

Metastable Legged-Robot Locomotion

by

Katie Byl

Submitted to the Department of Mechanical Engineering
in partial fulfillment of the requirements for the degree of

Doctor of Philosophy in Mechanical Engineering

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

September 2008

© Massachusetts Institute of Technology 2008. All rights reserved.

Author
Department of Mechanical Engineering
August 28, 2008

Certified by
Russell L. Tedrake
Assistant Professor, Department of Electrical Engineering and Computer
Science
Thesis Supervisor

Certified by
Neville Hogan
Professor, Department of Mechanical Engineering and Department of
Brain and Cognitive Sciences
Thesis Chair

Accepted by
Lallit Anand
Chairman, Department Committee on Graduate Students

Metastable Legged-Robot Locomotion

by

Katie Byl

Submitted to the Department of Mechanical Engineering
on August 28, 2008, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy in Mechanical Engineering

Abstract

A variety of impressive approaches to legged locomotion exist; however, the science of legged robotics is still far from demonstrating a solution which performs with a level of flexibility, reliability and careful foot placement that would enable practical locomotion on the variety of rough and intermittent terrain humans negotiate with ease on a regular basis. In this thesis, we strive toward this particular goal by developing a methodology for designing control algorithms for moving a legged robot across such terrain in a qualitatively satisfying manner, without falling down very often. We feel the definition of a meaningful metric for legged locomotion is a useful goal in and of itself. Specifically, the mean first-passage time (MFPT), also called the mean time to failure (MTTF), is an intuitively practical cost function to optimize for a legged robot, and we present the reader with a systematic, mathematical process for obtaining estimates of this MFPT metric.

Of particular significance, our models of walking on stochastically rough terrain generally result in dynamics with a fast mixing time, where initial conditions are largely "forgotten" within 1 to 3 steps. Additionally, we can often find a near-optimal solution for motion planning using only a short time-horizon look-ahead. Although we openly recognize that there are important classes of optimization problems for which long-term planning is required to avoid "running into a dead end" (or off of a cliff!), we demonstrate that many classes of rough terrain can in fact be successfully negotiated with a surprisingly high level of long-term reliability by selecting the short-sighted motion with the greatest probability of success. The methods used throughout have direct relevance to machine learning, providing a physics-based approach to reduce state space dimensionality and mathematical tools to obtain a scalar metric quantifying performance of the resulting reduced-order system.

Thesis Supervisor: Russell L. Tedrake

Title: Assistant Professor, Department of Electrical Engineering and Computer Science

Thesis Chair: Neville Hogan

Title: Professor, Department of Mechanical Engineering and Department of Brain and Cognitive Sciences

Acknowledgments

First and foremost, thanks to my thesis advisor, Russ Tedrake. Russ' mentorship and encouragement have been invaluable during my three years in the Robot Locomotion Group at CSAIL. Russ has been a role model in elucidating the key ideas from the details of results. He has also directed our LittleDog quadruped research throughout Phases 1 and 2 of the DARPA Learning Locomotion work at MIT. In particular, the decision to take a higher-risk route by designing more dynamic motions for LittleDog throughout Phase 2 paid off both in achieving program metrics and in making my own research with LittleDog a lot more fun and interesting.

The input from each of the members of my committee has also contributed significantly toward reorganization of the material in my thesis, all for the better. Neville Hogan has provided insightful guidance over the years as I searched for a thesis topic. His advice on my thesis has significantly improved the focus, tying together the individual pieces of work throughout. As co-PI with Russ on LittleDog, Nick Roy also gets my sincere thanks for the opportunity to work on this project and has also given important encouragement and advice both inside and outside of committee meetings. His editing suggestions have been thorough and have improved readability throughout. Finally, Dan Frey has consistently provided feedback on related work on optimization techniques and has encouraged further research directions in which to extend the work presented here.

Many of my fellow labmates in the Robot Locomotion Group and neighboring research groups at MIT have contributed useful suggestions and perspectives through the years, and all have my appreciation for putting up with intermittent periods of intense focus. Alec Shkolnik, Sam Prentice, Khash Rohanimanesh, Steve Proulx, Olivier Chatot, Michael Price, John Roberts, Emma Brunskill, and Lauren White have all contributed directly to our LittleDog efforts at MIT, and Abe Bachrach provided useful, late-night suggestions toward some of the more ambitious motions we planned with the robot. Olivier contributed significantly to our initial testing of various double-support motions – most particularly pacing – during the summer of 2007. Sam and Khash each endured tenures managing our code repository, and Alec developed our original “dynamic lunge”. All three have

devoted countless hours and sleepless nights toward our success with LittleDog. Thanks also to those who have worked to implement real-world hardware inspired by some of the simulations throughout this thesis: in particular, to Tim Villabona, Rick Cory and Thrish Nanayakkara for work building and obtaining motion capture data for realworld models of the rimless wheel and to Fumiya Iida, Steve Proulx and Ian Manchester for implementing a working version of the compass gait on a boom, to walk on rough terrain. Thanks also for encouragement and advice from the rest of our research group, past and present, including Vanessa Hsu, Elena Glassman, Zack Jackowski, A. J. Meyer, Mario Bollini, and Arlis Reynolds. And finally, thanks to Nira Manokharan for help organizing various food, supplies and travel reimbursements over the years.

Thanks for the support of The Infrastructure Group (TIG), who help everything run smoothly at the Computer Science and Artificial Intelligence Lab (CSAIL) and have dealt with various network, software and hardware issues relating to LittleDog with expedience and grace. Thanks in particular to Ron Wiken, who keeps the machine shop at CSAIL well-stocked and well-running; that is a resource which has saved many a roboticist countless hours over the years!

Finally, everyone knows immediate family put up with a lot when anyone is finishing a thesis; thanks to you all. Special thanks to my husband, Marty, who actually put up with me as a labmate writing a Masters thesis before putting up with me as a spouse writing a PhD thesis. Hopefully, I'll eventually start pulling my weight at home, particularly in cooking and in earning a non-student income. Special thanks also to our dog, Murphy: noone I know better illustrates the idea of legged locomotion with "exceptional performance most of the time and only occasional failures (falling)". We got Murphy as a young pup within about a week of the time I started studying legged locomotion in Russ' lab, and so he has been the subject of extensive (though rather informal) observations on path planning and gait adaptation on rough terrain throughout this research.

