 5% of its final value.

curves), the settling time is much larger in one case, while there is overshoot in the other. In real systems, it is tedious to identify the open-loop parameters with more than 20% of accuracy. Moreover, there is usually unexpected (or unmodeled) perturbations in the dynamics that cause variability of (at least) the same order of magnitude. Feedback is consequently mandatory to make a system robust to such perturbations.

The classical closed-loop approach to achieve asymptotic tracking, despite perturbations, is to do *proportional-integral* (PI) control:

$$\begin{aligned} u(t) &= k_p (b_{sp}r(t) - y(t)) + k_i \int_0^t (r(\tau) - y(\tau)) d\tau \\ &= k_p (r(t) - y(t)) + k_i \int_0^t (r(\tau) - y(\tau)) d\tau + k_p (b_{sp} - 1)r(t). \end{aligned} \quad (\text{A.9})$$

The first term of (A.9) is proportional to the difference between the reference and the actual output. The second term of (A.9) is proportional to the *integral* of the difference between the reference and the actual output. The third term of (A.9) is directly proportional to the reference and is therefore a (static) feedforward part. The “magic” of the integral action is that it always cancels the static error, if the closed-loop system is stable. Indeed, when $u(t)$ has reached its steady-state, the integrand of (A.9) has to be zero, and thus $y(t) = r(t)$, as $t \rightarrow +\infty$.

Assuming first that $b_{sp} = 1$ (no feedforward control), such that the closed-loop design corresponds to Fig. 1.2(a), the transfer function of the controller is given by:

$$C_{FB}(s) = k_p + \frac{k_i}{s} = \frac{k_p s + k_i}{s} \quad (\text{A.10})$$

and the corresponding closed-loop transfer function equals:

$$\frac{C_{FB}(s)G_m(s)}{1 + C_{FB}(s)G_m(s)e^{-\Delta s}} = \frac{K(k_p s + k_i)}{\tau s^2 + (1 + K k_p e^{-\Delta s})s + K k_i e^{-\Delta s}}. \quad (\text{A.11})$$

Neglecting the delay ($\Delta \simeq 0$), the gains k_p and k_i are used to *place the poles* of the closed-loop transfer function at the desired location. One may choose them to equal the denominator of (A.11) with the following characteristic polynomial: $s^2 + 2\zeta\omega_0 s + \omega_0^2$, where ζ is the desired damping factor and ω_0 is the desired natural frequency. This gives:

$$k_p = \frac{2\zeta\omega_0\tau - 1}{K}, \quad (\text{A.12})$$

$$k_i = \frac{\omega_0^2\tau}{K}. \quad (\text{A.13})$$

The design objectives invite to choose $\zeta = 1$ (critical damping factor, i.e. no overshoot) and obviously $\omega_0 = 20\text{rad/s}$. The corresponding gains are $k_p = 0.15$ and $k_i = 2$.

The step response of the so-designed closed-loop DC motor is given in Fig. A.2(b) (with $\Delta = 0.01$ s). As expected, PI control ensures a zero static error whatever the perturbations. In general, the five curves are close to each other.

The 5% overshoot of the closed-loop step response (black curve in Fig. A.2(b)) is due both to the delay — not accounted for in the design — and to the *zero* (i.e. the first order term of the numerator) of the closed-loop transfer function (A.11). This turns the dynamics to be more reactive than the “pure” second order we have designed. Nevertheless, this zero can also be freely placed by considering *feedforward* control ($b_{sp} \neq 1$), above the feedback loop (Fig. 1.2(c)). By considering simply a static feedforward gain, i.e. $C_{FF}(s) = k_p(b_{sp}-1)$, the loop transfer function becomes:

$$\frac{\frac{b_{sp}k_p s + k_i}{s} G_m(s)}{1 + \frac{k_p s + k_i}{s} G_m(s) e^{-\Delta s}} = \frac{K(b_{sp}k_p s + k_i)}{\tau s^2 + (1 + Kk_p e^{-\Delta s})s + Kk_i e^{-\Delta s}} \quad (\text{A.14})$$

and the zero of (A.14) can be placed with b_{sp} , i.e. the feedforward gain.

As an example, we may place the zero to cancel one of the closed-loop poles, i.e. $k_i/(b_{sp}k_p) = \omega_0$. This gives $b_{sp} = \omega_0\tau/(2\zeta\omega_0\tau - 1) = 0.67$, assuming that the other gains remain unchanged. The corresponding step responses to the nominal and perturbed motor transfer functions are given in Fig. A.2(c). There is no overshoot anymore in the step response of the nominal transfer function, since it reduced to first order (neglecting the delay). Asymptotic tracking is still achieved whatever the perturbation, and the settling time corresponds to the requested one. Note that more complex designs could be made by adding dynamical terms to the feedforward path.

In summary, the feedback is highly efficient to make the design robust to unexpected variations in the open-loop transfer function, or to external perturbations. The feedforward is complementary to the feedback, since it adds new degrees of freedom to the design. Moreover, since the feedforward part does not rest on sensory measurements, it may considerably facilitate the system design if robust enough: some state variables are indeed very hard to measure, while sensory processing is inherently subject to delays...

A.3 Robustness vs. performance

The previous section has already emphasized the role of feedback to strongly increase the robustness of a given system. This section objective is to rapidly overview some concepts which *quantify* the robustness of any LTI closed-loop system. Moreover, we shall see that increasing the robustness of the previous design can only be made by degrading its *performance*.

Two important transfer functions to study the robustness of a closed-loop system are called the *sensitivity function* $S(s)$ and the *complementary sensitivity function*

$T(s)$:

$$S(s) = \frac{1}{1 + C_{FB}(s)G(s)}, \quad (\text{A.15})$$

$$T(s) = \frac{C_{FB}(s)G(s)}{1 + C_{FB}(s)G(s)}. \quad (\text{A.16})$$

They capture respectively the closed-loop sensitivity w.r.t. to small variations in the process ($S(s)$), and large variations in the process ($T(s)$). They are said to be complementary since $S(s) + T(s) = 1$. In general, it is important to minimize these functions, in order to keep the closed-loop system robust to both the unexpected perturbations (typically at low frequencies) and the noise (typically at high frequencies). In order to minimize these functions, the problem can be reduced to minimize their maximum M_S and M_T .

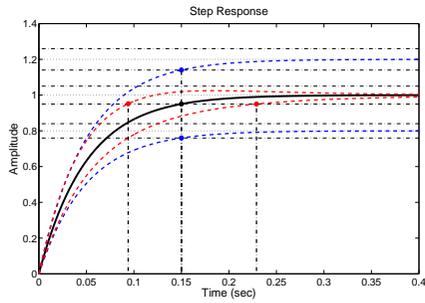
Fig. A.2(d) displays the maximum of the sensitivity functions, namely M_S and M_T^2 , for different closed-loop designs corresponding to different natural frequencies ω_0 , while considering $b_{sp} = 1$ (no feedforward control, for simplicity). The delay is still equal to $\Delta = 0.01$ s. As depicted in this figure, the initial design, i.e. $\omega_0 = 20$ rad/s, is not optimal to minimize the maximum of the sensitivity functions. Conversely, we may choose $\omega_0 \simeq 7$ rad/s to minimize M_S , while keeping M_T reasonably bounded. Note however that this closed-loop natural frequency is even smaller than the open-loop one, since the delay has been detrimental for the sensitivity functions.

The second design is then more robust than the first one. This is further confirmed by Fig. A.2(e) which displays the Nyquist plots of both loop functions. The stability margins, i.e. the gain and phase distances between the curve and the point of instability $(-1, 0)$, are larger for the second design (blue).

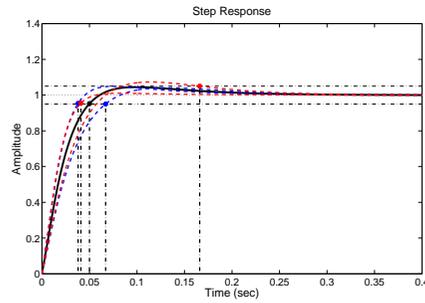
Nevertheless, the increase in robustness has been obtained by reducing the natural frequency (i.e. the closed-loop bandwidth), resulting in a decrease in performance. Indeed, the settling time corresponding to the second design (small bandwidth) is much larger than the settling time corresponding to the first design (see Fig. A.2(f)).

In general, performance and robustness of closed-loop systems cannot be arbitrarily increased simultaneously and result from a design trade-off.

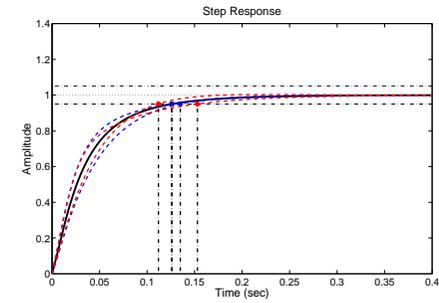
²For computational simplicity, we have appended the sensory delay directly to the open-loop transfer function: $G_m(s) = Ke^{-\Delta s}/(\tau s + 1)$.



(a) Forward controller.

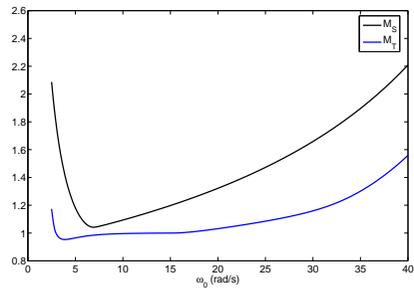


(b) Closed-loop controller.

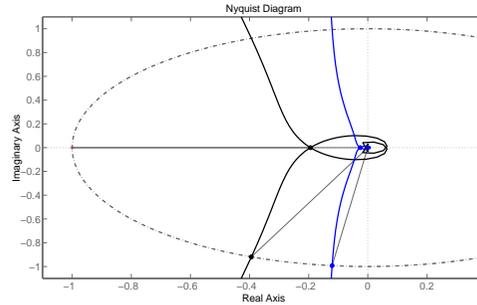


(c) Combined feedback+feedforward controller.

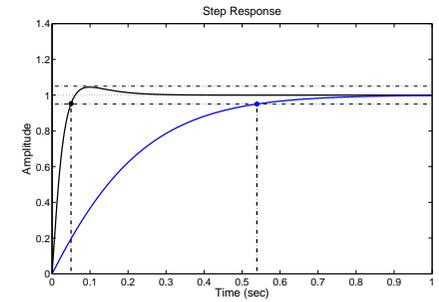
(a), (b) and (c) Step responses to the nominal transfer function (A.7) (black), the transfer function with an error of $\pm 20\%$ in the gain K (blue), and the transfer function with an error of $\pm 20\%$ in the time constant τ (red). The corresponding settling times are represented by colored dots, and vertical lines.



(d) Maximum of the sensitivity function $S(s)$ (M_S , black) and the complementary sensitivity function $T(s)$ (M_T , blue) for different closed-loop natural frequencies ω_0 .



(e) Nyquist plots and stability margins of $C_{FB}(s)G_{m2}(s)$ where the design is made for $\omega_0 = 20\text{rad/s}$ (black) and $\omega_0 = 7\text{rad/s}$ (blue).



(f) Step responses of the closed-loop controller to the nominal transfer function (A.7) with the designed closed-loop bandwidth at $\omega_0 = 20\text{rad/s}$ (black) and $\omega_0 = 7\text{rad/s}$ (blue). The corresponding settling times are represented by colored dots, and vertical lines.

Figure A.2: Figures related to Appendix A.

Appendix B

Major Technical Characteristics of the Experimental Setup

Technical skill is mastery of complexity, while creativity is mastery of simplicity.

Erik Christopher Zeeman

This appendix provides the main technical details about the practical implementations of the Wiper robot, presented in Chapter 5. Section B.1 describes the technical configuration of the setup in the “robotics” version. This configuration was used for experimental validations, reported in Chapters 6 and 7. Section B.2 refers to the “human” configuration, used for a motor control investigation on human juggling, reported in Chapter 8.

B.1 Technical details about the robotics configuration of Wiper

In the present section, we will give the main details about the technical implementation of Wiper which led to robotics experiments. These experiments are reported in Chapters 6 (sensorless control) and 7 (feedback control). A general picture of the setup in this configuration is provided in Fig. 5.3.

B.1.1 Hardware description

Air hockey table

The air-hockey table dimensions are given in Fig. B.1. A supporting structure has been manufactured such that the table can be tilted between 0° and 45° thanks to an electrical power jack (Servomech, see Fig. B.2). The air-hockey table is pierced with

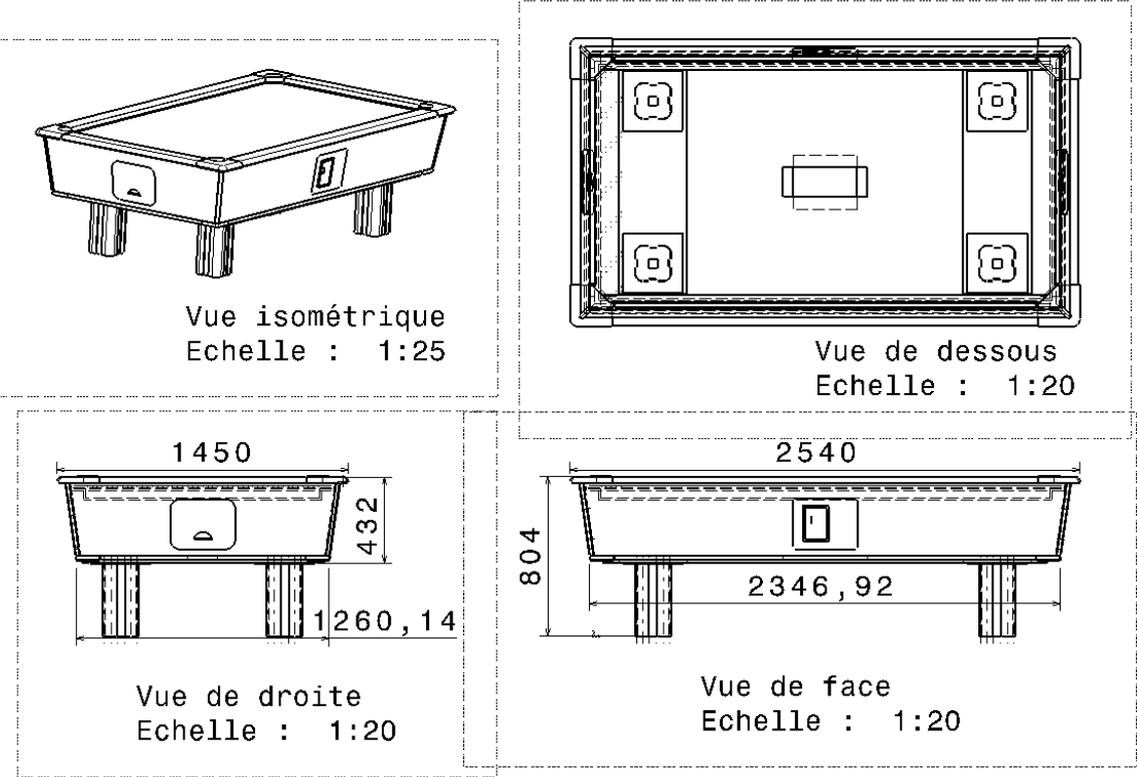


Figure B.1: Dimensions of the air-hockey table [mm].



