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

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: biocompatible 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.

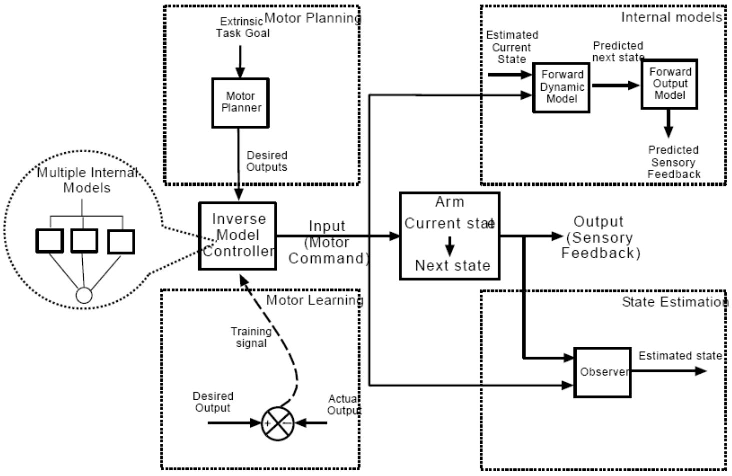

Figure 1.1: The motor system is shown schematically along with related themes of computational motor control. The motor system (center) has inputs — the motor commands — which cause it to change its states and produce an output — the sensory feedback. For clarity not all lines are shown. Reprinted from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

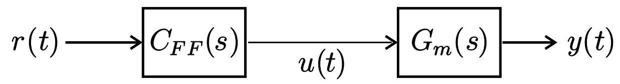

(a) Feedforward control.

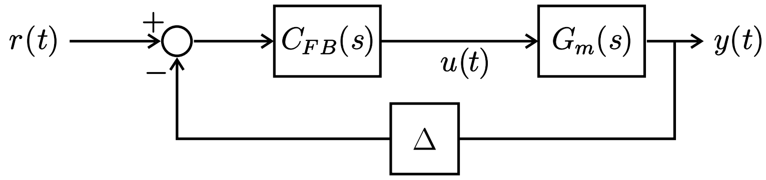

(b) Feedback loop. Δ is a time delay.

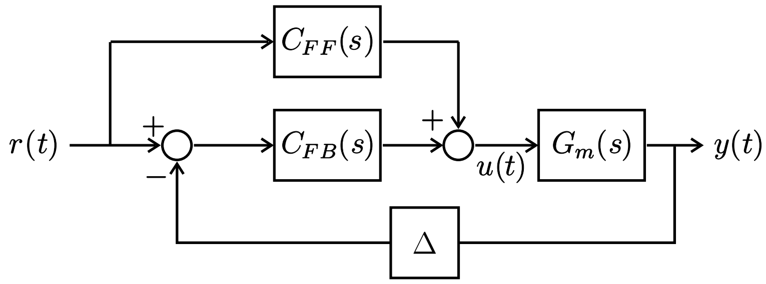

(c) Combined feedforward/feedback loop.

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 reactiveness? ; 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 wedge-billiard. 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) (CPGs[1]) (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

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

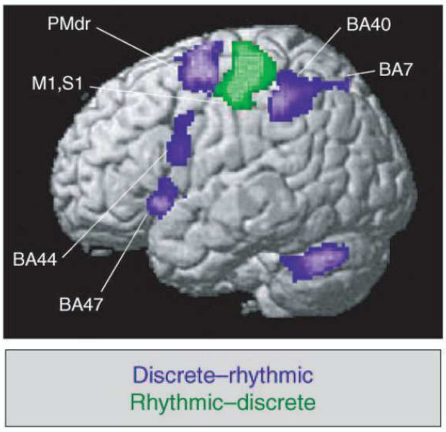

Figure 2.1: In blue, the brain areas involved in the production of a discrete wrist movement, but not in the production of a rhythmic wrist movement: the rostral part of the dorsal premotor cortex (PMdr), Broca's area (BA44), parietal cortex (BA7 and BA40) and the area B47. Widespread activation was 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-