Contents

1	Introduction and Motivation	17
1.1	Motivations for practical legged robots	21
1.1.1	Anthropomorphic fascination	22
1.1.2	Understanding human walking	23
1.1.3	Legged-robot locomotion in real-world environments	24
1.1.4	So, why have legs?	25
1.2	State-of-the-art in legged locomotion	25
1.2.1	ASIMO	26
1.2.2	Passive-dynamic based walkers	28
1.2.3	RHex	29
1.2.4	BigDog	30
1.2.5	Summary of state of the art	31
1.3	Roadmap of thesis	32
2	Metrics and Methods Toward Highly Dynamic Locomotion	35
2.1	On metrics for legged locomotion	36
2.1.1	Literature review on metrics for walking	37
2.1.2	Static stability margin (SSM)	37
2.1.3	Zero-moment point (ZMP) margin	38
2.1.4	Velocity- and energy-based metrics	45
2.2	The mean first-passage time (MFPT) metric	47
2.3	Metastability: Beyond stability	49
2.3.1	Walking as a metastable limit cycle	50

2.4	Discussion and conclusions	51
2.4.1	Efficiency of walking	51
2.4.2	The right answer?	52
3	Kinodynamic Planning for LittleDog	53
3.1	Introduction: Control solutions for locomotion	53
3.2	Toward dynamic maneuvers on rough terrain	56
3.2.1	Compliance versus stiffness	57
3.2.2	Current state-of-the-art in kinodynamic planning	58
3.3	The Learning Locomotion program	59
3.3.1	LittleDog hardware and environment	59
3.3.2	Metrics for LittleDog	60
3.4	Locomotion planning on rough terrain	61
3.4.1	Path planning	62
3.4.2	Motion planning strategy	66
3.5	Crawl gaits	66
3.5.1	Statically-stable walking	69
3.5.2	Dynamic, ZMP-based walking	69
3.5.3	Design methodology	72
3.5.4	Results and discussion	74
3.6	Double-support walking	74
3.7	Double-support lunging and pacing motions	75
3.7.1	Modeling double-support motions	77
3.7.2	Design of double-support trajectories	80
3.7.3	Result summary for 3 double-support motions	81
3.7.4	Diagonal trot-walk	83
3.7.5	Dynamic lunge	85
3.7.6	Pacing motions	86
3.7.7	Conclusions	92
3.8	Future work	93

3.8.1	Piecewise motion planning on rough terrain	94
3.8.2	Toward dynamic bounding across rough terrain	94
3.9	Summary and conclusions	98
4	Metastable Legged Locomotion	101
4.1	Approach	101
4.2	Metastable walking	102
4.2.1	Suggested reading on metastability	103
4.2.2	Metastable limit cycle analysis	103
4.3	Numerical modeling results	107
4.3.1	Rimless wheel	107
4.3.2	Passive compass gait walker	114
4.4	Discussion	118
4.4.1	Impacts on control design	122
4.4.2	Implementation issues	122
4.4.3	Multiple stable limit cycles	125
4.5	Conclusions	125
5	Compass Gait Model on Rough Terrain	127
5.1	Introduction	127
5.2	Related work for CG models on rough terrain	128
5.3	Stability versus agility for compass gait walking	130
5.4	Passive biped model on rough terrain	131
5.4.1	Passive compass gait model	131
5.4.2	Methods	132
5.4.3	Discussion of passive performance	135
5.5	Optimal control of the biped on rough terrain	136
5.5.1	Actuated compass gait model	136
5.5.2	Method for approximate optimal control	137
5.6	Controlled CG on wrapping rough terrain	142
5.6.1	Wrapping rough terrain models	143

5.6.2	Results on wrapping terrain	144
5.7	Controlled CG on stochastically rough terrain	149
5.8	One-step policy on wrapping terrain	154
5.9	Discussion	155
5.9.1	Policy interpolation	157
5.9.2	Efficiency versus stability	159
5.10	Conclusions	159
6	Toward Metastable Bounding for LittleDog	161
6.1	Bounding gait development	162
6.2	Experimental results	162
6.2.1	Rimless wheel motion trials	163
6.2.2	Triple-rocking trials	164
6.2.3	Open-loop continuous bounding	166
6.3	Discussion	169
7	Conclusions and Future Work	171
7.1	Contributions	171
7.1.1	Kinodynamic motion planning in an underactuated regime	172
7.1.2	Stochastic methods to quantify walking stability	173
7.1.3	Policy optimization for CG walker on rough terrain	174
7.2	Suggested directions for future work	175
7.2.1	LittleDog: Ongoing work	175
7.2.2	Verification of compass gait results on a real robot	177
7.2.3	Efficiency of passive walking on rough terrain	178
7.3	Implications for development of highly dynamic robots	179
A	Pendulum Foot Model	181
B	Compass Gait Implementation Details	185

List of Figures

1-1	Wheels versus legs on rough terrain	18
1-2	ASIMO	27
1-3	Passive-dynamic based walkers	29
1-4	RHex hexapod negotiating rough terrain.	29
1-5	The BigDog robot, by Boston Dynamics.	30
2-1	Inverted pendulum (2D) balancing on a weighted “foot” platform	40
2-2	Minimalist fully-actuated legged robot	42
2-3	Cartoon of a particle subject to Brownian motion in a potential $U(x)$ with two metastable states, A and B	50
3-1	The LittleDog robot	54
3-2	LittleDog executing a dynamic lunge	56
3-3	Global path planning on rough terrain	65
3-4	Center of pressure during level-ground walking	68
3-5	LittleDog walking on pegs	69
3-6	Planar ZMP model for fully-actuated motion	71
3-7	Rocky terrain, negotiated with a dynamic gait	71
3-8	COM vs. ZMP during pacing	73
3-9	COM vs. ZMP during trot-walking	74
3-10	The inverse (left) and forward (right) problems ZMP trajectory.	76
3-11	Predicted and observed pitch during underactuation in trot-walk	76
3-12	Planar model of LittleDog in a double-support lunge.	78
3-13	Simulated dynamics for a planar dynamic lunge.	82

3-14	Snapshot of double-support during dynamic, “diagonal” walk	83
3-15	Four terrain types using a diagonal trot-walk	84
3-16	Comparison of experimental results and theory for a dynamic lunge	85
3-17	LittleDog lunges across Jersey barrier (left) and gap (right)	86
3-18	Climbing steps through sequential dynamic lunging	87
3-19	LittleDog employing a dynamic lunge to clear a tall obstacle.	88
3-20	Repeatability of long, open-loop playback of various double-support motions	90
3-21	Demonstration of accuracy in positioning pacing steps.	91
3-22	Sensitivity of pitch angle, α , to initial conditions (x_o)	92
3-23	LittleDog climbs onto rough terrain	95
3-24	Dynamic lunge executed with feet at varying initial heights	96
3-25	Dynamic lunge initiated on rough rocks	97
3-26	Pitch data for dynamic lunge on rough terrain	98
4-1	A toy example of a Markov chain and corresponding transition matrix.	105
4-2	The rimless wheel (RW) (left) and compass gait (CG) walker (right) models	107
4-3	Return map and fixed point for an 8-spoke rimless wheel on flat, downhill terrain with a constant slope of 8 degrees	109
4-4	Return distribution and metastable “neighborhood” for an 8-spoke rimless wheel on downhill terrain with a mean step-to-step slope of 8 degrees and standard deviation of 1.5 degrees	109
4-5	3D view of the return distribution for the rimless wheel system	110
4-6	Mean first-passage time as a function of the initial condition, ω_o	111
4-7	Quasi-stationary probability density functions for the stochastic rimless wheel for each of several values of terrain noise, σ	112
4-8	3D view of the metastable “neighborhood” of state-to-state transitions	113
4-9	Mean first-passage time (MFPT) for the rimless wheel, as a function of terrain variation, σ	114
4-10	Mean first-passage time for a passive compass gait, as a function of terrain variation	115