Figure B.2: Electrical power jacks.

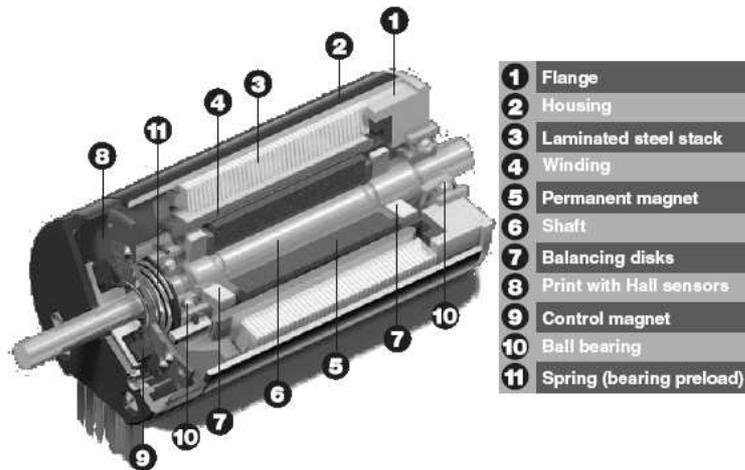


Figure B.3: Scheme of a brushless EC motor (Maxon Motor company).

a lattice of little holes, separated about 4cm from each other. Constant air blowing through the holes provides frictionless motion of the sliding bodies. The robotics experiments have been conducted with 12.5° of table inclination. This provided comfortable cycles frequencies for experimental supervision. Note moreover that the gravity constant g plays no role in the analysis further derived.

Motors and edges

The 1m-long impacting edges were manufactured in aluminium, to be both light and rigid. They are actuated around their lower extremity by two independent DC motors. The motors are mounted on the rigid frame of the table, close to each other (about 10cm) since the model assumes that both edges rotate around the same point. Given the typical velocity and acceleration of the edges trajectory, we estimated the maximum rotational velocity of the motor at 4500RPM, producing a maximal torque of 55mNm (the corresponding power is then 25W). These requirements are met with by the EC40-118896 motor (electronically commuted, see Fig. B.3) + GP42C-203129 planetary gearhead manufactured by the Maxon Motor company (www.maxonmotor.com, Sachseln, Switzerland).

The motor is also equipped with a digital-encoder (incremental), in order to reconstruct its rotational velocity. The encoders are HEDL 5540 (Maxon Motor company).

The primary motor control loop is velocity control. It is insured through two independent servoamplifiers, one for each motor (model DES50/5, 4 quadrants operational, Maxon Motor company, see Fig. B.4). The PI velocity controllers are factory-designed to provide excellent dynamical performances with no static error. The servoamplifier received the digital-encoder signals as inputs and provided the

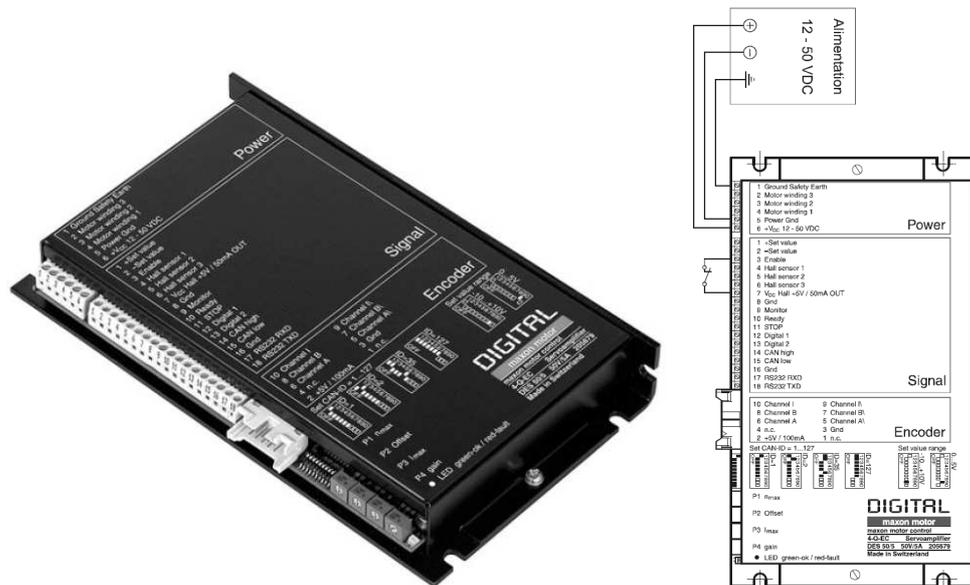


Figure B.4: DES50/5 servoamplifier (Maxon Motor company).

motor voltage as output. The external input, provided to the servoamplifier, is obviously the desired motor velocity.

Position sensor units

In order to realize *position* control of the juggling edges, we superimposed a second loop above each servoamplifier. The position sensors have been custom-made. Their input is the digital-encoder signals, while they outputted the current motor position. The block diagram of the position control loop is given in Fig. B.5. The reference position is derived and directly provided as voltage set value (referenced velocity), through a feedforward (FF) gain (upward path). The feedback part compares the reference with the real measured position and updates the controller output (voltage set value) accordingly, through proportional (P) and derivative (D) gains. No integral gain is necessary to insure zero static error in the output/reference transfer function, since the “plan” is an integrator: the measured output is the position while the plan input is the desired velocity.

Accelerometers

In order to validate closed-loop mechanism described in Chapter 4 with Wiper, we had to detect the impact times. This has been realized with two accelerometers mounted at the top of the edges to record the high-frequency small vibrations generated by the impacts. The accelerometers raw signals were high-pass filtered to remove the accelerations induced by regular edge motion (low frequencies). The

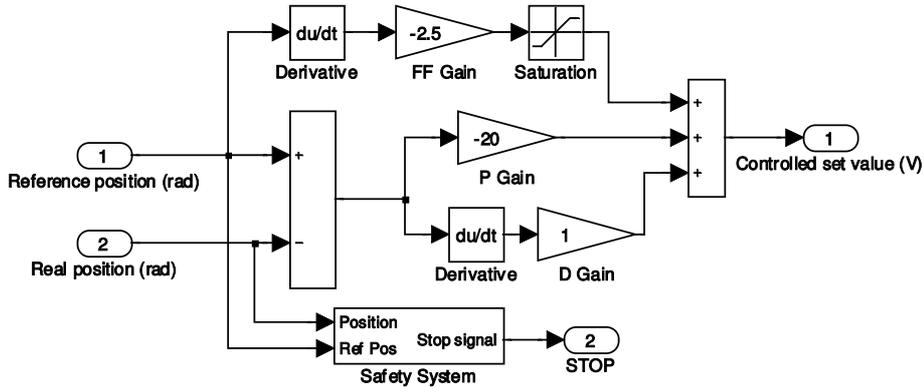


Figure B.5: Block diagram (Simulink) of the position unit.

design has been made on a Chebyshev digital filter (type II) with MATLAB (The Mathworks Inc., Natick, MA). This filter was designed to lose no more than 1dB in the passband [150Hz, . . .] and to have at least 80dB of attenuation in the stopband [0Hz, 1Hz]. The attenuation was smooth between 1Hz and 150Hz, such that the filter order was limited (second order). Consequently, its transfer function was:

$$0.83 \frac{z^2 - 2z + 1}{z^2 - 1.63z + 0.69}.$$

Computer control

The processing of the accelerometers high-pass filter, the motors position controllers, and, more importantly, the planning of the actuation profiles, were implemented in a dedicated computer. These programs were written with XPCTARGET, a Simulink-oriented software (The Mathworks Inc., Natick, MA). This software self-generated the executable code, and uploaded and executed it in another dedicated (target) computer, running the real-time operating system.

B.1.2 Programs architecture

Sensorless control

Open-loop sinusoidal actuation of the arms conducted to sensorless stabilization of periodic orbits in Wiper (see Chapter 6). The Simulink diagram corresponding to this “controller” is represented in Fig. B.6. The sinusoidal signal, at the desired amplitude and frequency, is directly provided as reference to two position controllers (pictured in Fig. B.5): one for each arm. The measured position signals are also provided. Note moreover that several security and protection boxes have been added to provide major damages to the setup, due for example to large differences between the referenced and the actual position. The algorithm sampling frequency was 1kHz (1ms as sampling interval).

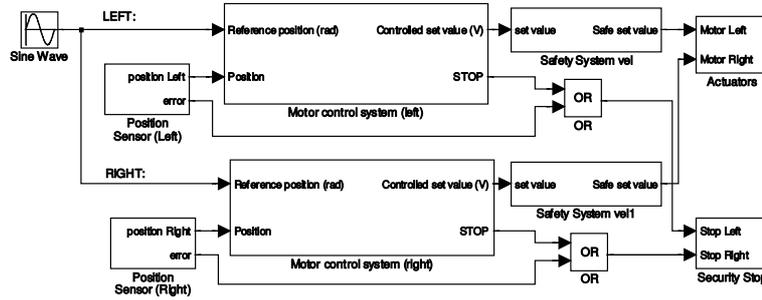


Figure B.6: Block diagram (Simulink) of the sensorless controller.

Feedback control

The closed-loop controller, based on impact times detection, is pictured in Fig. B.7. The figure is divided in 5 colored parts. The gray part is the actual setup, whose inputs are the voltage set values of the motors servoamplifiers, and outputs (in red) are the measured positions of the motors, and accelerometers raw signals. The white part is the position control unit of both edges, already described previously (the security blocks have been removed for clarity). The red part is the signal processing unit of the accelerometers data. It outputs two triggering signals (one for each arm) that rise when an impact is detected. The blue part is the observer (see (7.3)), which provides an estimate of the puck state at impact, on the basis of the motors position signals, and the detected impact times. Finally the green part is the planner, which constructs the reference signals of the tracking controller (7.8), on the basis of the estimated state variables. The algorithm sampling frequency was 333Hz (3ms as sampling interval).

B.2 Technical details about the “human” configuration of Wiper

In the present section, we will give the main details about the technical implementation of Wiper which led to human motor control experiments. These experiments are reported in Chapter 8. A general picture of the setup in that configuration is provided in Fig. 5.2.

The air-hockey table was the same as described in Section B.1.1, also tilted at 12.5° w.r.t. the ground. We replaced the actuated edges by two steel-plain arms. These arms rotated freely around their lower extremity, being axial rotating joints, and the subjects controlled the arms angular position through direct catching. For prehension facility, a thin metallic profile was fixated on the arms top surface. The position of the arms and the puck were recorded by an Optotrak for off-line processing, as reported in Section 8.1. A digital metronome was used to beep the

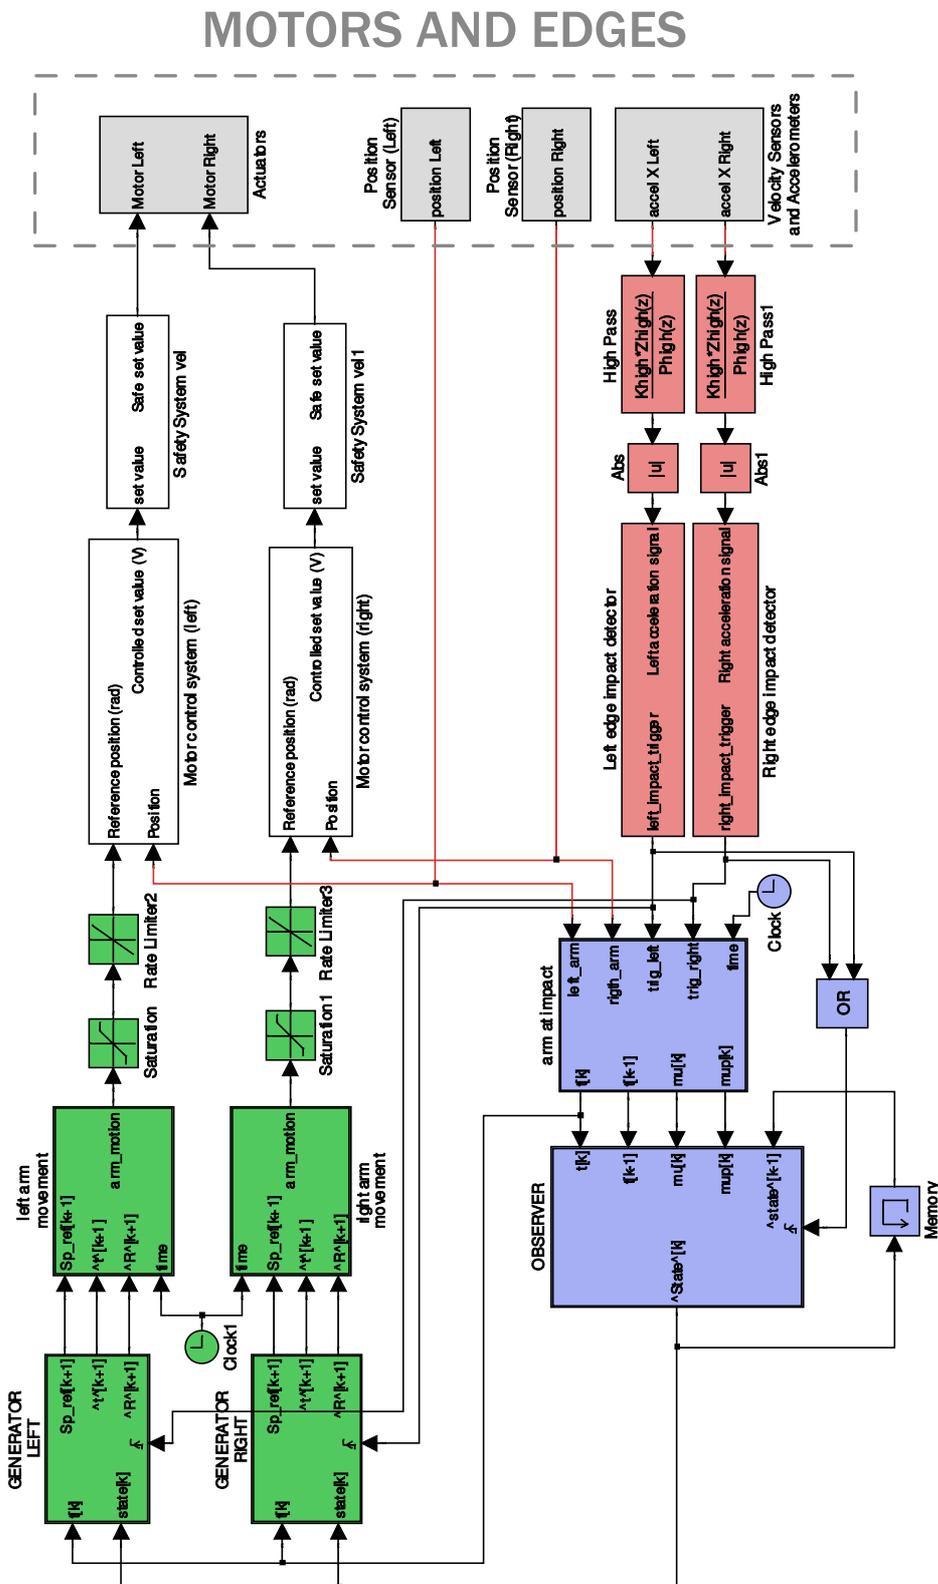


Figure B.7: Block diagram (Simulink) of the closed-loop controller.

requested tempo of the periodic cycles to stabilize for the subjects.