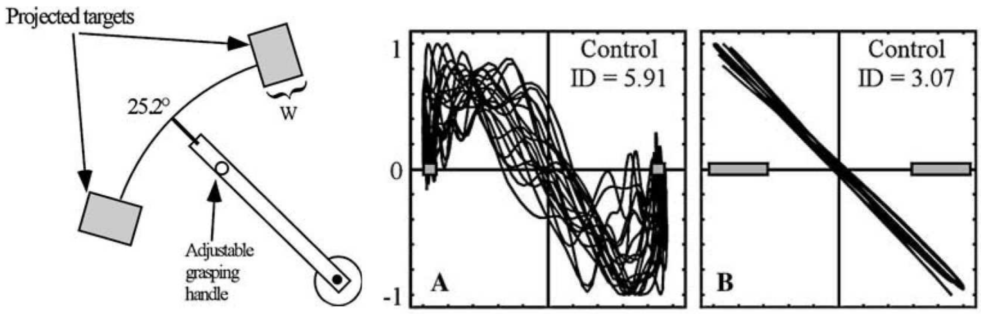

Figure 2.2: The left-hand side picture represents the experimental setup, showing the rotating handle and the targets of width W. The right-hand side panels show two typical phase portraits (position vs. acceleration) of the handle between narrow targets (large index of difficulty — ID = 5.91) and between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

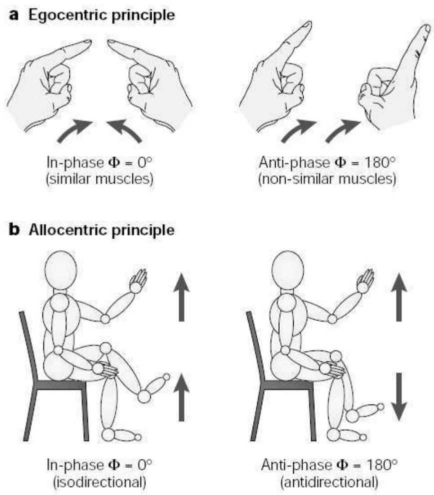

Figure 2.3: Basic coordination constraints: the egocentric and allocentric principles. The egocentric principle refers to a preference for moving according to mirror symmetry, which involves activating similar muscle groups simultaneously (a). The allocentric principle refers to a preference for moving 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:

$$V = -a \cos \phi - b \cos (2\phi) \tag{2.1}$$

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.

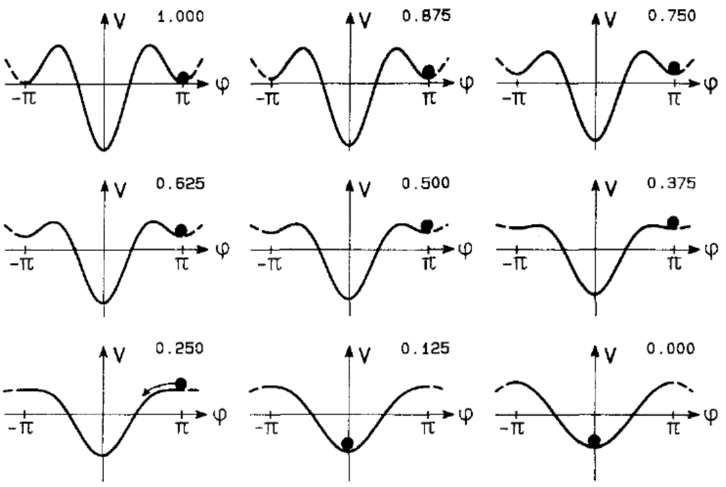

Figure 2.4: The potential V/a for the varying values of b/a, as referred by the numbers. The bottom-left panel corresponds to the critical value, where the anti-phase ($\phi = \pi$ or $-\pi$) loses stability and where the system is "forever" attracted into the in-phase coordination mode. Reprinted from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

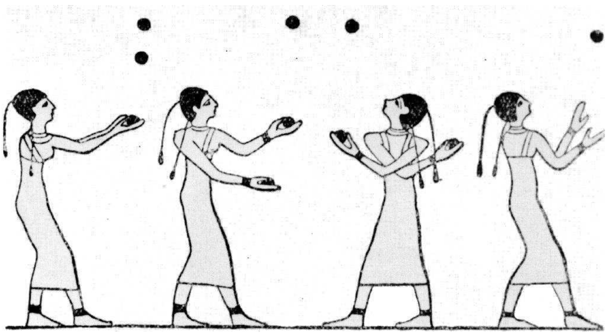

Figure 2.5: One of the earliest representations of juggling, at the ancient Egyptian age. The image source comes from Lewbel (2002).

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. Shannon[2] (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:*

$$(F + D)H = (V + D)N \tag{2.2}$$

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

[2] *Claude 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`.