4-11	Basin of attraction versus a map of MFPT	117
4-12	Basins of attraction (blue regions) and fixed point for two compass gait walkers, each on terrain of constant slope (4 deg)	119
4-13	Metastable neighborhoods for a passive compass gait walker	120
4-14	Metastable system: Contours of the stochastic “basin of attraction”	121
4-15	Controlled compass gait walker, with torque at the hip	124
5-1	The compass gait biped model.	131
5-2	Deterministic basin of attraction (top) and stochastic map (below) of mean first-passage time	134
5-3	Actuated compass gait model with torque source at hip.	136
5-4	Illustration of meshing approximation	138
5-5	Compass gait on wrapping terrain	143
5-6	Intermittent-foothold terrain.	144
5-7	Examples of terrain which were successfully negotiated using PD control alone.	145
5-8	A close-up of the bottommost terrain in Figure 5-7.	146
5-9	Foothold patterns from optimal control.	147
5-10	Footsteps taken using impulsive toe-off control, only.	148
5-11	Dynamic states using impulsive toe-off control, only.	149
5-12	Footsteps taken during a 60-second trial using the optimal control policy from value iteration (top) and using a one-step time horizon (bottom).	150
5-13	Dynamic states during a using the optimal control policy from value itera- tion (top) and using the one-step control strategy (bottom).	150
5-14	Comparison of step width and step height for the optimal control solution vs. for one-step control.	151
5-15	Smoothness of cost function over state space.	151
5-16	Compass gait on stochastic terrain	152
5-17	MFPT for one-step vs. no lookahead on stochastic terrain	154
5-18	One-step policy shows convergence on wrapping terrain.	156

5-19	Policy interpolation	158
6-1	LittleDog video frames during rocking	162
6-2	Data for rimless-wheel rocking trials	164
6-3	Stochastic transition map for rimless-wheel rocking trials	165
6-4	Data for metastable bounding, with stabilizing pauses	166
6-5	Zoom view of pitch	167
6-6	Apex-to-apex data for initial rocking motion	167
6-7	Data during continuous bounding	168
7-1	Dynamic motions for LittleDog.	173
7-2	LittleDog lunging onto rough terrain.	176
7-3	Compass gait robot posed on rough terrain.	178
A-1	Inverted pendulum (2D) balancing on a weighted “foot” platform	181
B-1	Four states of the CG walker during continuous dynamics	185
B-2	Compass gait meshing parameters for stochastic terrain	186
B-3	Compass gait meshing parameters for wrapping terrain	186

List of Tables

3.1	DARPA metrics for the Learning Locomotion project	61
5.1	MFPT for three different passive CG models	135
B.1	Meshing for compass gait model on stochastic terrain	188
B.2	Meshing for compass gait model on wrapping terrain	189

Chapter 1

Introduction and Motivation

This thesis addresses the interrelated problems of controlling and evaluating machines that locomote with legs. Significant interest in walking robots developed in the late 1970's and throughout the 1980's, as advances in microprocessor technology made the vision of autonomous robotics seem a likely (if not inevitable) direction for future research. Today, over two decades later, there are many exciting examples of walking machine prototypes, but the full promise of reliable legged locomotion has yet to be achieved.

Wheeled vehicles require far less sophisticated coordination of motion to achieve locomotion and can handle a wide variety of terrains. Recent examples of successful wheeled autonomous robots include many of the robot cars and trucks participating in the Grand Challenge competitions, sponsored by DARPA¹, and the Roomba home cleaning robot appliance, developed and produced by iRobot. Of particular note on rough terrain is the "rocker-bogie" design used recently by NASA for the Pathfinder robot. The rocker-bogie essentially consists of six independently-driven wheels. Three wheels on either side are arranged in a clever, rocking differential structure to allow the vehicle to negotiate obstacles on the order of the wheel diameter while balancing the body of the vehicle (to avoid toppling). Capability and stability on extreme terrain come at the price of relatively low traversal speeds however (to minimize dynamic effects).

Furthering the development of *legged* robot locomotion seems prudent only if it has obvious potential to provide clear advantages over wheeled locomotion in some practi-

¹Defense Advanced Research Projects Agency



(a)



(b)

Figure 1-1: Wheels versus legs on rough terrain

The rocker-bogie wheel design used on the Sojourner robot ((a) left) and on the larger Mars Exploration Rover design ((a) right) can both negotiate impressively rough terrain. Using this wheeled design requires significantly slower (and less nimble) traversal than many legged animals can achieve; the jumping mountain goat (at bottom) provides an excellent proof of the performance that legs can attain on extreme rough terrain. Upper image from Wikipedia. This image is copyrighted by the NASA/Caltech Jet Propulsion Laboratory. The JPL allows anyone to use it for any purpose, provided that the copyright notice Courtesy NASA/JPL-Caltech is displayed. Lower image shows goat jumping in Glacier National Park; photo taken by Cristal Jones, 2008.

cal regimes. A typical justification often given for legged locomotion is to enable better flexibility on rough terrain. Two additional, convincing motivations for developing legged machines are (1) to provide sophisticated tools (e.g., for rehabilitation), prosthetics and/or robot “partners” who can interact seamlessly in the same environment in which humans work, live and play and (2) to better understand the science of animal and human locomotion itself.

An additional advantage of using legs for robot locomotion which is seldom mentioned specifically, however, is that legs allow for more instantaneous control than wheels, since we can push off and/or brake very effectively during the intermittent contact with the ground at each step. In this thesis, we find that effective control solutions for walking on rough terrain often require only a short lookahead on the order of a step or two along the terrain, and we speculate that this is in large part due to the inherent ability legs provide for fast, effective energy transfer at each step. The problem still remains of *how* to exploit this potential advantage; finding solutions to this control problem is a motivating focus through this thesis.

To develop the science of walking robotics toward flexible locomotion on rough terrain, our first goal should be to agree upon what constitute qualitatively and quantitatively “good” performance. Clearly these definitions will be colored by the specific type of task or terrain a robot must negotiate, but there are some fundamentally common goals, as well. For example, we can say that we wish for our robots to be highly reliable, in the sense of not falling down often, and that we would like them to operate with sufficient repeatability (e.g., with behavior that agrees with an underlying, scientific model) that we can provide some estimate of their performance on the types of terrain which humans typically negotiate with ease.

A central premise throughout this thesis is that traditional, *deterministic* stability metrics and margins are not well-suited to the problems of control and evaluation in walking. Compared with traditional robotics, walking involves significant underactuation and stochasticity, and we should expect significant progress in legged machines only after we can appropriately *quantify* their performance. We contend that appropriate metrics to capture walking performance should inherently be based on *stochastic* analyses. A significant

contribution of this thesis is the adaptation of well-established concepts and mathematical tools used in other domains toward discussing and evaluating the performance of legged machines.

The author strongly believes that this formal introduction of appropriate mathematical language and definitions of stochastic stability can only aide in furthering the development of legged robots beyond the level of particular case examples to a more formal branch of engineering science, i.e., to advance legged robotics beyond the realm of the hobbyist to that of a more clean science, where results can be reproduced and built upon.