We still used XPCTARGET (The Mathworks Inc., Natick, MA), under the target/host computer configuration previously mentioned, to generate the metronome signal and to synchronize the external acquisition devices. The algorithm sampling frequency was 1kHz (1ms as sampling interval).

Appendix C

Computation of Gaze Orientation Under Unrestrained Head Movements

The soul that can speak through
the eyes, can also kiss with a
gaze.

Gustav Adolfo Becquer

The present appendix describes the computational technique we used to calculate the gaze orientation of human subjects in the data reported in Chapter 8.

C.1 Introduction

The accurate measurement of eye movements is crucial to oculomotor research. These movements are commonly expressed in terms of their horizontal, vertical, and torsional components, assuming a ball-in-socket model for the eye with three degrees of freedom (DOF). There is an enormous body of literature investigating eye movements under the head fixed condition, and describing accurate methods of measuring these rotations. Such methods are mainly based on either the *dual search coil technique* (Robinson, 1963; Collewyn et al., 1985) or *video image processing* devices (e.g. Nakayama, 1974; Haslwanter, 1995; Moore et al., 1996). Both coil- and video-based techniques are widely used by the oculomotor community in behavioral and clinical studies (e.g. Orban de Xivry et al., 2006; Yuksel et al., 2005).

The dual search coil technique is based on the measurement of electric fields induced in a coil placed directly on the subject's eye. The intensity of the electric field in the coil depends on the orientation of the coil (i.e., of the eye) with respect to an alternating magnetic field. Video image processing devices are designed to measure eye movements with camera-based systems. These techniques are more difficult to

apply, however, when gaze orientation arises from a combination of head and eye movements. Video-based systems are usually based on processing images from a camera fixed to the head. It follows that they do not capture any head movement, but only the eye-in-head component of gaze orientation. The coil technique captures the true eye-in-space orientation, on the other hand, but cannot cope with *translations* of the head since the magnetic field in the recording chamber is supposed to be uniform. In addition, *rotations* of the head always induce a translation of the eye center that is not captured by search coil algorithms. This induces a non-linear bias in the computed signals.

In the study described in this appendix, we measure the eye-in-head and head-in-space positions independently. Our technique is therefore valid for very large head movements, or even displacements of the body. Using the coil technique with an unrestrained head would require the measurement of head position, and the integration of this information using an algorithm similar to that presented in this appendix. This study provides a robust geometrical basis for computing the gaze orientation with no restrictions on head movement.

The mathematical developments of this appendix are based on common linear algebra operations. Head rotations are represented as a 3x3 matrix, according to the well-known sequence of Fick angles (Fick, 1874; Haslwanter, 1995), and head translations as a 3x1 position vector. The combination of head position and orientation, hereafter referred to as the head *pose*, therefore has 6 DOF. The eye-in-head orientation is similarly represented by a 3x3 matrix with 3 DOF, since the center of the eye is assumed to be fixed with respect to the head. Similar formalisms have been used by Allison et al. (1996) in testing the vestibular system, and by Wang and Sung (2002) to estimate gaze orientation in front of a computer screen. These translations and rotations could be represented in other ways, e.g. dual-number quaternions. There are even some papers dedicated to comparing these methods, originally from the perspective of robot kinematics (e.g. the survey by Aspragathos and Dimitros, 1998) and later with respect to the computation of eye rotations (Tweed et al., 1990; Haslwanter, 1995). Note that for the sake of simplicity this appendix does not take into account eye-in-head *torsion*, which corresponds to the third Fick angle and captures eye rotation around the optical axis. This angle can be measured by both search coil (e.g. Tweed et al., 1990) and video-based devices (e.g. Moore et al., 1996), however, and can easily be integrated into the eye-in-head orientation matrix. Eye torsion does not change the line of sight.

This appendix also addresses the issue of calibration. We present an efficient calibration protocol based on gaze fixation during self-paced, smooth head movements. This protocol can be easily adapted to a broad range of environments, since it only requires knowing the location of the fixation target in a 3D, ground-based coordinate system.

We will particularly stress the algorithmic sequence followed in computing gaze orientation from the head-in-space and eye-in-head components, via translations and rotations that are simply described using linear algebra. The method is therefore

simple, robust, and computationally efficient. Its main hardware requirement is a device that can measure the position of three points on the subject’s head in a ground-based coordinate system. This appendix also describes experimental results validating our algorithm, obtained by using a video-based device to measure the eye-in-head position. The method essentially consists of two steps:

1. determining the eye orientation in the head coordinate system; and
2. rotating the eye orientation vector into a ground-based coordinate system, using information on the head orientation provided by the head measurement device.

The rest of this appendix is organized as follows. The “Materials and Methods” section (C.2.1) describes the geometrical relationships required to compute the eye-in-head and head components of gaze, and describes their mutual interaction. Section C.2.2 discusses the calibration required to integrate measurements from both acquisition devices. Section C.2.3 describes the validation tasks. The results of validation are summarized in Section C.3, and the characteristics of this new method are discussed in Section C.4.

C.2 Materials and Methods

C.2.1 Geometrical developments

A geometric basis for measurement of eye position under the head fixed condition have been derived by Moore et al. (1996). The center of the eye is assumed to be fixed with respect to the head throughout the task. This section explains how Moore’s method can be extended to unrestrained head conditions, when the subject’s head is free to move without constraint. This technique provides *separate* signals for the head and eye orientations, and therefore distinguishes the relative contributions of each component to gaze orientation.

Gaze is treated as a vector in space, with its origin at the eye center and its direction following from both the head’s orientation in space and the eye’s orientation in the head. According to this method, gaze orientation has to be derived independently for each eye. This appendix therefore focuses on derivations involving one eye, except for a short section dedicated to the calculation of vergence (in Section C.2.1). In the following discussion, matrices are represented by bold uppercase characters (e.g. \mathbf{R}), vectors and points in 3D space by normal uppercase characters (e.g. \mathbf{P}), and scalar quantities by lowercase italic characters (e.g. x).

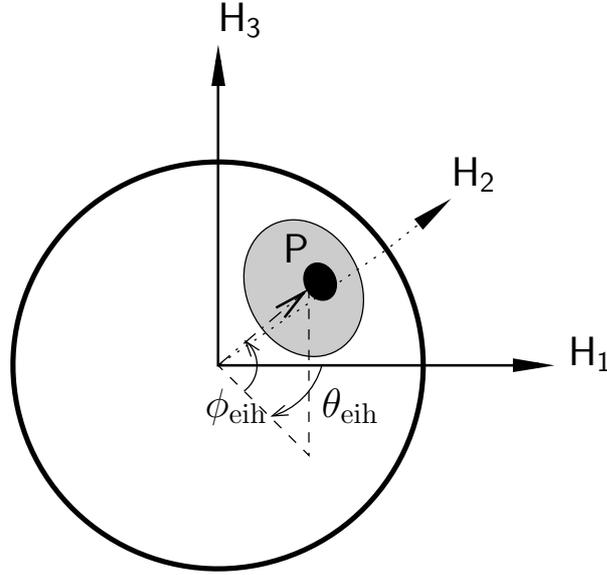


Figure C.1: Diagram of the eye's orientation in the head coordinate system $[H_1, H_2, H_3]$. P denotes the pupil center, and $(\theta_{\text{eih}}, \phi_{\text{eih}})$ are the horizontal and vertical Fick angles of the eye's orientation.

Eye position through image processing

According to Moore et al. (1996), the vertical and horizontal Fick angles of eye-in-head orientation as measured by a video-based device are given by

$$\phi_{\text{eih}} = \arcsin(-a'_{21}x - a'_{22}y - a'_{23}) \quad (\text{C.1})$$

$$\theta_{\text{eih}} = \arcsin\left(\frac{a'_{11}x + a'_{12}y + a'_{13}}{\cos(\phi_{\text{eih}})}\right), \quad (\text{C.2})$$

where the coefficients ('gains') a'_{ij} are determined by calibration (see Section C.2.2) and (x, y) are the coordinates of the pupil's center in the camera image. Both angles are depicted in Fig. C.1.

The main geometrical developments of Moore et al. (1996) are summarized in Appendix C.A of this chapter. The eye orientation *vector* in the head coordinate system follows directly from (C.27).

Head pose through image processing

Head pose is defined in terms of a *ground-based* (i.e., motionless with respect to the laboratory) coordinate system $[G_1, G_2, G_3]$ (see Fig. C.2). To efficiently compute the head pose one must measure the position of three points on the head, which must not be collinear. Let us denote these points by T_a , T_b and T_c . They are represented by grey dots in Fig.s C.2 and C.6, and define a plane parallel to the

frontal plane $\mathbf{H}_2 - \mathbf{H}_3$. Since the head is assumed to be a rigid body, the position of these points completely determines the head pose. It is of particular interest to determine the position of the eye center, i.e., the origin of the gaze. We assume that the position of this point can be deduced from the positions of $\mathbf{T}_{\{a,b,c\}}$ and prior knowledge of the head's anthropomorphic characteristics. The position of the eye center, $\mathbf{E} = (e_1, e_2, e_3)^T$, is taken as the origin of the $[\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3]$ coordinate system as depicted in Fig. C.2.

The head orientation is defined as the orientation of the vector \mathbf{H}_1 with respect to the coordinate system $[\mathbf{G}_1, \mathbf{G}_2, \mathbf{G}_3]$. This unit vector is computed using the cross product (hereafter denoted by \times) of two vectors between different pairs of points in $\mathbf{T}_{\{a,b,c\}}$, e.g.

$$\mathbf{H}_1 = \frac{(\mathbf{T}_c - \mathbf{T}_a) \times (\mathbf{T}_b - \mathbf{T}_a)}{|(\mathbf{T}_c - \mathbf{T}_a) \times (\mathbf{T}_b - \mathbf{T}_a)|}. \quad (\text{C.3})$$

The head orientation angles follow from this vector in a straightforward manner:

$$\theta_h = \arctan\left(\frac{h_{12}}{h_{11}}\right) \quad (\text{C.4})$$

$$\phi_h = -\arcsin(h_{13}), \quad (\text{C.5})$$

where $(h_{11}, h_{12}, h_{13})^T = \mathbf{H}_1$.

The torsional component of the head orientation must also be computed. This does not influence the line of sight, but it does modify the relationship between the eye-in-head $[\mathbf{H}_1, \mathbf{H}_2, \mathbf{H}_3]$ and ground-based $[\mathbf{G}_1, \mathbf{G}_2, \mathbf{G}_3]$ coordinate systems. The head rotation must be computed to know how the two gaze components should be combined. The head axis \mathbf{H}_1 is made parallel to \mathbf{G}_1 by left-multiplying the vectors \mathbf{T} with the following orthogonal rotation matrix:

$$\begin{aligned} & \begin{pmatrix} \cos(\theta_h) \cos(\phi_h) & -\sin(\theta_h) & \cos(\theta_h) \sin(\phi_h) \\ \sin(\theta_h) \cos(\phi_h) & \cos(\theta_h) & \sin(\theta_h) \sin(\phi_h) \\ -\sin(\phi_h) & 0 & \cos(\phi_h) \end{pmatrix}^{-1} \\ = & \begin{pmatrix} \cos(\theta_h) \cos(\phi_h) & -\sin(\theta_h) & \cos(\theta_h) \sin(\phi_h) \\ \sin(\theta_h) \cos(\phi_h) & \cos(\theta_h) & \sin(\theta_h) \sin(\phi_h) \\ -\sin(\phi_h) & 0 & \cos(\phi_h) \end{pmatrix}^T \end{aligned} \quad (\text{C.6})$$

Geometrically, head torsion corresponds to the angle between this rotated vector $(\mathbf{T}_c - \mathbf{T}_b)$, hereafter referred to as the *forehead* vector $\mathbf{F} = (f_1, f_2, f_3)^T$, and the \mathbf{G}_2 axis. This angle is equal to

$$\psi_h = \arctan\left(\frac{\cos(\theta_h) \sin(\phi_h) f_1 + \sin(\theta_h) \sin(\phi_h) f_2 + \cos(\phi_h) f_3}{-\sin(\theta_h) f_1 + \cos(\theta_h) f_2}\right). \quad (\text{C.7})$$

The gaze in space

This section integrates the previously defined eye-in-head and head components of the gaze orientation into a single vector giving the gaze orientation in the $[\mathbf{G}_1, \mathbf{G}_2, \mathbf{G}_3]$

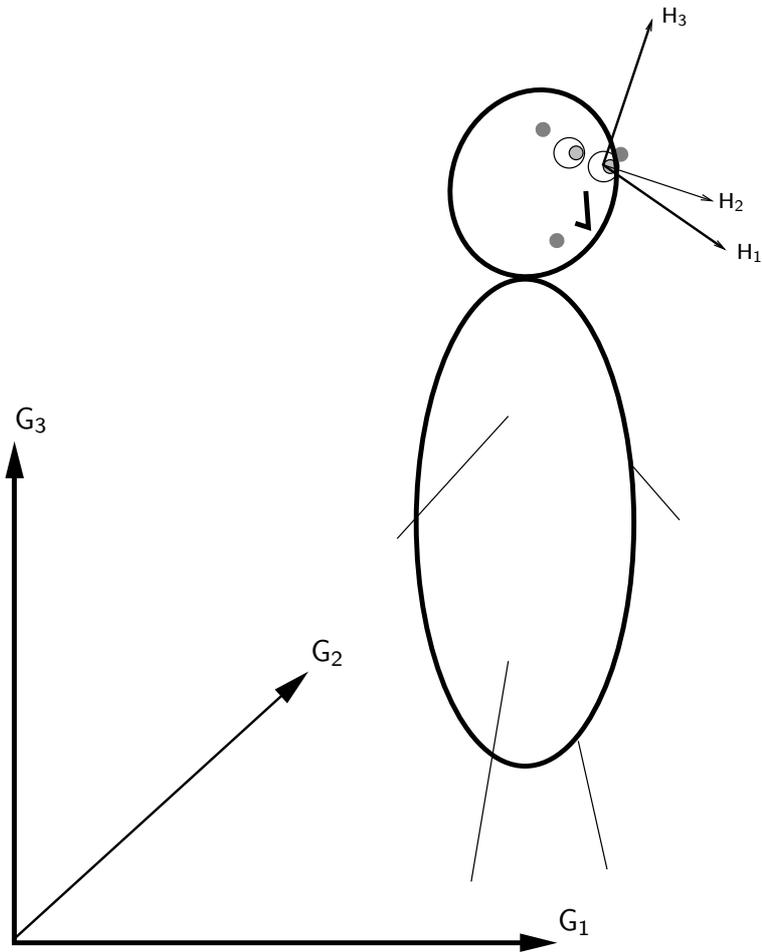


Figure C.2: Diagram of the whole body. This figure emphasizes the components of the gaze orientation: the head component is due to the head pose in the ground-based coordinate system $[G_1, G_2, G_3]$, and the eye-in-head component is due to the eye's orientation in the head coordinate system $[H_1, H_2, H_3]$. The grey dots denote the points T_a , T_b and T_c (see Fig. C.6 for more details), which are measured to determine the head pose.