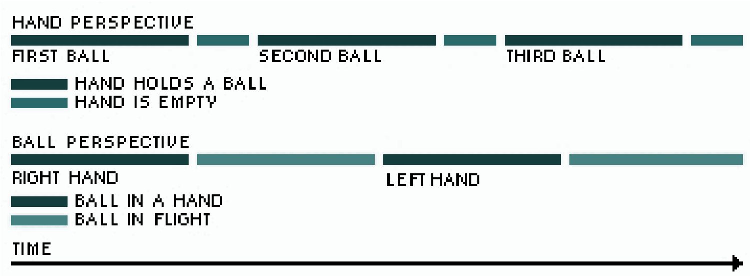

Figure 2.6: The juggling theorem proposed by Shannon (1993) is schematically represented for the three-ball cascade. The theorem is proved by following one complete cycle of the juggle from the point of view of the hand and of the ball and then equating the two. Reprinted from Beek and Lewbel (1995),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

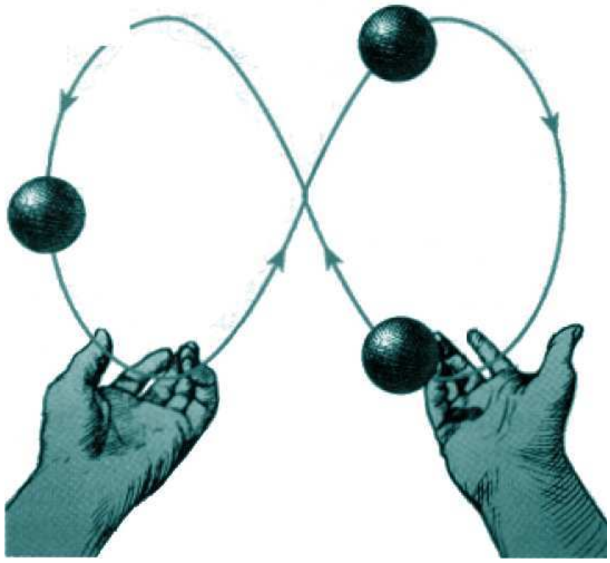

Figure 2.7: The three-balls cascade juggling pattern. Reprinted from Beek and Lewbel (1995).

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.`

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

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

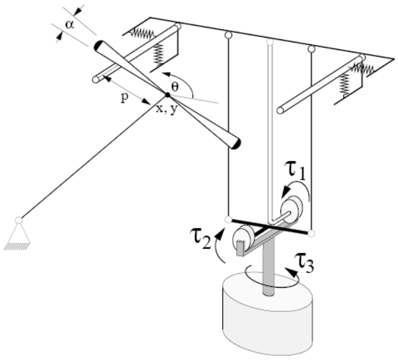

(c) 1991: van Zil's devil sticking robot, from Schaal and Atkeson (1993)

(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

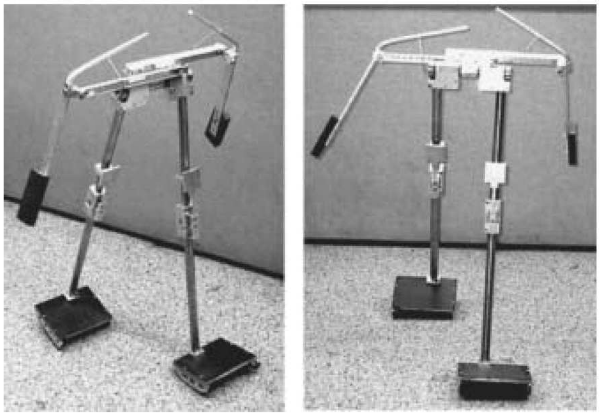

Figure 2.9: Collins, Wisse and Ruina's bipedal passive-walking robot. Reprinted from Collins et al. (2001).

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 juggling, namely *impact* juggling. In impact juggling, the contact between the hand and the object is supposed to be instantaneous[3], 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.

[3]In 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 dynamics[1]. 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:

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

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:

$$v^-(t[k+1]) = v^+(t[k]) - g\,(t[k+1] - t[k]). \tag{3.2}$$

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

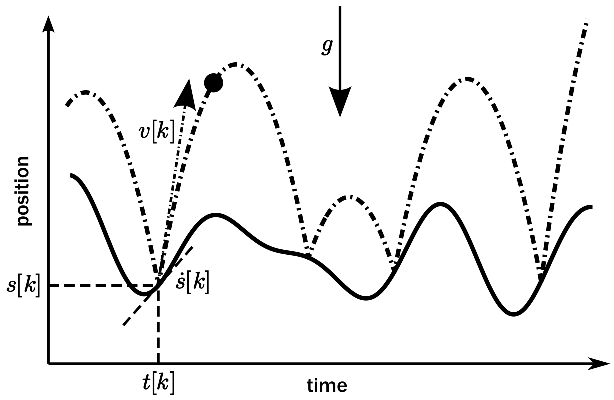

Figure 3.1: 1D bouncing ball. The racket (respectively the ball) trajectory is depicted with solid (respectively dash-dotted) lines over time. At time $t[k]$ (k^{th} impact), the actuator (and ball) position is $s(t[k])$, the actuator velocity is $\dot{s}(t[k])$ and the ball post-impact velocity is $v[k]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