Robot locomotion is inherently subject to much more significant levels of both underactuation and stochasticity than are traditional robot manipulation tasks, such as machining or factory assembly in a manufacturing setting. Unlike a robot leg, a factory arm can be rigidly mounted to the ground. This eliminates the underactuated limitation of unidirectional (push-but-not-pull) forces at the “base” of the robot: instead of a moveable foot, a factory robot can simply be bolted to the ground. A factory arm is still underactuated in its grasping contacts with objects to be manipulated. However, unlike the contact of a foot with the ground, one can create a force closure (i.e., a *wrench*) when grasping an object. These two factors both result in a strong tendency toward underactuation in legged locomotion when compared to factory assembly. There is also typically more stochasticity in legged robotics than in factory robotics, because a practical factory robot is contained in a more controlled environment. These two factors (underactuation and stochasticity) make the field particularly challenging and the use of stochastic tools particularly appropriate in its analysis.

To quantify the effects of stochasticity on walking stability, we develop mathematical tools to evaluate a walking system as a discrete Markov process: approximated as discrete in state space, to produce a transition matrix of the dynamics naturally discretized in time by the impact nature of taking each, particular step. We propose that “good” walking systems have what we call *metastable limit cycle* behaviors. Metastability will be discussed more formally in Section 2.3, but for now, consider metastable to mean “long-living, but destined to eventually end”. An eigenanalysis of the transpose of the transition matrix can determine if a metastable limit cycle exists. For such systems, the second-largest eigenvalue, λ_2 of

the discrete approximation of the system provides an estimate of the length of continuous walking (in steps) expected on average, given initial conditions within a particular (well-defined) range of state space, while the eigenvector for λ_2 identifies the regions in state space most frequently visited during this metastable limit cycle.

To demonstrate the practical use of these tools, we investigate several particular walking systems. The two primary systems studied in this work intentionally span a broad range in walking. One is a computer simulation of a simple model of underactuated biped walking, called the compass gait (CG) walker. The other is a small, high-impedance 18-DOF quadruped with 12 actuated degrees of freedom. However, despite the differences in these two mechanical systems, both controlled walking systems do share two important commonalities: (1) In each case, legged locomotion inherently entails significant underactuation, and (2) the goal for each system is to walk on rough terrain.

1.1 Motivations for practical legged robots

Why have legs? The wheel is a remarkable invention. Vehicles with wheels can be highly efficient, and wheel-based locomotion can be designed to achieve an arbitrarily small turning radius² or to negotiate extreme terrain³.

Three broad classifications of explanation for mechanized, legged locomotion are as follows: (1) anthropomorphic fascination, (2) to understand human walking, and (3) to negotiate rough terrain and other environmental factors effectively and dynamically. This thesis focuses on the third motivation listed.

Below, we describe each of these three inspirations for having a robot with legs and conclude with a summary of why legs are so appropriate to rough terrain. This is followed by a selective review of the state of the art in legged locomotion, with an emphasis on what opportunities for advancement seem most relevant for research in the near future.

²Consider a wheelchair, or the Segway transport. The Segway inverted pendulum design provides a platform which is agile enough to play soccer on grass [19].

³Consider the rocker-bogie suspension strategy used on various Mars rover expeditions [59, 80].

1.1.1 Anthropomorphic fascination

Many researchers use a variety of phrases to describe our intrinsic interest in creating robots which attempt to mimic the particular range of motion and general, physical appearance of humans. My favorite expression for this motivation is “anthropomorphic fascination” [186]. For any number of practical or psychological reasons, we, as humans, seem driven to create robots in our own image. Robots created out of this apparent motivation seem concerned, primarily, with attaining motions which are kinematically (rather than dynamically) similar to human motions. A large number of humanoids have been developed recently, particularly in Japan [70] and throughout Asia. Robots such as Honda’s ASIMO [65, 150], AIST’s HRP-2 [85, 126], WABIAN-2R [134, 136] (Waseda University), Sony’s QRIO [45, 53], PINO [191] (Kitano Symbiotic Systems) and KIST’s MAHRU [95] are remarkable engineering accomplishments. However, these humanoids are generally designed with high-impedance actuators, intended to execute commanded joint trajectories with high fidelity but little compliance. As a result, these robots can “perform” well-planned motions, such as dance-like movements based on data recorded from humans using motion capture [139, 130, 95]. However, there are few (if any) examples of such humanoids walking robustly outside of highly controlled environments, such as a laboratory or stage.

Research interest has increased in recent years in developing control solutions to achieve dynamic gaits for such stiff robots [88]. This is an intrinsically challenging if not overwhelmingly task, since the physical designs of these high-impedance humanoid bipeds are intended to allow for motions approaching the range of **kinematic** degrees of freedom of a human, rather than for the range of **dynamic** impedance characteristic to humans. At some level, the concept of copying the nominal degrees of freedom and general structure of an actual animal which successfully locomotes (such as a human) seems a reasonable approach toward creating robots with similar capabilities. However, one is still confronted with the large task of deciding how to *control* all those degrees of freedom for stable walking.

The eventual goal of creating humanoids which are both kinematically *and* dynamically similar to humans seems much more appealing and practical, if we want to eventually develop robots which can operate out in the real world.

1.1.2 Understanding human walking

An additional reason to study walking mechanisms is to learn more about human walking. There are several applications for creating and validating models of human walking. Understanding human walking has direct application in areas which include: design of prosthetics (active or passive), evaluation of the effects of any of a number of gait pathologies may have on stability (for stroke patients, those with cerebral palsy, etc.) , and development of effective rehabilitation.

Another use for improved models of walking is in developing increasingly realistic animation of humanoid characters, for use in animated movies, video games and a variety of interactive software. Modeling human dynamics has helped researchers to develop animations which mimic the general patterns of humanoids while also obeying physics [1, 140, 106]. One goal of such technology is to automate the process of developing more natural-looking motions by using truly generative models, rather than simply replaying particular motions which are painstakingly captured from human models in motion capture.

A large variety of approaches have been taken in identifying and understanding the key elements of human locomotion. A few examples are listed below, and many more examples exist – particularly within the last ten years or so. [31] present a 3D model of the human gait, based on passive dynamic principles. [132] study the biomechanics of human running. [6] look into the dynamic optimization of a model of human walking. [104] examine models of kinematic control of a variety of animal locomotion, including human walking. [176] look into actively-adjustable leg compliance in creating smoother hip trajectories of a passive walking model, toward application to either rehabilitation equipment or robot design. [36] create a mechanical model to understand stumble recovery. [124] model human walking with a compass gait model and examine the effect of replacing a point foot with a curved one. [164] analyze a simple passive walking model on rough terrain toward estimating risk of fall in humans. Others have looked for the differences and similarities between human walking and that of other animals in bipedal gaits [64, 3], or at recovery strategies from perturbation while walking [77, 43].

In many such studies, ground reaction force data from real humans in walking or run-

ning are compared with the force profiles predicted by models. For example, many of the tools used in designing and analyzing legged robot gaits, such as the zero-moment point (ZMP) and the related foot rotation indicator (FRI), both described in Chapter 2, have also been used to analyze data on human walking [141]. Models allow us to test hypotheses about the magnitude and complexity of control action required to stabilize various aspects of dynamic, limit-cycle walking in a safe and systematic way. This understanding may in turn guide development of ways to help people walk with greater capability and stability: e.g., in guiding rehabilitation from injury; design of appropriate orthotic devices, regimens or appropriate surgical correction for chronic issues like cerebral palsy; and software tools to diagnose walking performance from a set of human gait data.

1.1.3 Legged-robot locomotion in real-world environments

An additional motivation is to create autonomous machines which are **dynamically** appropriate to interact with humans and their environment. This topic is becoming increasingly relevant as engineers develop robots for human environments. Some of the work to date in this area is surveyed very briefly below.