coordinate system. The origin of this vector is the center of the eye \mathbf{E} . The orientation vector follows directly from the eye-in-head vector \mathbf{P} , and can be obtained by substituting (C.1) and (C.2) into (C.27), then applying three rotations to the result: first by the horizontal angle θ_h , then by the “meridian” angle ϕ_h , and finally by the torsional angle ψ_h obtained from (C.4), (C.5) and (C.7) respectively:

$$\begin{aligned}
 \mathbf{P}_G &= \begin{pmatrix} \cos(\theta_h) & -\sin(\theta_h) & 0 \\ \sin(\theta_h) & \cos(\theta_h) & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \cos(\phi_h) & 0 & \sin(\phi_h) \\ 0 & 1 & 0 \\ -\sin(\phi_h) & 0 & \cos(\phi_h) \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(\psi_h) & -\sin(\psi_h) \\ 0 & \sin(\psi_h) & \cos(\psi_h) \end{pmatrix} \mathbf{P} \\
 &= \begin{pmatrix} \cos(\theta_h)\cos(\phi_h) & \cos(\theta_h)\sin(\phi_h)\sin(\psi_h) - \sin(\theta_h)\cos(\psi_h) & \cos(\theta_h)\sin(\phi_h)\cos(\psi_h) + \sin(\theta_h)\sin(\psi_h) \\ \sin(\theta_h)\cos(\phi_h) & \sin(\theta_h)\sin(\phi_h)\sin(\psi_h) + \cos(\theta_h)\cos(\psi_h) & \sin(\theta_h)\sin(\phi_h)\cos(\psi_h) - \cos(\theta_h)\sin(\psi_h) \\ -\sin(\phi_h) & \cos(\phi_h)\sin(\psi_h) & \cos(\phi_h)\cos(\psi_h) \end{pmatrix} \mathbf{P} \\
 &= \mathbf{R}_h \mathbf{P} \tag{C.8}
 \end{aligned}$$

\mathbf{P}_G therefore denotes the vector from \mathbf{E} to \mathbf{P} in the ground-based coordinate system, and defines the orientation of the line of sight.

Vergence

In general, the lines of sight of both eyes intersect at a specific distance from the subject called the depth of focus. The angle defining the difference between the two lines of sight is called the *vergence*. This point of intersection is easily determined with our method, since the origins (\mathbf{E}_l and \mathbf{E}_r , for the left and right eyes respectively) and directions (\mathbf{P}_{Gl} and \mathbf{P}_{Gr}) of both lines are known. While the two lines of sight should belong to a common plane, measurement noise may skew their apparent orientations. In this case, the best approximation to their point of intersection is halfway along the segment connecting their two closest points¹. This point, denoted \mathbf{P}_{Gv} , is equal to (Goldman, 1990)

$$\mathbf{P}_{Gv} = \frac{(\mathbf{E}_l + \mathbf{P}_{Gl}s_l + \mathbf{E}_r + \mathbf{P}_{Gr}s_r)}{2}, \tag{C.9}$$

where s_l and s_r are given by

$$s_l = \frac{\det(\mathbf{E}_r - \mathbf{E}_l, \mathbf{P}_{Gr}, \mathbf{P}_{Gl} \times \mathbf{P}_{Gr})}{|\mathbf{P}_{Gl} \times \mathbf{P}_{Gr}|^2} \tag{C.10}$$

$$s_r = \frac{\det(\mathbf{E}_r - \mathbf{E}_l, \mathbf{P}_{Gl}, \mathbf{P}_{Gl} \times \mathbf{P}_{Gr})}{|\mathbf{P}_{Gl} \times \mathbf{P}_{Gr}|^2}. \tag{C.11}$$

\mathbf{P}_{Gv} can therefore be interpreted as the subject’s point of interest.

¹Geometrically, this segment is *orthogonal* to both lines of sight.

Summary of the procedure

After calibration (see Section C.2.2) it is possible to compute the angular orientation of the eye in space, given the position of the pupil center $(x, y)^T$ in the head coordinate system and the positions of the points \mathbb{T} , as follows:

- (i) The head orientation angles θ_h , ϕ_h , and ψ_h are computed from the positions of \mathbb{T}_a , \mathbb{T}_b and \mathbb{T}_c using Equations (C.4), (C.5) and (C.7).
- (ii) The position of the eye center \mathbb{E} is computed from the positions of \mathbb{T}_a , \mathbb{T}_b and \mathbb{T}_c and known anthropomorphic parameters.
- (iii) Vertical and horizontal eye angles are computed from (C.1) and (C.2), using coefficients determined through calibration.
- (iv) The eye-in-head orientation vector \mathbb{P} is then determined from these angles and (C.27).
- (v) The eye orientation vector is expressed in terms of the ground coordinate system by applying the rotation matrix (C.8).
- (vi) Optionally, the eye-in-space horizontal and vertical Fick angles of the gaze are computed by the equations

$$\theta_G = \arctan\left(\frac{p_{G2}}{p_{G1}}\right) \quad (\text{C.12})$$

$$\phi_G = -\arcsin(p_{G3}), \quad (\text{C.13})$$

where $(p_{G1}, p_{G2}, p_{G3})^T = \mathbb{P}_G$.

- (vii) The point of intersection between the lines of sight and the angle of vergence can also be computed, according to the equations derived in the corresponding section.

C.2.2 Calibration method

As addressed earlier, the points \mathbb{T} fixed to the head are assumed to be measured in the ground-based coordinate system. We also assume that they define a plane parallel to the frontal plane, such that $\theta_h = \phi_h = \psi_h = 0$ when the subject is looking straight ahead (the primary position). If the points \mathbb{T} cannot be accurately fixed with respect to the subject's head, they can still be calibrated to ensure that $\theta_h = \phi_h = \psi_h = 0$ in the primary position by a method outlined in Appendix C.B of this chapter.

Calibration protocols for the measurement of eye-in-head orientation via pupil detection by image processing devices are well documented in the literature (e.g. Haslwanter and Moore, 1995; Moore et al., 1996; Clarke et al., 2002; Schreiber and

Haslwanter, 2004). They are generally based on a small number of fixations at *known* horizontal and vertical Fick angles in the head coordinate system. The objective of calibration is to identify the unknown coefficients a'_{ij} in (C.1) and (C.2). The torsional component does not need to be calibrated for image processing devices, since the polar cross-correlation technique provides an angular measurement directly. It is only necessary to define an iral reference signature indicating the eye's primary position in the orbit.

This appendix considers tasks in a framework where the head is unrestrained, and it is inconvenient to ask a subject to keep their head fixed during calibration. The identification of fixation targets at known horizontal and vertical positions with respect to the head coordinate system is therefore impractical. To cope with these restrictions, we propose a calibration protocol that integrates the head pose measurement. The subject is asked to move their head while keeping their gaze fixed on a specific point. The results of the calibration have to match the eye displacements generated by this procedure, given the horizontal and vertical Fick angles of the target in the head coordinate system $[H_1, H_2, H_3]$ (Fig. C.2).

In the head coordinate system, the Fick angles of the fixation target C vary with head position as follows:

$$\theta_{\text{tar}} = \arctan\left(\frac{c_{f2}}{c_{f1}}\right) \quad (\text{C.14})$$

$$\phi_{\text{tar}} = -\arcsin(c_{f3}), \quad (\text{C.15})$$

where

$$(c_{f1}, c_{f2}, c_{f3})^T = \mathbf{C}_f = \mathbf{R}_h^T(\mathbf{C} - \mathbf{E}) \quad (\text{C.16})$$

is the position of C in the head coordinate system whose origin is the eye center. The calibration is performed by inverting (C.1) and (C.2), and adding a third equation that represents the translation offset (see Maxwell, 1951; Denavit and Hartenberg, 1955, for references about general homogeneous coordinates):

$$\begin{pmatrix} \sin(\theta_{\text{tar}}) \cos(\phi_{\text{tar}}) \\ -\sin(\phi_{\text{tar}}) \\ 1 \end{pmatrix} = \mathbf{A} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} = \begin{pmatrix} a'_{11} & a'_{12} & a'_{13} \\ a'_{21} & a'_{22} & a'_{23} \\ a'_{31} & a'_{32} & a'_{33} \end{pmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix}. \quad (\text{C.17})$$

The best “gain” matrix \mathbf{A} for this overdetermined system can be determined by any numerical processing software. The computation gives the best-fitting solution matrix, in the least squares sense, to the series of eye signals $(x, y, 1)^T$ and target signals $(\sin(\theta_{\text{tar}}) \cos(\phi_{\text{tar}}), -\sin(\phi_{\text{tar}}), 1)^T$ generated by the head movements. This problem is well-conditioned, since the horizontal and vertical eye positions generated during the calibration task are highly independent. The gains a'_{31} and a'_{32} are identified to 0, and a'_{33} to 1, to agree with the form of the vectors used in (C.17).

The subjects we tested reported no difficulty in maintaining gaze fixation during the calibration task, since the velocity of their head displacements was only about 25 deg/s on average. In this range the gain of smooth pursuit is very close to 1

with negligible phase lag (see Lisberger et al., 1981), validating the accuracy of gaze fixation for the expected behavior.

Summary of the procedure To calibrate the eye-in-head orientation with a video-based unit, the subject is asked to maintain gaze fixation on a known point (C) while moving their head first horizontally and then vertically. A series of pupil center displacements $(x, y)^T$ and head pose movements are thereby generated. The gain matrix \mathbf{A} is obtained through the following procedure:

- (i) If necessary, the points \mathbf{T} are artificially corrected such that $\theta_h = \phi_h = \psi_h = 0$ in the primary position by asking the subject to maintain this position for a few seconds. This procedure is detailed in Appendix C.B.
- (ii) The head Fick angles θ_h , ϕ_h and ψ_h are calculated from the \mathbf{T} positions and Equations (C.4), (C.5) and (C.7).
- (iii) The position of the eye center \mathbf{E} is computed from the positions of \mathbf{T}_a , \mathbf{T}_b and \mathbf{T}_c and known anthropomorphic parameters.
- (iv) The target position in the head coordinate system is computed from (C.16).
- (v) The Fick angles of the target in the head coordinate system are computed by (C.14) and (C.15).
- (vi) The components a'_{ij} of (C.17) are calculated, using θ_{tar} and ϕ_{tar} from the previous step and $(x, y)^T$ values from the eye-in-head measurement device.

This calibration method could be easily adapted to clinical studies of patients with oculomotor disorders who are not able to maintain gaze fixation while moving the head. In this case, a “discretized” version of the calibration task could be proposed: for several steady head positions, the patient would be asked to stabilize their gaze toward the calibration target C. Only steady eye-head orientations could be used to compute the linear regression (C.17).

C.2.3 Experimental validation

Data have been collected on 5 human subjects (S1 is female, and S2-S5 are male) between 24 and 27 years of age (mean 25.4). They provided informed written consent, and reported no history of neurological or musculoskeletal disorder. All had normal vision, either natural or corrected. All but one subject was naive regarding the goals of the experiment. All the procedures conducted were approved by the local ethics committee, in compliance with the Helsinki declaration.

A sketch of the experimental setup is given in Fig. C.3. Initially, the subject was asked to stand upon a cross marked on the ground (point A). He or she had to keep their gaze fixed on an IRED marker (point C) placed on a camera tripod 1.3 m away.

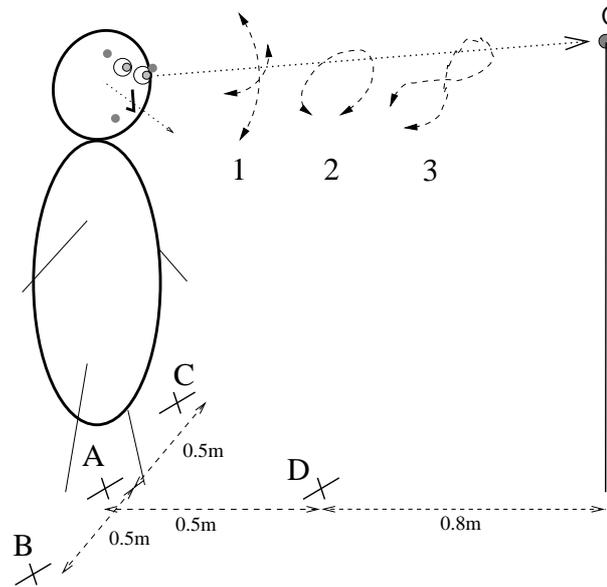


Figure C.3: Experimental setup. Standing on one of the fixed positions (crosses A, B, C and D), the subject is asked to maintain gaze fixation on the grey spot (point C) while moving their head in a cross pattern (1), a circular pattern (2), or an ∞ -shaped pattern (3).

Following the method described in Section C.2.2, each subject was asked to move their head first with a pure horizontal movement, then with a pure vertical movement (Pattern 1 on Fig. C.3). Two other patterns were used to *validate* the calibration so obtained. Pattern 2 consisted of a circular head motion, and was used to validate the *coupling* between horizontal and vertical components. Pattern 3 consisted of a ∞ -shaped head motion, the subject being asked to generate significant head *torsion*. The robustness of the calibration to lateral displacement was evaluated by asking the subject to repeat these patterns while standing on positions B and C (0.5 m to the right and left of point A). The robustness of the calibration to distance was evaluated by repeating the procedure at position D (0.5 m in front of point A). The task was easy to perform under all conditions, resulting in unambiguous gaze orientations. Except for pattern 1 on position A, which was used in calibration, every other pattern and position was used only to validate the method. The method's accuracy is determined by comparing the computed gaze orientation to the actual position of C. The subjects went through the following sequence of patterns: pos. A, pattern 1 (calibration) then patterns 2 and 3; pos. B, patterns 1, 2 and 3; pos. C, patterns 1, 2 and 3; pos. D, patterns 1, 2 and 3; and finally pos. A, pattern 1 once more for a second calibration.

Two-dimensional (horizontal and vertical) recordings of both eyes were made simultaneously using a Chronos head-mounted, video-based eye tracker (Clarke et al.,

2003, CHRONOS VISION GmbH, Berlin, Germany). The calculation of eye positions was based on determination of the pupil center (see Zhu et al., 1999, and references therein). The recording frame rate was 200 Hz. The system is linear in the horizontal and vertical planes for deflections up to ± 20 deg, and has a resolution better than $5'$. System noise was measured to be 0.11 deg for the horizontal plane and 0.09 deg for the vertical plane (Clarke et al., 2002, 2003). A bite-bar was mounted on the helmet frame to prevent slippage between the head and the helmet. This bite-bar was not mandatory, however, and could be removed for subject comfort. In this case the calibration task would be performed at regular intervals to compensate for any slippage between the head and the helmet. The second video-based device used was a 3D position measurement system. The positions of infrared light-emitting diodes (IREDs) on the Chronos helmet and at the fixation target (the four grey dots on Fig. C.3) were measured using an OptoTrak 3020 system (Northern Digital, Ontario, Canada). The OptoTrak was mounted on the ceiling about 3m in front of the subject. The positions of the IREDs were rotated so that they could be expressed in a coordinate system with axes parallel to the floor and centered on the fixation target. The axes G 's of this ground coordinate system are shown in Fig. C.2. The position of each IRED was sampled with a frequency of 200 Hz and resolution of about 0.1 mm within this working environment. The eye signals and IRED signals were filtered at 48 Hz by a zero-phased digital filter (autoregressive, forward and backward). The Chronos eye tracker and the OptoTrak were synchronized by an external computer. Each pattern was executed over a period of 20 s, and its data recorded in a separate file.