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:

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

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:

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

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

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):

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

3.2.2 Sinusoidal actuation

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

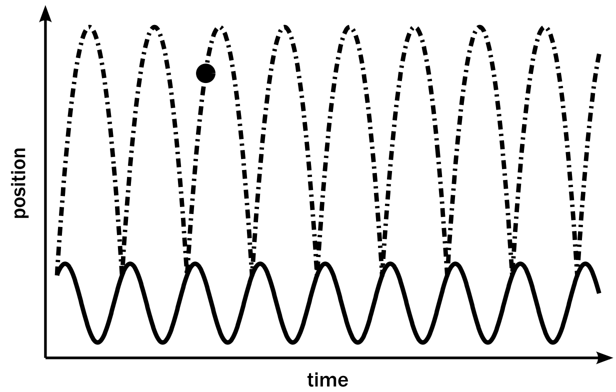

Figure 3.2: Period-one orbit of the 1D bouncing ball under sinusoidal actuation. The racket (respectively the ball) trajectory is depicted with solid (respectively dash-dotted) lines over time.

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):

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

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

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

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

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:

$$\begin{aligned} v^\star &= \frac{n\pi g}{\omega}, & (3.11) \\ \dot{s}^\star &= A\omega\cos\phi^\star \\ &= \frac{1-e}{1+e}\frac{n\pi g}{\omega} & (3.12) \end{aligned}$$

and the steady-state impact phase is equal to:

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

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

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:

$$\begin{aligned} \delta t[k+1] &= \delta t[k] + \frac{1+e}{g}\delta v[k], & (3.14) \\ \delta v[k+1] &= e^2\delta v[k] - (1+e)A\omega^2\sqrt{1-\left(\frac{(1-e)\pi n g}{(1+e)A\omega^2}\right)^2}\,\delta t[k+1] & (3.15) \end{aligned}$$

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:

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

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):

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

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.

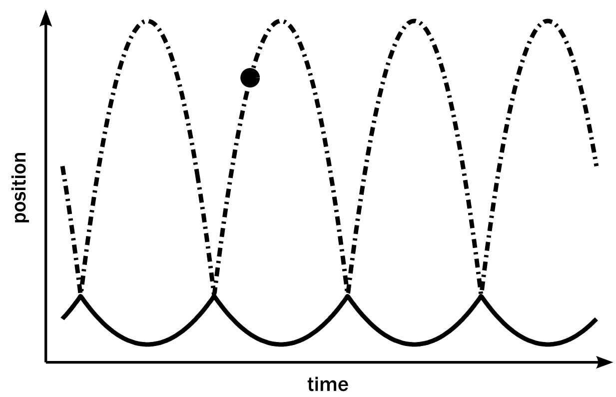

Figure 3.3: Period-one orbit of the 1D bouncing ball under mirror law control. The racket (respectively the ball) trajectory is depicted with solid (respectively dash-dotted) lines over time.

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:

$$s(t) = \frac{-(1-e)}{1+e}\beta(t) - \kappa_1 \left(E^\star_\rho - E(t)\right)\beta(t). \tag{3.18}$$

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: $E^\star_\rho = gs^\star_\rho + 0.5(v^\star_\rho)^2$, 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 $E(t) = E^\star_\rho$, 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.

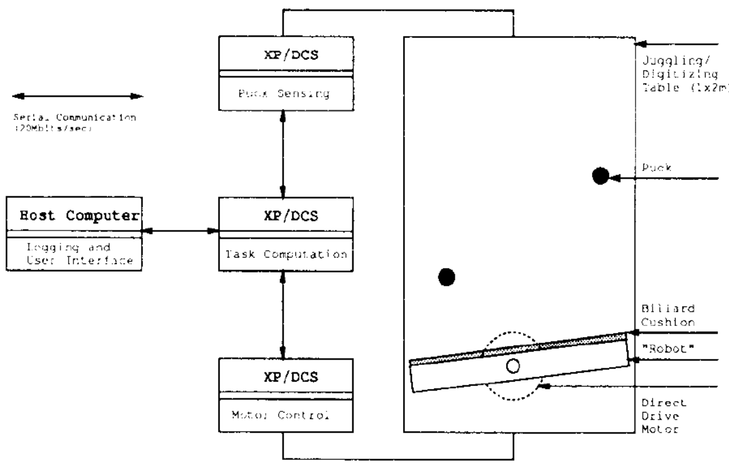

Figure 3.4: The planar juggler with simultaneous juggling of two pucks. Reprinted from Buehler et al. (1994).