For robots to interact with humans and with many household items successfully⁴, they need to respond to a variety of often-unknown impedances they encounter in nature. This makes the requirement for robots in the real world quite different from those of stiff, factory robots. In a factory, precision is a primary concern, and overpowering the environment is generally an advantage (e.g., in machining, or in ensuring assembly is done accurately). By contrast, robots outside of a factory or laboratory will encounter far more variance when performing a given task. For example, in walking, the ground profile and ground properties (including friction, compliance, etc.) will often be both unknown and time-varying.

There are several possible solutions to regulate the force interaction between a robot and its world. One approach is to adjust the impedance actively *through control*. Several now-classic force-control solutions were introduced over two decades ago and include active stiffness control [151] and impedance control [73, 22]. Modeling the end effector dynamics, e.g., [93], is an important general concept which has continued to influence the

⁴i.e., picture a robot in a china shop...

design of robots and their control in human environments [92].

The performance characteristics of the actuators clearly have important consequences on the capabilities of robot systems as well [75]. Another direction in achieving desired force control is to design the actuators for legged locomotion with inherently low impedance. Actuator designs of this type include series-elastic actuators (SEAs) [143, 142] and the distributed macro-mini actuation approach (DM^2) [193]. The general concept of using variable stiffness in actuators specially designed for either manipulation or legged locomotion has become increasingly popular over time, with a variety of particular design approaches emerging in the last few years [89, 175, 172, 13].

1.1.4 So, why have legs?

Although there are a variety of motivations to study mechanical devices which walk, as discussed above, the focus throughout this work is on achieving robust locomotion for a robot. In particular, we have aimed toward developing legged locomotion which can provide some advantages over wheeled locomotion. Legged locomotion should ideally allow for both **(1) careful foot placement** (e.g., on intermittent terrain) and **(2) rapid traversal over a variety of mild obstacles**⁵. In Section 1.2 below, we examine the state of the art in legged locomotion today, with an emphasis upon how well each of these two particular goals is achieved.

1.2 State-of-the-art in legged locomotion

This section provides a brief introduction to the state of the art in legged robotics today, surveying a select set of exemplary solutions. Each style of robot (by which we mean its natural dynamics and accompanying control methodology) accomplishes a particular style of locomotion well. The reader should note ahead of time, however, that a critical goal for legged locomotion still remains unsolved: robots which can both (1) exploit their natural dynamics (e.g. make use of stored energy from the pendular walking or from spring

⁵By “mild”, we essentially mean not so extreme as to require a rock-climbing style, where force closure about the obstacle is required.

storage elements in the design, etc.) and also (2) negotiate obstacles (such as irregular steps, gaps or other terrain features) and other, unexpected environmental perturbations with high reliability. These two goals must both be satisfied to make legged machines practical in a real-world environment.

There are, of course, a large number of relevant examples of legged locomotion which have emerged in recent years [170, 33, 20, 138, 144, 189, 62, 182, 98, 71, 111, 37, 96, 103]. The combination of mechanical leg design and CPG⁶-based control in the robot “Tekken” [46] demonstrates a successful, dynamic approach to negotiate rough terrain with a quadruped with compliant legs. Of note in achieving flexibility to environment type is a quadruped salamander with a gait which can adapt to transition continuously from traveling on land to swimming [82]. Recently, much effort in legged locomotion has also focused on gecko-inspired climbing robots with feet which can produce adhesion forces [10, 153, 94]. Although a complete discuss of the State of the Art is outside the scope of this thesis, we will discuss four particular examples (in Sections 1.2.1 - 1.2.4) in some detail to survey the range of notable achievement to date in legged locomotion.

1.2.1 ASIMO

After a decade and a half of development in humanoid robotics, Honda unveiled the first version of the now-famous ASIMO humanoid robot in 2000. The design and capabilities of this robot have continued to advance over the last eight years. It is arguably the most advanced humanoid robot ever created, demonstrating both exceptional engineering in overall design and impressive capabilities in areas including locomotion, balance control, and vision processing. The robot is scaled to be the size of a child and is pictured in Figure 1-2.

In addition to ASIMO [150, 65], there are many examples of similar, highly-actuated humanoid robots in recent years [85, 126, 134, 136, 45, 53, 191, 95]. These robots all share particular characteristics. In particular, although they are all underactuated by virtue of the unidirectional (push-but-not-pull) contact between foot and ground, they are designed to operate as though they are fully actuated. They are designed with a large number of

⁶center pattern generator




Figure 1-2: ASIMO

Honda's impressive humanoid robot stands about 4'3" (130 cm) tall, weighs 119 pounds, has 34 actuated degrees of freedom, and can operate for up to 1 hour on internal batteries. Battery weight is about 13 pounds. Promotional diagram of ASIMO from Honda, circa 2005 (at left). Photo of ASIMO running (at right) is from <http://thefutureofthings.com/pod/121/the-rise-and-fall-of-asimo.html>

actuators, with degrees of freedom which roughly approximate the degrees of freedom of human motion. Although this does theoretically give these robots an incredible range of motion, it also makes the problem of *how* to control the motion quite complex. In general, the baseline aim is to maintain complete control authority over all degrees of freedom – these robots are not gymnastic, for instance.

To remain “fully” actuated, such humanoids classically depend on zero-moment point (ZMP) methods, described further in Section 2.1.3, which ensure that at least one foot remains essentially “anchored” by the weight of the robot throughout all motion. If the robot encounters unexpected terrain, obstacles or external interactions, however, it may lose balance and begin to topple over. Full actuation is a convenient assumption for these robots, but they still remain underactuated in reality.

To achieve high accuracy in motion trajectories, the design of these humanoids is quite intricate and their actuators have relatively high impedance. These high-precision robots are also quite expensive. The types of motions which are executed tend to be conservative, to prevent damage to the robot or to humans nearby due to falling or colliding with the environment, and the overall efficiency in locomotion or other motion is quite poor, compared

with humans [35]. Although ASIMO can “perform” intricate motions, including walking, turning, stair-climbing, stylized dancing (which emphasizes upper, rather than lower, body motions), and a highly-publicized “run-walk”⁷, these demonstrations must occur in a highly controlled and well-characterized environment; they are highly choreographed.

For all these reasons, the potential for these stiff, ZMP-based humanoids to negotiate rough terrain is seemingly limited. A large inspiration for the development of humanoids is to have them interact with us in our everyday environments, but this will clearly require significant advances – most likely both in control *and* inherent physical design.

1.2.2 Passive-dynamic based walkers

Passive dynamic walkers are devices which exhibit stable limit cycles to walk down a shallow slope with no external actuation or power. In these limit cycles, energy lost at each foot collision is balanced by the conversion of potential energy to kinetic energy in going downhill. Tad McGeer inspired a new direction in bipedal robot design with his original McGeer Walker [115]. Perhaps the most compelling example to date of a purely passive walker comes from the effort of Steve Collins, Andy Ruina, and Martijn Wisse [34]; this device is shown on the left half of Figure 1-3.