C.3 Results

Experimental results are presented in two separate sections. Section C.3.1 describes the calibration results, and Section C.3.2 validates the calibration matrices using data from the patterns described in Section C.2.3.

C.3.1 Calibration

Calibration of the eye-in-head measurement device relies on the algorithm described in Section C.2.2. It assumes that the eye position E is known perfectly with respect to the points T . To keep the method as simple as possible, we assume that this distance is similar for all subjects. Empirical observations of all the subjects wearing the helmet allow us to estimate the eye positions as

$$E = \frac{2B_T + T_c}{3} - 0.09m H_1 + 0.01m F \quad (C.18)$$

for the left eye, and

$$E = \frac{2B_T + T_b}{3} - 0.09m H_1 - 0.01m F \quad (C.19)$$

Table C.1: Mean and standard deviation of the error by this calibration method.

		first calibration				second calibration			
		Left eye [deg]		Right eye [deg]		Left eye [deg]		Right eye [deg]	
		mean	SD	mean	SD	mean	SD	mean	SD
S1	θ	0.06	1.16	0.01	0.37	0.06	1.05	-0.01	0.27
	ϕ	0	0.53	0	0.44	0	0.49	0	0.24
S2	θ	0.02	0.48	0	0.38	0.01	0.35	0	0.35
	ϕ	0	0.65	0	0.73	0	0.43	0	0.43
S3	θ	0.02	0.65	0	0.37	0.04	0.75	0	0.52
	ϕ	0	0.45	0	0.59	-0.01	1.32	-0.01	1.38
S4	θ	0.03	0.99	-0.01	0.65	0.01	0.28	-0.01	0.39
	ϕ	0	0.85	0	0.58	0	0.36	0	0.4
S5	θ	0.09	1.72	-0.01	0.96	0.02	0.43	0.01	0.82
	ϕ	0	0.38	-0.01	1.08	-0.01	1.16	0	0.6

for the right eye. In these equations B_T is the barycenter of the points T (see Appendix C.B), which gives an interocular distance of about 0.07 m. Note that any errors induced by this approximation are corrected to the zeroth and first order by the linear gains in (C.17) during the calibration procedure. A more complicated calibration algorithm could treat these distances as additional unknown parameters to be determined. In this case the calibration would become non-linear, however, and in addition to being less robust would require a more complicated and computationally costly implementation.

To calibrate the helmet unit, the subject executed pattern 1 in position A (see Fig. C.3). From the E position computed by (C.18) or (C.19), the *target* Fick angles are computed according to (C.14) and (C.15). The overall motion is slow enough to assume that the subjects maintained a permanent fixation on the point C. In the calibration data (a 20 s recording) the experimenter manually excluded eye blinks, small saccades, and eye movements outside the detection range to reduce signal distortion in the linear regression (C.17).

The real target angles and eye-in-head angles are compared after calibration in Fig. C.4 (left eye, S2). This diagram emphasizes the sinusoidal motion (alternating pure horizontal and pure vertical movements) of the eye-in-head orientation during the smooth head displacement of the calibration task. Fig. C.4 also illustrates the low frequency of the smooth head movements. The green lines denote the error, i.e., the difference between the target angles and the computed eye-in-head angles. The errors in θ over the test period have an average value of 0.01 deg and a standard deviation of 0.35 deg. For ϕ the average error is 0 deg, and the standard deviation is 0.43 deg. Well-conditioned linear regression of (C.17) gives the following matrix

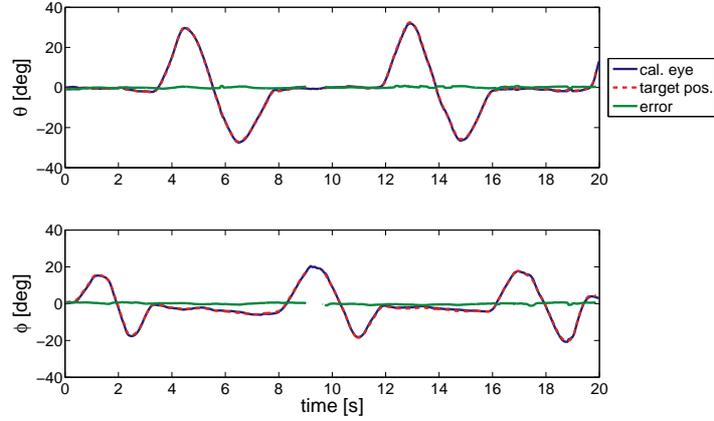


Figure C.4: The target position (dashed red lines) expressed in θ_{tar} (above) and ϕ_{tar} (below) is compared to the computed eye position after calibration (solid blue lines). The solid green lines represent the difference between the two orientations. For eccentric head orientations, in some cases one of the LEDs on the helmet could not be measured (the gaps in the green and red curves around 9s).

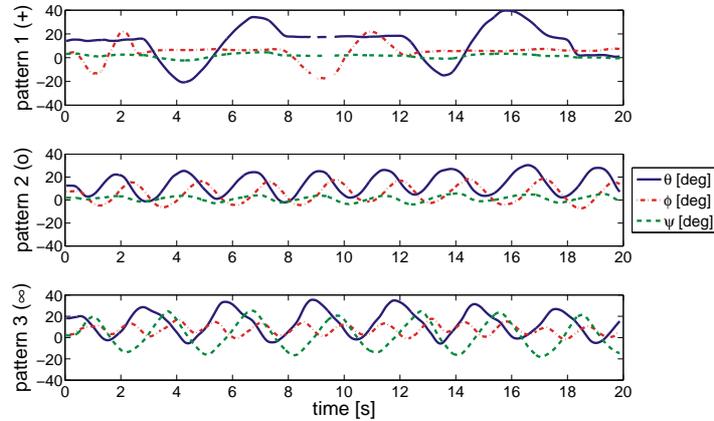


Figure C.5: Typical head motions for the three patterns. Pattern 1 (top) is characterized by alternating pure horizontal (θ , blue solid) and pure vertical (ϕ , red dash-dotted) movements, with no torsion (ψ , green dashed). Pattern 2 (middle) is characterized by coupled, sinusoidal horizontal and vertical movements with a 90 degree phase difference, and limited torsion. In this pattern, the head traces a circle. Pattern 3 (bottom) is characterized by harmonic horizontal and vertical movements in a 2 : 1 frequency ratio (the head follows an ∞ -shaped path), and a significant torsion component. For some head orientations, one of the LEDs on the helmet was not measured (the gaps in the curves).

equation for the orientation of the left eye (subject S2):

$$\begin{pmatrix} \sin(\theta_{\text{tar}}) \cos(\phi_{\text{tar}}) \\ -\sin(\phi_{\text{tar}}) \\ 1 \end{pmatrix} = \begin{pmatrix} 0.0076 & 0.0006 & -0.0005 \\ -0.0002 & 0.0089 & -0.0013 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ 1 \end{pmatrix} \quad (\text{C.20})$$

The average errors and standard deviations of the error, for both eyes and each subject, are listed in Table C.1. The mean error is never significantly different from 0, as should be the case for a linear regression method. The standard deviations, however, certainly reflect a real variability in eye positions (see Ott et al., 1990, and references therein) as well as system noise.

C.3.2 Validation tasks

After performing the calibration task, the subjects were asked to move their heads in other patterns and stand on different predefined points in the workspace. All these tests were to be performed while maintaining their fixation on the same point C (see Section C.2.3). Under these conditions, the known position of C was compared *a posteriori* to the gaze orientation obtained by our algorithm.

Typical head motions for all three patterns are depicted in Fig. C.5. All these data were obtained from subject S2, while standing on point B. The mean value of θ_h is therefore positive (the subject's gaze points to the left), while the mean value of ϕ_h and ψ_h are close to 0 for each pattern. Pattern performances were similar for all subjects, so their validation data were pooled in the analysis.

Several error parameters were computed for each eye, each pattern, and each standing position. $\Delta\theta_{\text{left,right}}$ is the difference between θ_{tar} and the gaze angle θ_G ; $\Delta\phi_{\text{left,right}}$ is the difference between ϕ_{tar} and the gaze angle ϕ_G . We also define the error $\Delta|\alpha_{\text{left,right}}|$ as the absolute value of the angle between the theoretical and the real gaze vectors:

$$\Delta|\alpha| = \left| \arccos \left(\frac{C_f P_G}{|C_f|} \right) \right| \geq 0. \quad (\text{C.21})$$

All these parameters are given in Table C.2, together with the typical ranges of θ_h , ϕ_h , and ψ_h (i.e., the minimum and maximum head angles for each subject).

The average absolute error is always less than 3.5 deg, which is satisfactory with respect to the range of head movements. In particular, the third pattern generates the most important errors. This is due to the relative difficulty of maintaining a fixed gaze while generating the significant head torsion required by this pattern.

As a supplement to this study, several animations of the three patterns have been uploaded to the publisher's web service². They are also available on the first author's web page.

²URL: <http://www.sciencedirect.com> (Ronsse et al., 2007d)

Table C.2: Error parameters for each standing position {A,B,C,D} and head pattern {1,2,3}. The three first columns present the overall range of head angles for each test. The following angular errors for the left and right eyes are averaged across all subjects: $\Delta\theta_{\text{left}}$ is the difference between θ_{tar} and θ_{G} , $\Delta\phi_{\text{left}}$ is the difference between ϕ_{tar} and ϕ_{G} , and $\Delta|\alpha_{\text{left}}|$ is the total angular error as defined by (C.21). All these data are expressed in degrees (deg). The listed ranges are the standard deviations of the angular errors.

		θ_{h}	ϕ_{h}	ψ_{h}	$\Delta\theta_{\text{left}}$		$\Delta\phi_{\text{left}}$		$\Delta \alpha_{\text{left}} $		$\Delta\theta_{\text{right}}$		$\Delta\phi_{\text{right}}$		$\Delta \alpha_{\text{right}} $	
					mean	SD	mean	SD	mean	SD	mean	SD	mean	SD	mean	SD
A	P1	± 35	± 24	± 4	-0.06	0.75	0.02	0.62	0.83	0.52	-0.02	0.58	0.03	0.61	0.72	0.44
	P2	± 28	± 24	± 7	0.83	1.64	0.52	1.8	2.26	1.33	0.27	1.14	0.71	1	1.5	0.79
	P3	± 36	± 23	± 30	0.6	1.69	-0.64	1.78	2.35	1.11	0.09	1.5	-0.5	1.71	2.02	1.17
B	P1	-22, 47	± 22	± 4	0.45	3.47	-0.36	1.08	1.69	3.23	-0.59	1.18	-0.03	0.86	1.32	0.86
	P2	-17, 41	± 24	± 9	0.23	1.09	-0.33	1.01	1.2	0.96	-0.35	1.28	-0.35	0.88	0.97	1.25
	P3	-37, 46	± 22	± 30	0.06	3.47	-0.5	1.56	1.61	3.47	-0.45	1.26	-0.64	1.48	1.72	1.18
C	P1	-50, 23	± 24	± 4	-0.45	0.99	-0.25	1.37	1.43	1.05	-0.58	1.07	-0.12	1.51	1.61	1.09
	P2	-37, 16	± 23	± 9	-0.31	1.2	-1.23	1.43	1.98	1.08	-0.45	1.1	-0.85	1.55	1.91	0.95
	P3	-53, 23	± 25	± 33	-0.12	1.16	-1.15	1.44	1.95	0.98	0.05	1.45	-1.21	1.6	2.22	1.09
D	P1	± 35	± 20	± 4	0.11	0.82	0.4	0.88	1.04	0.72	0.37	0.78	0.45	0.84	1.08	0.69
	P2	± 29	± 25	± 7	0.56	0.91	0.72	1.16	1.59	0.71	0.76	0.66	0.73	1.39	1.75	0.66
	P3	± 34	± 21	± 32	0.49	0.98	1.03	1.54	1.96	0.88	0.81	1.02	0.41	1.35	1.77	0.75
A	P1	± 36	± 24	± 4	0.07	2.31	0.03	0.91	0.76	2.33	-0.04	0.48	0.01	0.62	0.58	0.52

C.4 Discussion

This appendix investigates the geometry of human gaze orientations. To compute gaze orientation in a behavioral context where the head has complete freedom of movement, not only head rotations but also head *translations* must be taken into account. This appendix provides a simple calibration protocol that can integrate both eye-in-head and head-in-space orientations into a single coordinate system.

Our method rests on separate measurements of the head pose and the eye-in-head position. The head pose is determined by measuring the positions of three markers fixed to the head. In the present experiment the eye-in-head position is recorded with a video-based device, but any other classical eye recording method (corneal reflection methods, EOG, etc.) will do as well. The availability of independent head and eye-in-head orientation signals invites further investigation of their interaction. For example, if a search coil is used then eye-in-space angular components are recorded directly but a contribution due to head motion must be subtracted from this signal (see e.g. Crawford and Vilis, 1991; Crawford et al., 1999, 2003; Harris et al., 2001; Han et al., 2005).

Dual search coil and video-based devices are equally popular methods of measuring eye orientation, and each has its specific advantages and drawbacks. On the one hand, *dual search coils* (Robinson, 1963; Collewijn et al., 1985) are expensive and fragile. They are also invasive, since a human subject can endure the presence of a coil on their eye for at most 40 minutes, and require the supervision of an ophthalmologist. Moreover, because the coil measures eye orientation in a ground-based coordinate system it is not straightforward to isolate the eye-in-head component of the gaze orientation. The main advantages of the search coil technique are its excellent spatial accuracy and the high sampling frequencies that can be obtained. With horizontal, vertical, and torsional gaze orientations available in real time, it is easy to implement closed-loop protocols. On the other hand, recent *video image processing* devices have been able to accurately measure eye movements using camera-based systems (see the paper by Clarke et al., 2003, that presents the device used in this study). These techniques are currently approaching the search coil technique in terms of spatial and temporal accuracy. Most importantly, they are much less invasive.

With the development of more advanced hardware, a mathematical formalism based on 3D rotations has been established to relate facial images to gaze orientation (Nakayama, 1974; Tweed and Vilis, 1987; Tweed et al., 1990; Van Opstal, 1993; Haslwanter, 1995). Horizontal and vertical displacements of the eye are computed by tracking the pupil center (see Zhu et al., 1999, and references therein). Torsion may be computed either by tracking natural or artificial landmarks on the eye (Nakayama, 1974; Parker et al., 1985; Ott et al., 1990), or by tracking the iral intensity along a circular sampling path (the *polar cross-correlation* method) (Hatamian and Anderson, 1983; Vieville and Masse, 1987; Tweed et al., 1990; Clarke et al., 1991; Moore et al., 1991; Haslwanter, 1995; Haslwanter and Moore, 1995; Moore

et al., 1996; Haslwanter, 2000; Schreiber and Haslwanter, 2004).