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 $u = \omega - \omega^\star$ (where $\omega^\star$ denotes the steady-state value) and is computed on the basis of the measured quantities, i.e. the impact phase $x_1 = \phi - \phi^\star$ and the post-impact velocity $x_2 = v - v^\star$. 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

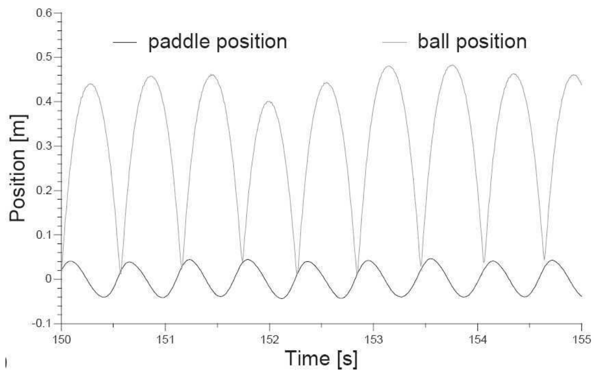

Figure 3.5: Typical strategy of a human subject playing the 1D bouncing ball task with a racket (paddle). Reprinted from Schaal et al. (1996).

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 $\ddot{s}^\star = -A\omega^2 \sin\phi^\star$:

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

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

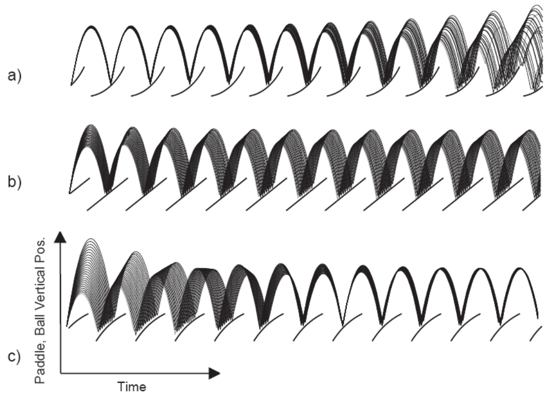

Figure 3.6: Illustration of the role of the acceleration at impact for open-loop stability of the period-one bouncing ball. 25 balls started at the same position but at different initial velocity. The racket trajectory was accelerating (a), at constant velocity (b) or decelerating (c) at impact. Reprinted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

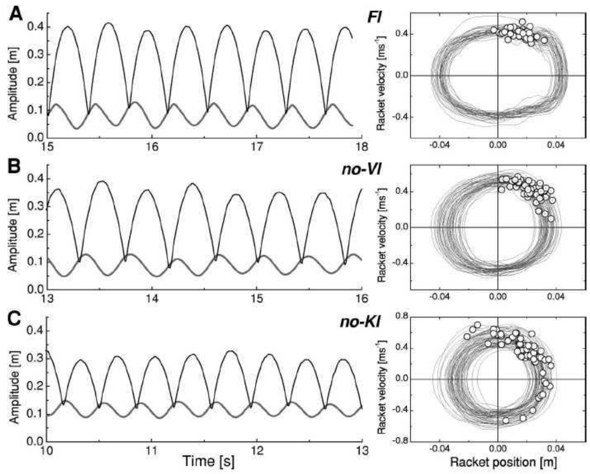

Figure 3.7: Three time series (left panels) and their respective phase portraits (right panels) of exemplary trials performed in the three perceptual conditions: both visual and haptic feedback available (top), no visual feedback (middle) and no haptic feedback (bottom). The dots in the phase portraits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

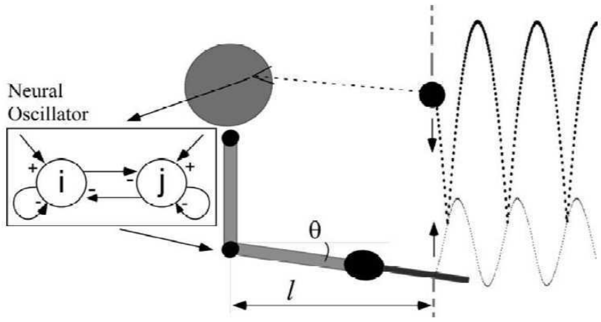

Figure 3.8: Neural-based model of a ball-bouncing controller. Reprinted from de Rugy et al. (2003).

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 another 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.