To enable locomotion on level ground, some source of actuation can be added. These “passivity-based” robots have sparked particular interest both because they use relatively low energy, compared with ZMP-based walkers like ASIMO [35], and because the dynamic coupling in the mechanical design of such walkers simplifies the control problem. Unfortunately, however, such limit-cycle walkers are notoriously sensitive to initial conditions and perturbations, severely limiting their practical applicability to date.

Much work has been done to analyze the stability of models of purely passive walking [56, 31, 101, 154, 78]. Producing active control to mimic but stabilize the passive dynamic motions of a walking model is a topic of much recent interest, both in modeling [135, 160, 137] and on real robots [190, 170, 33, 7, 148]. In this thesis, Section 4.3.2 introduces a “compass gait” biped model and explores its stability on stochastically-rough

⁷With a top speed of 6 km/hr and an airborne phase of .08 sec, this run, while technologically impressive, is certainly not practical for locomotion.




Figure 1-3: Passive-dynamic based walkers

The Cornell passive dynamic walker (left) and an actuated walker based on passive dynamics (right), both created by Steve Collins and Andy Ruina.

terrain, and Chapter 5 investigates minimal control strategies to allow a simple walker based on passive dynamic principles to negotiate rough terrain. Of particular note, our work in Chapter 5 addresses the foothold selection problem for passive-based walking robots – a problem which has not been adequately addressed in walking to date.

1.2.3 RHex




Figure 1-4: RHex hexapod negotiating rough terrain.

RHex, shown in Figure 1-4, is an impressive hexapedal robot capable of robust locomotion over a variety of rough terrain [5, 4]. This robot is the product of the collaborative

design effort of researchers from McGill University, Berkeley, U. Michigan and CMU. Much of its success depends on its carefully-designed compliant legs. Its six legs are each driven independently, allowing for a range of gait capabilities [63]. The legs rotate fully about the hip axis, so the robot can continue traveling even after it has been flipped over.

RHex clearly provides some advantages over a wheeled robot on many types of rugged terrain, traveling robustly with impressive speed⁸. [63] also discusses gait transitions for RHex, to allow it to transition between different classes of terrain and even to climb stairs. However, when RHex operates, it is essentially blind to the particular details of the upcoming terrain during operation, and so it is not capable of careful foot placement.

1.2.4 BigDog



Figure 1-5: The BigDog robot, by Boston Dynamics. Two BigDog robots, with a man in the background (for a sense of scale).

BigDog is a large, dynamic quadruped, produced by Boston Dynamics [21]. It is by far the most impressive robot which can negotiate rough terrain to date, although most of

⁸RHex can travel about one body length per second.

the details on its design and control remain unpublished. Some of the principles used in controlling BigDog have likely evolved from work in the 1980's and 90's by Marc Raibert and his students through research on a variety of dynamic robots at the Leg Lab at MIT [146, 147, 72]. The control approach for BigDog clearly relies on controlling forces and posture of the robot, which are central ideas discussed in detail both throughout legged robot research at the MIT Leg Lab and in [131]. BigDog is highly dynamic and has demonstrated recovery from significant perturbations, including being kicked or unexpectedly slipping on ice, using deliberate foot placement to control body attitude. However, BigDog is still largely blind to upcoming terrain and seemingly responds largely reflexively, to maintain dynamic balance.

1.2.5 Summary of state of the art

In Section 1.1.4, we stated two primary reasons to employ legs (rather than wheels) for locomotion: **(1) careful foot placement** and **(2) rapid traversal over a variety of mild obstacles**.

To compare the dynamic walking approach with that of humanoids such as ASIMO, we note that the passive-dynamic approach essentially mimics the *dynamics* of human walking, which is dominated by the inverted pendulum motion of the body over the stance foot. The resulting underactuated limit cycle gaits are quite fragile to perturbations from terrain variation or other interactions with the surroundings, however. By contrast, a highly-actuated humanoid robot like ASIMO employs a much fuller set of actuators and relies much less on natural, passive dynamics. Although adding more degrees of actuation seemingly increases the potential for greater robustness and flexibility in motion, it also increases the complexity of the *control* necessary, too. Humanoids such as ASIMO can achieve precise foot placement, but they are not yet capable on significantly rough terrain.

RHex and BigDog each demonstrate significantly more dynamic capability than either ASIMO or the dynamic walkers has demonstrated to date. However, neither RHex nor BigDog⁹ has solved the foothold selection problem for legged locomotion.

⁹nor the dynamic walkers...

1.3 Roadmap of thesis

This introductory chapter has attempted to motivate a practical desire for legged locomotion and has reviewed the state of the art in legged robotics. We hope the reader agrees at this point that there is a need in robotics for better solutions for walking and running, particularly to negotiate rough terrain at speeds comparable to human walking (or better). Although achieving this goal depends both on robot design (mechanical actuation, passive dynamics, sensing, etc.) and control, the remainder of this thesis focuses on control solutions.

Chapter 2 discusses the broader issues of how to go about designing and optimizing dynamic motions. This discussion begins by considering what aspect(s) of a locomotive motion on which we will focus, reviewing the most popular stability metrics for legged locomotion which have been proposed to date. Because legged locomotion involves such a high degree of underactuation and stochasticity, we argue that attempts at absolute guarantees of stability (i.e., never falling down) should instead be replaced by the *mean first-passage time* (MFPT) as an appropriate metric for stability. Our primary objective becomes that of obtaining highly-reliable motions and long-living locomotion gaits: simply said, our goal is exceptional performance from our robots most of the time while allowing occasional “failures” (i.e., falling down). This goal in turn motivates (1) the design of repeatable physical motions, particularly those which can exploit dynamics and underactuation for greater speed and agility, (2) a methodology for evaluating the long-livingness, or “metastability” of cyclical (locomotion) motions, ideally, through a quantifiable metric, and (3) an iterative process by which our control can be modified to obtain an optimal solution to maximize this metric. Each of these three elements, **design**, **evaluation**, and **optimization**, is then developed in more detail in each of the next three chapters.

Chapter 3 demonstrates the **design** of a range of motions for LittleDog, a small quadruped robot produced by Boston Dynamics. We describe several, related kinodynamic planning strategies, all of which are based upon considering external forces (ground reaction forces and gravity) as the essential means by which the body mass and inertia are pushed around in the world. Our technique borrows heavily upon those used in “ZMP” (zero-moment point)

control strategies; the ZMP approach has been used to develop successful walking motions for high degree-of-freedom bipeds such as ASIMO, which was introduced in Section 1-2. However, while the traditional goal of zero-moment walking is to remain in a *fully actuated* regime, our technique extends the approach to design underactuated motions, as well. The most dramatic of these motions is a highly-repeatable “dynamic lunge”, which has been successfully used to negotiate intermittent negative or vertical obstacles (essentially, gaps and vertical barriers) and step changes in terrain height. Our results yield fast, dynamic motions which are quite repeatable for execution of a step or two. As highly underactuated, open-loop motions are planned beyond a single step, however, they begin to demonstrate significant stochasticity. Our need to quantify and increase stochastic stability motivates the approach presented in Chapters 4 and 5.

In Chapter 4, we present a more careful, mathematical consideration of **evaluation** of the reliability of a walking system. We demonstrate a method for estimating the mean first-passage time, introduced in Chapter 2, by discretizing the dynamics of a system and performing an eigenanalysis. To illustrate this methodology, we analyze two simple walking systems on stochastically rough terrain: (1) the rimless wheel, and (2) a passive, compass-gait biped. The discretized dynamics of each system are well-represented as a Markov process. The main goals in this chapter are to apply statistical tools to analyze simple walking systems and to illustrate some of the properties typical of metastable systems more generally.