The main limitation of video-based acquisition systems is that they do not capture head motions, since the cameras recording eye motion are fixed to a helmet worn by the subject. This appendix proposes a method of integrating both translations and rotations of the head into any video-based system, based on the measurement of three points on the subject's head. This extra information allows the gaze orientation to be calculated accurately in a ground-based coordinate system. The first part of this appendix (Section C.2.1) covered the geometrical developments of our method. Eye-in-head orientation is first computed using classical techniques (Moore et al., 1996), then translated and rotated into the ground coordinate system once the head pose has been computed. For the sake of simplicity, this appendix did not summarize the existing methods of computing eye *torsion* from video-based devices. Eye torsion does not modify the line of sight, so is not a critical issue within the scope of this appendix.

In addition to the accurate video-based devices described above, we would like to discuss another fruitful aspect of gaze estimation in the literature. Several methods have been developed to estimate the gaze orientation of a head before a computer screen. The goal of such studies is to use gaze as part of the interface between humans and computerized devices, for applications both clinical and otherwise. The main difference between these methods and the technique described in this appendix lies in the hardware; the cameras that track the eye pupil must be fixed to a ground-based coordinate system instead of the subject's head. This has the advantage of measuring gaze orientation directly in the ground-based coordinate system. On the other hand, it also requires continuous monitoring of the distance between the camera and the eye, which is not easy to measure accurately. Even though the most recent of these methods take head translations into account, they cannot compute gaze orientation for very large head displacements (for example, about 1m). Furthermore, these techniques do not achieve the level of spatial and temporal resolution required in oculomotor research. Papers of interest in this field of study include that of Newman et al. (2000), which proposes a method of reconstructing the head pose in real time (30Hz) by tracking typical features of the head; and that of Wang and Sung (2002), which presents a similar method based on linear algebra with homogeneous coordinates. An overview of this topic can be found in recent papers (Shih and Liu, 2004; Yoo and Chung, 2005), as well as the "Computer Vision Bibliography" web page which refers extensively to this literature³. Finally, an alternative approach to determining the point of gaze has been developed by Yu and Eizenman (2004). This method is based on the detection of corresponding points in an image from a head-mounted scene camera and a reference image. This method is not designed to accurately locate the head and other objects in a 3D ground-based coordinate system.

In contrast with the restricted workspace of a computer screen, our method

³<http://iris.usc.edu/Vision-Notes/bibliography/people911.html>

allows not only large head movements but also subject displacements as long as the head markers T remain in the recording space.

In Section C.2.2, we detailed an algorithm for calibrating a video-based eye tracker unit by converting pixel outputs into angular displacements. The mathematical operations involved in this calibration rest on the rotation and translation of vectors and matrices, as was the case for computation of the gaze orientation.

The main advantages of our method can be summarized as follows:

- it is non-invasive for the subjects, since a video-based device is used to compute eye-in-head orientation;
- its algorithm is expressed in terms of linear algebra, and is therefore computationally efficient;
- it separates the eye-in-head and the head-in-space components of gaze, thus allowing the study of their mutual interaction;
- it can be used with a wide variety of acquisition devices, provided that the 3D positions of three points on the head can be independently measured;
- it is based on the geometry of the body, so is easy to understand and implement;
- it is robust under challenging experimental conditions; a preliminary version of this method has already been tested in parabolic flight campaigns by the European Space Agency.

Future improvements to our method mainly concern algorithmic issues: nonlinear estimation of the eye center position with respect to the head markers (this would, however, greatly complicate the calibration algorithm), more accurate pupil center detection (see Zhu et al., 1999), decoupling the horizontal and vertical axes of eye rotation (Schreiber and Haslwanter, 2004), etc.

Section C.3 was devoted to the experimental validation of our method. The calibration process has been strongly validated by data from five different human subjects, confirming that the proposed linear algorithm accurately reflects the dynamics of gaze orientation. Using the calibration matrices so obtained, we also performed other validation tasks. As expected, the smallest angular errors were measured for the first pattern in position A; i.e., for the data using in computing the calibration matrix. The errors remained within reasonable bounds, however, for all head patterns and standing positions. In general, the angular errors obtained in the validation tasks compared favorably to the natural variability of gaze orientation in typical fixation tasks (see Ott et al., 1990, and references therein). The errors incurred by this method are also similar to those obtained by previous experiments comparing the accuracy of video-based and search coil eye tracking techniques in human experiments (e.g. Ott et al., 1990; Moore et al., 1996; Clarke et al., 2002, 2003).

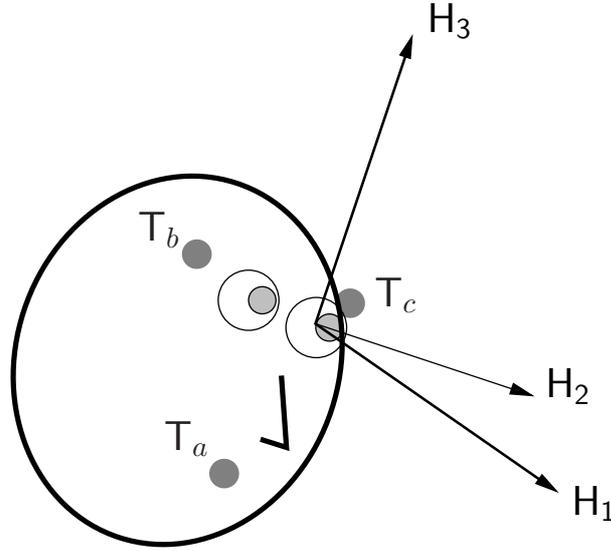


Figure C.6: Diagram of the head in space. The coordinate system $[H_1, H_2, H_3]$ is fixed to the head, with its origin at the center of the (arbitrarily chosen) left eye. This point is also the origin of the gaze orientation vector. The points $T_{\{a,b,c\}}$ are fixed on the head, and discussed in Section C.2.1.

This appendix therefore provides a computationally efficient procedure for computing eye orientation in a ground-based coordinate system. This method can be implemented rapidly in a variety of settings, since it is based on linear transformations related to the actual body configuration. This material has been published as such in Ronsse et al. (2007d).

C.A Eye-in-head orientation

We define an orthogonal, right-handed, head-based coordinate system $[H_1, H_2, H_3]$ with its origin at the center of the eyeball. The H_2 axis is parallel to the interaural axis, and the $H_2 - H_3$ plane is parallel to the frontal plane. The H_1 axis therefore points out of the face from the occiput (see Fig. C.6). The eye's *primary position* is defined as the position where its line of sight corresponds to the H_1 axis. We also define a *camera* coordinate system $[C_1, C_2, C_3]$, where C_2 and C_3 lie within the image plane and C_1 corresponds to the optical axis of the camera.

According to the definition of Moore et al. (1996), the coordinates of the pupil center with respect to the head coordinate system $P = (p_1, p_2, p_3)^T$ and the camera coordinate system $P' = (p'_1, p'_2, p'_3)^T$ are related by

$$P' = \mathbf{R}_{\text{cam}} P + T_{\text{cam}}, \quad (\text{C.22})$$

where T_{cam} and \mathbf{R}_{cam} are the translation vector and rotation matrix relating the

head coordinate system to the camera coordinate system. Angular positions of the eye are expressed in terms of Fick angles (Fick, 1874; Haslwanter, 1995), which are commonly used in oculomotor research. If the space around the eye is viewed as a sphere marked with parallels and meridians, then the sequence of Fick angles defining an orientation is as follows: first a horizontal rotation θ along the equator, then a vertical rotation ϕ along a meridian, and finally a torsional rotation ψ about the optical axis. According to the right-hand rule, eye movements to the left, down, and clockwise are positive from the subject's point of view.

Moore et al. (1996) showed that the horizontal (θ_{cam}) and vertical (ϕ_{cam}) components of the offset matrix \mathbf{R}_{cam} can be captured by calibration gains and do not influence the accuracy of measurement if they remain bounded within 5 deg. This assumption is reasonable, provided the camera is properly fixed to the acquiring device. Equation (C.22) therefore reduces to

$$\mathbf{P}' = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(\psi_{\text{cam}}) & -\sin(\psi_{\text{cam}}) \\ 0 & \sin(\psi_{\text{cam}}) & \cos(\psi_{\text{cam}}) \end{pmatrix} \mathbf{P} + \mathbf{T}_{\text{cam}}, \quad (\text{C.23})$$

where ψ_{cam} denotes an offset rotation of the camera around its optical axis.

If the distance between the lens plane and the center of the eye is large compared to the distance between the lens plane and the image plane of the camera (Nakayama, 1974; Moore et al., 1996), then the projection \mathbf{P}'' of \mathbf{P}' onto the image plane is given by

$$\mathbf{P}'' = \begin{pmatrix} 0 \\ x \\ y \end{pmatrix} = k \begin{pmatrix} 0 \\ p'_2 \\ p'_3 \end{pmatrix} \quad (\text{C.24})$$

where k is a scaling factor related to the image magnification and p'_i are the individual components of \mathbf{P}' . By inverting (C.23) and (C.24), we find the following relation between the actual eye position \mathbf{P} and the pupil center $(x, y)^T$ as measured by the camera:

$$p_2 = a_{11}x + a_{12}y + a_{13} \quad (\text{C.25})$$

$$p_3 = a_{21}x + a_{22}y + a_{23}, \quad (\text{C.26})$$

where the coefficients $a_{11} = a_{22} = \cos(\psi_{\text{cam}})/k$, $a_{12} = -a_{21} = \sin(\psi_{\text{cam}})/k$, $a_{13} = -\cos(\psi_{\text{cam}})t_2 - \sin(\psi_{\text{cam}})t_3$, and $a_{23} = \sin(\psi_{\text{cam}})t_2 - \cos(\psi_{\text{cam}})t_3$ all have to be determined by an appropriate calibration. The scalar variables p_i (t_i , etc.) denote the individual components of \mathbf{P} (\mathbf{T}_{cam} , etc.) respectively.

The pupil center \mathbf{P} follows from applying the Fick rotation $(\theta_{\text{eih}}, \phi_{\text{eih}})$ to the eye's primary position $(r_p, 0, 0)^T$, where r_p is the distance between the center of the eye and the center of the pupil.

$$\mathbf{P} = \begin{pmatrix} \cos(\theta_{\text{eih}})\cos(\phi_{\text{eih}}) & -\sin(\theta_{\text{eih}}) & \cos(\theta_{\text{eih}})\sin(\phi_{\text{eih}}) \\ \sin(\theta_{\text{eih}})\cos(\phi_{\text{eih}}) & \cos(\theta_{\text{eih}}) & \sin(\theta_{\text{eih}})\sin(\phi_{\text{eih}}) \\ -\sin(\phi_{\text{eih}}) & 0 & \cos(\phi_{\text{eih}}) \end{pmatrix} \begin{pmatrix} r_p \\ 0 \\ 0 \end{pmatrix} \quad (\text{C.27})$$

These rotations are depicted in Fig. C.1. Substituting this result into (C.25) and (C.26), we arrive at the eye-in-head orientation:

$$\phi_{\text{eih}} = \arcsin\left(-\frac{a_{21}x + a_{22}y + a_{23}}{r_p}\right) \quad (\text{C.28})$$

$$\theta_{\text{eih}} = \arcsin\left(\frac{a_{11}x + a_{12}y + a_{13}}{r_p \cos(\phi_{\text{eih}})}\right) \quad (\text{C.29})$$

This result gives (C.1) and (C.2), which define $a'_{ij} \triangleq a_{ij}/r_p$.

C.B Calibration of the points T

Our gaze estimation method rests on the fact that the points $T_{\{a,b,c\}}$ define a plane parallel to the frontal plane $H_2 - H_3$ (Fig. C.6), while the forehead vector $F = (T_b - T_c)$ is parallel to H_2 . In theory this ensures that $\theta_h = \phi_h = \psi_h = 0$ whenever the subject stands up and looks straight ahead (the primary position). The placement of these points on the subject's head, however, cannot be accurate enough to validate this assumption. Offset yaw (θ_{off}), pitch (ϕ_{off}), and torsion (ψ_{off}) angles must therefore be measured while a subject is maintaining the primary position. This appendix describes a calibration protocol that compensates for these errors by virtually moving the points T.

Prior to the calibration pattern, the subject is asked to maintain the primary position by looking straight ahead while their head orientation angles, i.e. the offset angles, are recorded. In this reference posture, the real points $T_{\{a,b,c\}}$ are related to the so-called corrected points $T'_{\{a,b,c\}}$ by the transformation $(T_a, T_b, T_c) = \mathbf{R}_{\text{off}}(T'_a, T'_b, T'_c)$, where \mathbf{R}_{off} is a rotation matrix with exactly the same structure as \mathbf{R}_h in (C.8). However, this last relation is only valid when the Fick angles of $T'_{\{a,b,c\}}$ are equal to zero. To put it another way, when the points $T_{\{a,b,c\}}$ are such that their *measured* Fick angles are equal to zero, we have $(T'_a, T'_b, T'_c) = \mathbf{R}_{\text{off}}(T_a, T_b, T_c)$. In this case, \mathbf{R}_{off} is the Fick rotation matrix with angles $-\theta_{\text{off}}$, $-\phi_{\text{off}}$ and $-\psi_{\text{off}}$. In order to make this relation valid for any head orientation, the sets of points $T_{\{a,b,c\}}$ and $T'_{\{a,b,c\}}$ must be pre-multiplied by $\mathbf{R}_h^{-1} = \mathbf{R}_h^T$ (\mathbf{R}_h is an orthogonal rotation matrix):

$$\mathbf{R}_h^T(T'_a, T'_b, T'_c) = \mathbf{R}_{\text{off}}\mathbf{R}_h^T(T_a, T_b, T_c). \quad (\text{C.30})$$

The new set of head points is therefore

$$(T'_a, T'_b, T'_c) = \mathbf{R}_h\mathbf{R}_{\text{off}}\mathbf{R}_h^T(T_a, T_b, T_c). \quad (\text{C.31})$$

Finally, for convenience this rotation is applied around the *barycenter* B_T of $T_{\{a,b,c\}}$, $(T_a + T_b + T_c)/3$, which is then conserved.

$$(T'_a, T'_b, T'_c) = \mathbf{R}_h\mathbf{R}_{\text{off}}\mathbf{R}_h^T(T_a - B_T, T_b - B_T, T_c - B_T) + (B_T, B_T, B_T) \quad (\text{C.32})$$

In summary, the set of points $\mathbb{T}'_{\{a,b,c\}}$ can be simply derived from the measured points $\mathbb{T}_{\{a,b,c\}}$. Their Fick angles are equal to zero when the subject is in the primary position; the plane they define is parallel to $\mathbf{H}_2 - \mathbf{H}_3$, and the forehead vector \mathbf{F} is parallel to \mathbf{H}_2 . The real Fick angles for other head poses are computed using these corrected points, according to (C.4), (C.5) and (C.7). For the sake of simplicity, however, the notation $(')$ specifying the use of corrected points is not used elsewhere in this appendix.

Bibliography

- R. S. Allison, M. Eizenman and B. S. K. Cheung.** Combined head and eye tracking system for dynamic testing of the vestibular system. *IEEE Transactions on Biomedical Engineering*, 43(11): 1073–1082, 1996.
- E. L. Amazeen, P. G. Amazeen, A. A. Post and P. J. Beek.** Timing the selection of information during rhythmic catching. *J Mot Behav*, 31(3): 279–289, 1999.
- N. A. Aspragathos and J. K. Dimitros.** A comparative study of three methods for robot kinematics. *IEEE Transactions on Systems Man and Cybernetics Part B-Cybernetics*, 28(2): 135–145, 1998.
- K. J. Astrom and R. M. Murray.** *Feedback Systems: An Introduction for Scientists and Engineers*. <http://www.cds.caltech.edu/~murray/books/AM05/wiki/index.php>, 2005.
- C. G. Atkeson, J. G. Hale, F. Pollick, M. Riley, S. Kotosaka, S. Schaal, T. Shibata, G. Tevatia, A. Ude, S. Vijayakumar and M. Kawato.** Using humanoid robots to study human behavior. *IEEE Intelligent Systems and Their Applications*, 15(4): 46–55, 2000.
- C. Bapat, S. Sankar and N. Popplewell.** Repeated impacts on a sinusoidally vibrating table. *J Sound Vibration*, 108(1): 99–115, 1986.
- P. J. Beek and A. Lewbel.** The science of juggling. *Scientific American*, 273(5): 92–97, 1995.
- P. J. Beek and M. T. Turvey.** Temporal patterning in cascade juggling. *J Exp Psychol Hum Percept Perform*, 18(4): 934–947, 1992.
- P. J. Beek and A. A. M. van Santvoord.** Learning the cascade juggle - a dynamic-systems analysis. *J Mot Behav*, 24(1): 85–94, 1992.
- N. Bernstein.** *The coordination and regulation of movements*. Pergamon, London, 1967.

- B. Brogliato, M. Mabrouk and A. Zavala-Rio.** On the controllability of linear juggling mechanical systems. *Systems & Control Letters*, 55: 350–367, 2006.
- B. Brogliato and A. Zavala-Rio.** On the control of complementary-slackness juggling mechanical systems. *IEEE Transactions on Automatic Control*, 45(2): 235–246, 2000.
- A. E. Bryson and Y. C. Ho.** *Applied optimal control: optimization, estimation, and control*. Blaisdell, Waltham, MA, 1969.
- J. J. Buchanan, J. H. Park and C. H. Shea.** Target width scaling in a repetitive aiming task: switching between cyclical and discrete units of action. *Exp Brain Res*, 175(4): 710–725, 2006.
- M. Buehler, D. Koditschek and P. Kindlmann.** A one degree of freedom juggler in a two degree of freedom environment. In *IEEE/RSJ Conf. intelligent Systems and Robots*, pages pp. 91–97. Tokyo, Japan, 1988.
- M. Buehler, D. Koditschek and P. Kindlmann.** A family of robot control strategies for intermittent dynamical environments. *IEEE Control Systems Magazine*, 10(2): pp. 16–22, 1990.
- M. Buehler, D. Koditschek and P. Kindlmann.** Planning and control of robotic juggling and catching tasks. *International Journal of Robotics Research*, 13(2): 101–118, 1994.
- A. Buschges.** Sensory control and organization of neural networks mediating coordination of multisegmental organs for locomotion. *J Neurophysiol*, 93(3): 1127–1135, 2005.
- D. E. Chang, R. Ronsse and R. Sepulchre.** Is bounce juggling easier in a parabolic billiard than in a wedge ? In **D. H. van Campen, M. D. Lazurko and W. P. J. M. van den Oever**, editors, *5th Euromech Nonlinear Dynamics Conference*, pages 321–328. Eindhoven, The Netherlands, 2005.
- A. H. Clarke, J. Ditterich, K. Druen, U. Schonfeld and C. Steineke.** Using high frame rate cmos sensors for three-dimensional eye tracking. *Behav Res Methods Instrum Comput*, 34(4): 549–560, 2002.
- A. H. Clarke, J. Ditterich, K. Druen, U. Schonfeld and C. Steineke.** The chronos eye tracker: Description and verification study, 2003.
- A. H. Clarke, W. Teiwes and H. Schrerer.** Videoculography - an alternative method for measurement of three-dimensional eye movements. In **R. Schmid and D. Zambarbieri**, editors, *Oculomotor control and cognitive processes*, pages 431–443. Elsevier, Amsterdam, 1991.

- A. Cohen, S. Rossignol and S. Grillner.** *Neural Control of Rhythmic Movements in Vertebrates.* Wiley, New York, 1988.
- H. Collewyn, J. Van der Steen, L. Ferman and T. C. Jansen.** Human ocular counterroll: assessment of static and dynamic properties from electromagnetic scleral coil recordings. *Exp Brain Res*, 59(1): 185–196, 1985.
- J. J. Collins and S. A. Richmond.** Hard-wired central pattern generators for quadrupedal locomotion. *Biol Cybern*, 71(5): 375–385, 1994.
- S. Collins, A. Ruina, R. Tedrake and M. Wisse.** Efficient bipedal robots based on passive-dynamic walkers. *Science*, 307(5712): 1082–1085, 2005.
- S. H. Collins, M. Wisse and A. Ruina.** A three-dimensional passive-dynamic walking robot with two legs and knees. *International Journal of Robotics Research*, 20(7): 607–615, 2001.
- J. D. Crawford, M. Z. Ceylan, E. M. Klier and D. Guitton.** Three-dimensional eye-head coordination during gaze saccades in the primate. *J Neurophysiol*, 81(4): 1760–1782, 1999.
- J. D. Crawford, J. Martinez-Trujillo and E. M. Klier.** Neural control of three-dimensional eye and head movements. *Curr Opin Neurobiol*, 13(6): 655–662, 2003.
- J. D. Crawford and T. Vilis.** Axes of eye rotation and listing’s law during rotations of the head. *J Neurophysiol*, 65(3): 407–423, 1991.
- A. de Rugy and D. Sternad.** Interaction between discrete and rhythmic movements: reaction time and phase of discrete movement initiation during oscillatory movements. *Brain Res*, 994(2): 160–174, 2003.
- A. de Rugy, K. Wei, H. Muller and D. Sternad.** Actively tracking ‘passive’ stability in a ball bouncing task. *Brain Res*, 982(1): 64–78, 2003.
- J. Denavit and R. S. Hartenberg.** A kinematic notation for the lower-pair mechanism based on matrices. *ASME Journal of Applied Mechanics*, pages 215–221, 1955.
- M. Desmurget and S. Grafton.** Forward modeling allows feedback control for fast reaching movements. *Trends Cogn Sci*, 4(11): 423–431, 2000.
- M. Desmurget, H. Grea, J. S. Grethe, C. Prablanc, G. E. Alexander and S. T. Grafton.** Functional anatomy of nonvisual feedback loops during reaching: a positron emission tomography study. *J Neurosci*, 21(8): 2919–2928, 2001.
- V. Dietz.** Do human bipeds use quadrupedal coordination? *Trends Neurosci*, 25(9): 462–467, 2002.

- T. M. H. Dijkstra, H. Katsumata, A. de Rugy and D. Sternad.** The dialogue between data and model: Passive stability and relaxation behavior in a ball bouncing task. *Nonlinear Studies*, 11(3): 319–344, 2004.
- J. Duysens and H. W. Van de Crommert.** Neural control of locomotion; the central pattern generator from cats to humans. *Gait Posture*, 7(2): 131–141, 1998.
- A. Fick.** Die bewegungen des menschlichen augapfels. *Zeitschrift für rationelle Medizin*, 4: 109–128, 1874.
- T. Flash and N. Hogan.** The coordination of arm movements: an experimentally confirmed mathematical model. *J Neurosci*, 5(7): 1688–1703, 1985.
- G. Franklin, J. D. Powell and A. Emami-Naeini.** *Feedback Control of Dynamic Systems*. Prentice Hall, 5th edition, 2005.
- M. Gerard.** *Stabilisation de tâches rythmiques: étude d'un jongleur plan*. Ph.D. thesis, Université de Liège, 2005.
- M. Gerard and R. Sepulchre.** Stabilization through weak and occasional interactions: a billiard benchmark. In **F. Allgower**, editor, *6th IFAC-Symposium on Nonlinear Control Systems (NOLCOS)*, pages 75–80. IFAC, Stuttgart, Germany, 2004.
- M. Gerard and R. Sepulchre.** Rhythmic stabilization of periodic orbits in a wedge. In **P. Horacek, M. Simandl and P. Zitek**, editors, *IFAC World Congress*. Prague, Czech Republic, 2005.
- R. Goldman.** Intersection of two lines in three-space. In **A. S. Glassner**, editor, *Graphics Gems*, volume I, page 304. Academic Press, San Diego, 1990.
- A. Goswami, B. Thuilot and B. Espiau.** A study of the passive gait of a compass-like biped robot: Symmetry and chaos. *International Journal of Robotics Research*, 17(12): 1282–1301, 1998.
- J. Guckenheimer and P. J. Holmes.** *Nonlinear oscillations, dynamical systems and bifurcations of vector fields*. Springer-Verlag, New York, 1986.
- H. Haken, J. A. Kelso and H. Bunz.** A theoretical model of phase transitions in human hand movements. *Biol Cybern*, 51(5): 347–356, 1985.
- Y. H. Han, A. N. Kumar, M. F. Reschke, J. T. Somers, L. F. Dell'Osso and R. J. Leigh.** Vestibular and non-vestibular contributions to eye movements that compensate for head rotations during viewing of near targets. *Exp Brain Res*, 165(3): 294–304, 2005.
- C. M. Harris and D. M. Wolpert.** Signal-dependent noise determines motor planning. *Nature*, 394(6695): 780–784, 1998.

- L. Harris, K. Beykirch and M. Fetter.** The visual consequences of deviations in the orientation of the axis of rotation of the human vestibulo-ocular reflex. *Vision Res*, 41(25-26): 3271–3281, 2001.
- M. Haruno, D. M. Wolpert and M. Kawato.** Mosaic model for sensorimotor learning and control. *Neural Comput*, 13(10): 2201–2220, 2001.
- T. Haslwanter.** Mathematics of three-dimensional eye rotations. *Vision Res*, 35(12): 1727–1739, 1995.
- T. Haslwanter.** *Computational and Experimental Aspects of Rotatory Eye Movements in Three Dimensions*. Habilitationsschrift, Universität Zürich, 2000.
- T. Haslwanter and S. T. Moore.** A theoretical analysis of three-dimensional eye position measurement using polar cross-correlation. *IEEE Trans Biomed Eng*, 42(11): 1053–1061, 1995.
- M. Hatamian and D. J. Anderson.** Design considerations for a real-time ocular counterroll instrument. *IEEE Trans Biomed Eng*, 30(5): 278–288, 1983.
- H. Hirai and F. Miyazaki.** Dynamic coordination between robots: Self-organized timing selection in a juggling-like ball-passing task. *IEEE Transactions on Systems Man and Cybernetics Part B-Cybernetics*, 36(4): 738–754, 2006.
- A. L. Hodgkin and A. F. Huxley.** A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol*, 117(4): 500–544, 1952.
- P. Holmes, R. J. Full, D. Koditschek and J. Guckenheimer.** The dynamics of legged locomotion: Models, analyses, and challenges. *Siam Review*, 48(2): 207–304, 2006.
- P. J. Holmes.** The dynamics of repeated impacts with a sinusoidally vibrating table. *J Sound Vibration*, 84(2): 173–189, 1982.
- R. Huys.** *Disentangling Perceptual-Motor Learning: assembly and embedding in juggling and rowing*. Ph.D. thesis, Vrije Universiteit Amsterdam, 2004.
- R. Huys and P. J. Beek.** The coupling between point-of-gaze and ball movements in three-ball cascade juggling: the effects of expertise, pattern and tempo. *J Sports Sci*, 20(3): 171–186, 2002.
- R. Huys, A. Daffertshofer and P. J. Beek.** Learning to juggle: on the assembly of functional subsystems into a task-specific dynamical organization. *Biol Cybern*, 88(4): 302–318, 2003.

- V. K. Jirsa and J. A. Kelso.** The excitator as a minimal model for the coordination dynamics of discrete and rhythmic movement generation. *J Mot Behav*, 37(1): 35–51, 2005.
- M. I. Jordan and D. M. Wolpert.** Computational motor control. In **M. Gazzaniga**, editor, *The Cognitive Neurosciences*. MIT Press, Cambridge, MA, 1999.
- J. F. Kalaska, S. H. Scott, P. Cisek and L. E. Sergio.** Cortical control of reaching movements. *Curr Opin Neurobiol*, 7(6): 849–859, 1997.
- H. Katsumata, V. Zatsiorsky and D. Sternad.** Control of ball-racket interactions in rhythmic propulsion of elastic and non-elastic balls. *Exp Brain Res*, 149(1): 17–29, 2003.
- M. Kawato.** Internal models for motor control and trajectory planning. *Curr Opin Neurobiol*, 9(6): 718–727, 1999.
- J. A. Kelso.** Phase transitions and critical behavior in human bimanual coordination. *Am J Physiol*, 246(6 Pt 2): R1000–4, 1984.
- J. A. Kelso, D. L. Southard and D. Goodman.** On the nature of human interlimb coordination. *Science*, 203(4384): 1029–1031, 1979.
- J. A. S. Kelso.** *Dynamic patterns. The self-organization of brain and behavior*. MIT Press, Cambridge, MA, 1995.
- A. D. Kuo.** Energetics of actively powered locomotion using the simplest walking model. *J Biomech Eng*, 124(1): 113–120, 2002a.
- A. D. Kuo.** The relative roles of feedforward and feedback in the control of rhythmic movements. *Motor Control*, 6(2): 129–145, 2002b.
- M. F. Land and P. McLeod.** From eye movements to actions: how batsmen hit the ball. *Nat Neurosci*, 3(12): 1340–1345, 2000.
- H. E. Lehtihet and B. N. Miller.** Numerical study of a billiard in a gravitational field. *Phys. D*, 21(1): 93–104, 1986.
- A. Lewbel.** Research in juggling history, 2002.
- S. G. Lisberger, C. Evinger, G. W. Johanson and A. F. Fuchs.** Relationship between eye acceleration and retinal image velocity during foveal smooth pursuit in man and monkey. *J Neurophysiol*, 46(2): 229–249, 1981.
- K. M. Lynch and C. K. Black.** Recurrence, controllability, and stabilization of juggling. *IEEE Transactions on Robotics and Automation*, 17(2): 113–124, 2001.
- E. Marder.** Motor pattern generation. *Curr Opin Neurobiol*, 10(6): 691–698, 2000.