Chapter 5 presents a hierarchical control approach to obtain approximate **optimization** of an actively-controlled (but still underactuated) version of the compass gait walker on rough terrain. We present results both on stochastically varying and on known wrapping terrains. The use of wrapping terrain allows us to obtain a near-infinite horizon solution for optimal control. Using the policy obtained for stochastic terrain to “seed” the value function for the wrapping terrain, we find that the policy obtained with a very short look-ahead, i.e., one or two steps ahead on terrain, is capable of continuous walking. Our results strongly imply that near-optimal walking solutions on rough terrain may often be found using a relatively short look-ahead.

In Chapter 6, the bounding style dynamic lunge presented in Chapter 3 is augmented

with the addition of a secondary rocking motion to produce a reciprocating bounding gait. This is a brief chapter, and the primary goal is to discuss the concepts of evaluation and optimization presented in Chapters 4 and 5, respectively, for control of a real-world robot. We hope this presentation inspires the reader to apply the ideas presented throughout toward developing metastable locomotion on highly dynamic mechanical devices. An essential element here is the reduction of a high degree-of-freedom robot to a system with much lower complexity, through low-dimensional physical models which capture the dominant dynamics.

Finally, Chapter 7 provides a summary and suggestions directions for future work. Of particular note, we highlight the fact that the techniques developed throughout are highly applicable to machine learning.

Chapter 2

Metrics and Methods Toward Highly Dynamic Locomotion

Engineering often balances reward against risk. For legged robots, we want both high performance (i.e., agility and speed on a variety of terrain) and high reliability (i.e., low risk of falling).

While the previous chapter argues that better control solutions for legged locomotion are needed, this chapter motivates the overall *approach* we develop throughout the remainder of the thesis toward achieving this. This chapter has two parts. **First**, we discuss what measure(s) should be “optimized” in design of a walking system. Here, we emphasize a particularly straight-forward goal: minimize the rate of failure (e.g., falling-down) for a given locomotion task. **Second**, we discuss possible methods of achieving this optimization. A particular contribution throughout this thesis is the development of methods which can produce desired motions which are both *highly dynamic*, as motivated in Chapter 1, and *highly reliable*, as motivated below. In the end, our goal here is essentially to minimize risk for a pre-determined performance requirement.

We begin below by reviewing some of the stability margins and stability metrics commonly used in evaluating the performance of robots with legs. We present arguments about what is both good and bad about each method of measuring performance, providing some particular toy examples as illustration. The review is intended both to familiarize the reader with the analysis tools roboticists typically use today and also to motivate the need for tools

presented in detail in Chapter 4 for more directly estimating the stochastic stability of walking systems.

In the second half of this chapter, we propose the well-known concept of metastability to capture the long-living (but not globally stable) nature of walking, and we outline a straightforward approach for achieving this. Chapters 3 through 6 provide particular examples of legged systems, including both computer-simulated models and a real quadruped robot.

2.1 On metrics for legged locomotion

Humans are, in general, exceptionally good at walking. Normal human walking is metabolically efficient and can be used to negotiate significant sources of terrain variability, such as ground compliance and height and terrain slope variations, while requiring little conscious control effort. A general hope has existed for some time now that we can design legged machines with appropriate intrinsic dynamics and low-level control [178, 147, 173, 122] to mimic such performance. Many of the existing solutions for legged robots, such as those described in Chapter 1, are sufficiently dissimilar in approach, however, that it is currently difficult to make quantitative comparisons between them.

Appropriately chosen metrics are important both in evaluating performance of our robots and in guiding their design. For example, a natural task definition in robot locomotion involves traversing a particular path on terrain, perhaps within a specified time limit. Given the constraints of the problem specification, we wish to design a control policy to maximize the likelihood of success. At other times, we may be free to plan speed as required and would like to know how speed affects our risk. In this second case, we wish not only to optimize likelihood of success but also to *quantify* it, so that we can compare risk and reward over a variety of possible speeds. The problem of estimating reliability is of course important in a wide range of engineering applications, including microelectronics [84], building survival during earthquakes [12, 167] and manipulation planning [112], but it has not been sufficiently emphasized within the walking community as yet. Emphasizing a viewpoint toward quantification of reliability is a significant contribution of this

thesis, overall.

2.1.1 Literature review on metrics for walking

Below is brief review of some metrics which have been proposed and used in evaluating walking robots; for interested readers, [67], [49] and [145] provide lengthier reviews. The velocity-based metrics and definitions for “stability” in walking proposed in [145] are of particular interest in this section.

2.1.2 Static stability margin (SSM)

The most elementary stability margin discussed here is the static stability margin (SSM). The static stability margin is the minimal distance which the projection of the center of mass onto the ground must travel before it exits the convex hull of the leg-to-ground contacts. This convex hull is known as the support polygon, because it defines the shape within which a stationary center of mass (COM) must be to prevent toppling.

In cases where the mass is *not* stationary, the force of gravity may no longer be the only force acting to produce a toppling moment on a legged entity. In the more general case of a robot with non-trivial accelerations over time, it is the location of the **center of pressure** (COP) with respect to the edge of the support polygon which indicates an instantaneous danger of initiating toppling. When the COP remains strictly within the support polygon, so there is no net toppling moment about any edges of the support polygon, then we refer to the COP as the “zero-moment point” (ZMP), as described in the next section. If a walker is moving slowly enough, however, then the static stability margin can often serve as a good approximation for the ZMP.

Thus, the static stability margin provides a geometric “safety margin” which we can guarantee is allowed in static displacement of the mass (in any arbitrary direction) without its projection exiting the support polygon: it is essentially a definition of the “wiggle room” we have before worrying that our stability definition may have been violated. The gain and phase margins of classical control theory are similarly defined to quantify the magnitude of allowable uncertainty we can accept while remaining stable. For stable systems, such

a margin has a positive value to indicate a positive amount of wiggle room before losing our stability guarantee. In cases where a system is *already* in a configuration defined as *unstable*, a stability margin is negative, and its magnitude indicates the minimal distance required before re-entering a stable regime.

The static stability margin is very appealing because it is simple to calculate, but its simplicity comes at the price of ignoring the dynamics of the system. As legged motions become more dynamic (gazelle vs. turtle, say), the static stability margin loses its utility, and we need better indicators of stability during locomotion. For example, a robot can have a net moment about an edge of its support polygon and still demonstrate static stability. Similarly, a robot can be in complete violation of the static stability margin while maintaining its center of pressure strictly inside the support polygon. The static stability margin is not an accurate predictor of robot “stability” because (1) it is neither necessary nor sufficient to avoid toppling moments as robot motions become more dynamic and (2) stable limit cycles may exist even if we *do* have intermittent “toppling” moments during walking, as is the case for the passive walkers described in Section 1.2.2. The zero-moment point (ZMP) margin, discussed next, addressed the first (but not the second) of these serious limitations of the SSM.