- E. Marder and D. Bucher.** Central pattern generators and the control of rhythmic movements. *Curr Biol*, 11(23): R986–996, 2001.
- K. Matsuoka.** Sustained oscillations generated by mutually inhibiting neurons with adaptation. *Biol Cybern*, 52(6): 367–376, 1985.
- K. Matsuoka.** Mechanisms of frequency and pattern control in the neural rhythm generators. *Biol Cybern*, 56(5-6): 345–353, 1987.
- E. A. Maxwell.** *General Homogeneous Coordinates in Space of Three Dimensions*. Cambridge Univ. Press, Cambridge, U.K., 1951.
- T. McGeer.** Passive dynamic walking. *International Journal of Robotics Research*, 9(2): 62–82, 1990.
- B. Mehta and S. Schaal.** Forward models in visuomotor control. *J Neurophysiol*, 88(2): 942–953, 2002.
- L. Menini and A. Tornambe.** Control of (otherwise) uncontrollable linear mechanical systems through non-smooth impacts. *Systems & Control Letters*, 49(4): 311–322, 2003.
- R. C. Miall, D. J. Weir, D. M. Wolpert and J. F. Stein.** Is the cerebellum a smith predictor? *J Mot Behav*, 25(3): 203–216, 1993.
- S. T. Moore, I. S. Curthoys and S. G. McCoy.** Vtm—an image-processing system for measuring ocular torsion. *Comput Methods Programs Biomed*, 35(3): 219–230, 1991.
- S. T. Moore, T. Haslwanter, I. S. Curthoys and S. T. Smith.** A geometric basis for measurement of three-dimensional eye position using image processing. *Vision Res*, 36(3): 445–459, 1996.
- K. Nakayama.** Photographic determination of the rotational state of the eye using matrices. *Am J Optom Physiol Opt*, 51(10): 736–742, 1974.
- R. Newman, Y. Matsumoto, S. Rougeaux and A. Zelinsky.** Real-time stereo tracking for head pose and gaze estimation. In **J. L. Crowley**, editor, *4th IEEE Internat. Conf. on Automatic Face and Gesture Recognition*. IEEE, Grenoble, France, 2000.
- J. J. Orban de Xivry, S. J. Bennett, P. Lefevre and G. R. Barnes.** Evidence for synergy between saccades and smooth pursuit during transient target disappearance. *J Neurophysiol*, 95(1): 418–427, 2006.
- D. Ott, F. Gehle and R. Eckmiller.** Video-oculographic measurement of 3-dimensional eye rotations. *J Neurosci Methods*, 35(3): 229–234, 1990.

- J. A. Parker, R. V. Kenyon and L. R. Young.** Measurement of torsion from multitemporal images of the eye using digital signal processing techniques. *IEEE Trans Biomed Eng*, 32(1): 28–36, 1985.
- F. Plestan, J. W. Grizzle, E. R. Westervelt and G. Abba.** Stable walking of a 7-dof biped robot. *IEEE Transactions on Robotics and Automation*, 19(4): 653–668, 2003.
- A. A. Post, A. Daffertshofer and P. J. Beek.** Principal components in three-ball cascade juggling. *Biol Cybern*, 82(2): 143–152, 2000.
- M. H. Raibert.** *Legged robots that balance*. The MIT Press series in artificial intelligence. MIT Press, Cambridge, Mass., 1986.
- A. A. Rizzi and D. E. Koditschek.** Progress in spatial robot juggling. In *IEEE International Conference on Robotics and Automation*, pages 775–780. Nice, France, 1992.
- A. A. Rizzi and D. E. Koditschek.** Further progress in robot juggling : The spatial two-juggle. In *IEEE International Conference on Robotics and Automation*, volume 3, pages 919–924. Atlanta, GA, 1993.
- A. A. Rizzi, L. L. Whitcomb and D. E. Koditschek.** Distributed real-time control of a spatial robot juggler. *Computer*, 25(5): 12–24, 1992.
- D. A. Robinson.** A method of measuring eye movement using a scleral search coil in a magnetic field. *IEEE, Trans Biomed Eng*, BME-10: 137–145, 1963.
- R. Ronsse, P. Lefevre and R. Sepulchre.** Open-loop stabilization of 2d impact juggling. In **F. Allgower**, editor, *6th IFAC-Symposium on Nonlinear Control Systems (NOLCOS)*, pages 1157–1162. IFAC, Stuttgart, Germany, 2004.
- R. Ronsse, P. Lefevre and R. Sepulchre.** Timing feedback control of a rhythmic system. In **E. F. Camacho, P. J. Fleming and S. Yurkovich**, editors, *44th IEEE Conference on Decision and Control and European Control Conference ECC 2005*, pages 6146–6151. IEEE, Seville, Spain, 2005.
- R. Ronsse, P. Lefevre and R. Sepulchre.** Sensorless stabilization of bounce juggling. *IEEE Transactions on Robotics*, 22(1): 147–159, 2006.
- R. Ronsse, P. Lefevre and R. Sepulchre.** Rhythmic feedback control of a blind planar juggler. *IEEE Transactions on Robotics*, in press, 2007a.
- R. Ronsse, P. Lefevre and R. Sepulchre.** Robotics and neuroscience: a rhythmic interaction. *Neural Netw*, submitted, 2007b.

- R. Ronsse and R. Sepulchre.** Feedback control of impact dynamics: the bouncing ball revisited. In **P. Mistra and R. Middleton**, editors, *45th IEEE Conference on Decision and Control*, pages 4807–4812. San Diego, CA, 2006.
- R. Ronsse, J.-L. Thonnard, P. Lefevre and R. Sepulchre.** Trading performance for robustness in rhythmic movements. submitted, 2007c.
- R. Ronsse, O. White and P. Lefevre.** Computation of gaze orientation under unrestrained head movements. *J Neurosci Methods*, 159(1): 158–169, 2007d.
- P. Sabes.** The planning and control of reaching movements. *Curr Opin Neurobiol*, 10: 740–746, 2000.
- U. Saranli, M. Buehler and D. Koditschek.** Rhex - a simple and highly mobile hexapod robot. *International Journal of Robotics Research*, 20(7): 616–631, 2001.
- S. Schaal and C. G. Atkeson.** Open loop stable control strategies for robot juggling. In *IEEE International Conference on Robotics and Automation*, volume 3, pages 913–918, 1993.
- S. Schaal, C. G. Atkeson and D. Sternad.** One-handed juggling: A dynamical approach to a rhythmic movement task. *J Mot Behav*, 28(2): 165–183, 1996.
- S. Schaal and N. Schweighofer.** Computational motor control in humans and robots. *Curr Opin Neurobiol*, 15(6): 675–682, 2005.
- S. Schaal, D. Sternad, R. Osu and M. Kawato.** Rhythmic arm movement is not discrete. *Nat Neurosci*, 7(10): 1136–1143, 2004.
- G. Schoner.** A dynamic theory of coordination of discrete movement. *Biol Cybern*, 63(4): 257–270, 1990.
- G. Schoner and J. A. Kelso.** Dynamic pattern generation in behavioral and neural systems. *Science*, 239(4847): 1513–1520, 1988.
- K. Schreiber and T. Haslwanter.** Improving calibration of 3-d video oculography systems. *IEEE Trans Biomed Eng*, 51(4): 676–679, 2004.
- S. H. Scott.** Optimal feedback control and the neural basis of volitional motor control. *Nat Rev Neurosci*, 5(7): 532–546, 2004.
- S. H. Scott and K. E. Norman.** Computational approaches to motor control and their potential role for interpreting motor dysfunction. *Curr Opin Neurol*, 16(6): 693–698, 2003.
- R. Sepulchre and M. Gerard.** Stabilization of periodic orbits in a wedge billiard. In **D. A. Lawrence and T. Parisini**, editors, *42nd IEEE Conference on Decision and Control*, pages 1568–1573. IEEE, Hawaii, 2003.

- C. E. Shannon.** Scientific aspects of juggling. In **N. J. A. Sloane and A. D. Wyner**, editors, *Claude Elwood Shannon : collected papers*, page 924. IEEE Press, New York, 1993.
- S. W. Shih and J. Liu.** A novel approach to 3-d gaze tracking using stereo cameras. *IEEE Trans Syst Man Cybern B Cybern*, 34(1): 234–245, 2004.
- M. W. Spong.** Bipedal locomotion, robot gymnastics, and robot air hockey: A rapprochement. In *Super-Mechano Systems (SMS'99) Workshop*. Tokyo, Japan, 1999.
- M. W. Spong.** Impact controllability of an air hockey puck. *Systems and Control Letters*, 42(5): 333–345, 2001.
- D. Sternad.** Juggling and bouncing balls: Parallels and differences in dynamic concepts and tools. *International Journal of Sport Psychology*, 30(4): 462–489, 1999.
- D. Sternad, W. J. Dean and S. Schaal.** Interaction of rhythmic and discrete pattern generators in single-joint movements. *Hum Mov Sci*, 19(4): 627–664, 2000.
- D. Sternad, M. Duarte, H. Katsumata and S. Schaal.** Bouncing a ball: tuning into dynamic stability. *J Exp Psychol Hum Percept Perform*, 27(5): 1163–1184, 2001a.
- D. Sternad, M. Duarte, H. Katsumata and S. Schaal.** Dynamics of a bouncing ball in human performance. *Physical Review E*, 6301(1): art. no.–011902, 2001b.
- S. H. Strogatz.** *Sync: The Emerging Science of Spontaneous Order*. Hyperion, New York, 2003.
- S. P. Swinnen.** Intermanual coordination: from behavioural principles to neural-network interactions. *Nat Rev Neurosci*, 3(5): 348–359, 2002.
- S. P. Swinnen and N. Wenderoth.** Two hands, one brain: cognitive neuroscience of bimanual skill. *Trends Cogn Sci*, 8(1): 18–25, 2004.
- M. Tlili, D. Mottet, M. A. Dupuy and B. Pavis.** Stability and phase locking in human soccer juggling. *Neurosci Lett*, 360(1-2): 45–48, 2004.
- E. Todorov.** Optimality principles in sensorimotor control. *Nat Neurosci*, 7(9): 907–915, 2004.
- E. Todorov.** Optimal control theory. In **K. Doya, S. Ishii, A. Pouget and R. P. Rao**, editors, *Bayesian Brain*. MIT Press, Cambridge, MA, 2006.

- E. Todorov and M. I. Jordan.** Optimal feedback control as a theory of motor coordination. *Nat Neurosci*, 5(11): 1226–1235, 2002.
- A. Tornambe.** Modeling and control of impact in mechanical systems: Theory and experimental results. *IEEE Transactions on Automatic Control*, 44(2): 294–309, 1999.
- N. Tufillaro, T. Abbott and J. Reilly.** *An experimental approach to nonlinear dynamics and chaos*. Addison-Wesley studies in nonlinearity. Addison-Wesley, Redwood City, CA, 1992.
- N. B. Tufillaro and A. M. Albano.** Chaotics dynamics of a bouncing ball. *Am J Phys*, 54(10): 939–944, 1985.
- M. T. Turvey.** Coordination. *Am Psychol*, 45(8): 938–953, 1990.
- D. Tweed, W. Cadera and T. Vilis.** Computing three-dimensional eye position quaternions and eye velocity from search coil signals. *Vision Res*, 30(1): 97–110, 1990.
- D. Tweed and T. Vilis.** Implications of rotational kinematics for the oculomotor system in three dimensions. *J Neurophysiol*, 58(4): 832–849, 1987.
- A. M. van Mourik and P. J. Beek.** Discrete and cyclical movements: unified dynamics or separate control? *Acta Psychol (Amst)*, 117(2): 121–138, 2004.
- J. Van Opstal.** Representation of eye position in three dimensions. In **A. Berthoz**, editor, *Multisensory control of movement*, pages 27–41. Oxford University Press, Oxford, 1993.
- A. A. M. van Santvoord and P. J. Beek.** Spatiotemporal variability in cascade juggling. *Acta Psychologica*, 91(2): 131–151, 1996.
- T. Vieville and D. Masse.** Ocular counter-rolling during active head tilting in humans. *Acta Otolaryngol*, 103(3-4): 280–290, 1987.
- T. L. Vincent and A. I. Mees.** Controlling a bouncing ball. *International Journal of Bifurcation and Chaos*, 10(3): 579–592, 2000.
- J. G. Wang and E. Sung.** Study on eye gaze estimation. *IEEE Transactions on Systems Man and Cybernetics Part B-Cybernetics*, 32(3): 332–350, 2002.
- K. Wei, T. M. H. Dijkstra and D. Sternad.** Passive stability, perturbations, and active control in a rhythmic task. In *Society for Neuroscience 36th Annual Meeting*. Atlanta, GA, 2006.

- E. R. Westervelt, G. Buche and J. W. Grizzle.** Experimental validation of a framework for the design of controllers that induce stable walking in planar bipeds. *International Journal of Robotics Research*, 23(6): 559–582, 2004.
- O. White, Y. Bleyenheuft, R. Ronsse, A. M. Wing, P. Lefevre and J.-L. Thonnard.** Altered gravity highlights central pattern generators mechanisms. submitted, 2007.
- D. Wolpert, R. Miall and M. Kawato.** Internal models in the cerebellum. *Trends Cogn. Sci.*, 2(9): 338–347, 1998.
- D. M. Wolpert and Z. Ghahramani.** Computational principles of movement neuroscience. *Nat Neurosci*, 3 Suppl: 1212–1217, 2000.
- D. M. Wolpert and M. Kawato.** Multiple paired forward and inverse models for motor control. *Neural Netw*, 11(7-8): 1317–1329, 1998.
- D. M. Wolpert and R. C. Miall.** Forward models for physiological motor control. *Neural Netw*, 9(8): 1265–1279, 1996.
- D. H. Yoo and M. J. Chung.** A novel non-intrusive eye gaze estimation using cross-ratio under large head motion. *Computer Vision and Image Understanding*, 98(1): 25–51, 2005.
- L. H. Yu and M. Eizenman.** A new methodology for determining point-of-gaze in head-mounted eye tracking systems. *IEEE Transactions on Biomedical Engineering*, 51(10): 1765–1773, 2004.
- D. Yuksel, L. M. Optican and P. Lefevre.** Properties of saccades in duane retraction syndrome. *Invest Ophthalmol Vis Sci*, 46(9): 3144–3151, 2005.
- A. Zavala-Rio and B. Brogliato.** On the control of a one degree-of-freedom juggling robot. *Dynamics and Control*, 9(1): 67–90, 1999.
- A. Zavala-Rio and B. Brogliato.** Direct adaptive control design for one-degree-of-freedom complementary-slackness jugglers. *Automatica*, 37(7): 1117–1123, 2001.
- E. P. Zehr, T. J. Carroll, R. Chua, D. F. Collins, A. Frigon, C. Haridas, S. R. Hundza and A. K. Thompson.** Possible contributions of cpg activity to the control of rhythmic human arm movement. *Can J Physiol Pharmacol*, 82(8-9): 556–568, 2004.
- D. Zhu, S. T. Moore and T. Raphan.** Robust pupil center detection using a curvature algorithm. *Comput Methods Programs Biomed*, 59(3): 145–157, 1999.