2.1.3 Zero-moment point (ZMP) margin

The zero-moment point (ZMP) for a legged robot is the unique point on the ground about which the total torque produced by robot-ground interactions is precisely zero. In other words, it is the *center of pressure* (COP) exerted by the robot on the ground. The concept of the zero-moment point was given its familiar name in the late 1960’s by Miomir Vukobratovic [179]. In its original context, the ZMP was intended as a way of insuring that desired motions of the internal degrees of freedom (center of mass position and accelerations, plus any rotations of inertias) could be supported by the inherently underactuated¹ interface between the “feet” of a robot and the ground. The term “ZMP” is most typically used when designing for *planned* center of pressure locations, as a starting point toward developing COM and joint angle trajectories. Vukobratovic originally termed this motion

¹i.e., the vertical force achievable at the foot is unilateral.

design process the *synthesis of artificial synergy*; in his words:

The basic concept used in the synthesis of artificial synergy is that the law of the change in the reaction and friction forces is given in advance, or that their interrelationship has been prescribed. For example, the ZMP motion is given and the point through which the resulting friction force passes is prescribed. This determination of dynamic quantities imposes dynamic relations which will result in additional dynamic constraints into the model of the system [178].

By contrast, the term COP is more typically used when discussing actual measurements, as when measuring forces at the feet to maintain stability in a feedback loop. Somewhat implicit in the use of term “ZMP” is a concept of maintaining an intentionally *actuated* degree of freedom at the ankle, to avoid toppling moments over on the edge of the foot.

The ground contacts during legged locomotion are inherently unilateral and underactuated [55]. As a robot begins to topple over, the center of pressure it exerts on the ground goes to the edge of its support polygon. Because the foot can only push and not pull on the ground, it becomes more challenging, if not impossible, to stabilize the motion of the robot as it tips. Dynamically, the robot becomes underactuated, because it can no longer use ankle torque at the ground contact in the tipped state.

Direct control of the ZMP has an intuitive appeal, such that we maintain the center of pressure far enough from the edge of the foot and can respond to any destabilizing perturbations or joint trajectory errors by pushing against the ground to apply appropriate ground reaction forces. Maintaining a desired ZMP trajectory over time requires a careful strategy, however. It is possible to balance the ZMP of a robot exactly at the center of its support polygon, for instance, while crashing the body to the ground. A trivial example of this is presented below to illustrate the inherent conflict between regulating the instantaneous location of the ZMP and ensuring its long-term stability.

Figure 2-1 shows a planar example of a simple mechanism with a large base (or “foot”) and a single source of torque actuation, to control the motion of a pendulum. To treat our simple “robot” as a fully-actuated device, we must guarantee that the foot maintains both point contacts with the ground. Once it begins to roll onto a single point of contact, the




Figure 2-1: Inverted pendulum (2D) balancing on a weighted “foot” platform

At a particular instant in time, we may solve for the required torque which results in a ZMP at the center of its support polygon or “foot”. However, this direct strategy will move the center of mass further and further from its equilibrium position, resulting in failure. Successful balancing requires a better, long-sighted strategy.

system dynamics become nearly identical to those of the (underactuated) acrobot²; control of the full state of the robot becomes less straight-forward; and the robot may topple over.

To maintain positive pressure at both F_1 and F_2 , let us write a force balance at each point contact, valid only when $F_1, F_2 \geq 0$. Assuming the height of the point-contact “feet” on the bottom of the base is negligible in our model, we obtain:

$$F_z = \frac{-\tau}{a} \sin \theta + mg \cos^2 \theta - ma\dot{\theta}^2 \cos \theta \quad (2.1)$$

$$F_1 = \frac{-\tau + b_2 F_z + (b_2 - a_f) m_f g}{b_1 + b_2} \quad (2.2)$$

$$F_2 = \frac{\tau + b_1 F_z + (b_1 + a_f) m_f g}{b_1 + b_2} \quad (2.3)$$

where the term F_z represents the external force acting on the pendulum mass, m , which acts at the pivot point and g is $9.8 (m/s^2)^3$. A full derivation of these equations appears in

²...except for friction limits in the horizontal forces and the fact that ground reaction forces are only unilateral in the vertical direction.

³i.e., it is the *magnitude* of gravity, $9.8 m/s^2$, not the z-directional value of -9.8.

Appendix A. When both F_1 and F_2 are strictly positive in magnitude, we can calculate the location of the ZMP (equivalently, COP) with respect to the toe at F_1 as:

$$x_{zmp} = \frac{F_2(b_1 + b_2)}{F_1 + F_2} \quad (2.4)$$

To attain the maximum ZMP margin, with the ZMP exactly midway between the two toes, we would simply have to select a value for τ which would instantaneously result in identical forces, F_1 and F_2 at the toes, i.e., $F_1 = F_2$. Equating the right-hand sides of 2.2 and 2.3 and then using 2.1 to substitute for F_z , we can derive an instantaneous solution for τ trivially as:

$$\tau + b_1 F_z + (b_1 + a_f) m_f g = -\tau + b_2 F_z + (b_2 - a_f) m_f g \quad (2.5)$$

$$\tau \left(2 + (b_2 - b_1) \frac{\sin \theta}{a} \right) = (b_2 - b_1) \left[mg \cos^2 \theta - ma\dot{\theta}^2 \cos \theta \right] + (b_1 - b_2 + 2a_f) m_f g \quad (2.6)$$

$$\tau = \frac{(b_2 - b_1) \left[mg \cos^2 \theta - ma\dot{\theta}^2 \cos \theta \right] + (b_1 - b_2 + 2a_f) m_f g}{\left(2 + (b_2 - b_1) \frac{\sin \theta}{a} \right)} \quad (2.7)$$

Here, given any small perturbation from equilibrium, the effect of this torque is negligible⁴ on the total motion of the pendulum mass, m , and may even accelerate how rapidly it arcs dramatically toward the ground. Maintaining the ZMP at the center of the support polygon is entirely compatible here with “falling down”. Only long-term regulation of ZMP makes any sense, since the ZMP location depends on accelerations, which are set instantaneously for this idealized second-order dynamic system.

The dilemma between achieving instantaneous regulation of the ZMP and long-range stability is perhaps more clear for the toy example of a planar legged robot shown in Figure 2-2. We can eliminate the intermediary step of calculating the x and z force components at each foot in deriving an equation to relate the overall motion of the (point) mass to the location of the ZMP on the ground. By definition, the ZMP is the location where the sum

⁴We will justify this comment more precisely in the pages ahead...




Figure 2-2: Minimalist fully-actuated legged robot

All internal joints are fully actuated. The center of mass motion can be treated as fully actuated as long as kinematic and ZMP constraints are not violated.

of all moments is zero, i.e.:

$$0 = z_m \ddot{x}_m - (x_m - x_{zmp})(\ddot{z}_m + g) \quad (2.8)$$

$$x_{zmp} = x_m - \frac{z_m}{\ddot{z}_m + g} \ddot{x}_m \quad (2.9)$$

To remain in the regime where we can assume full actuation, we require a net positive vertical force, F_z , at each foot contact, as described for the previous example from Figure 2-1. Correspondingly, this requires that the sum of the vertical acceleration of the mass and gravity remain positive: $\ddot{z}_m + g > 0$. We shall also assume remaining “not fallen” requires a positive height: $z_m > 0$. Given these physical restrictions, it is clear that the sign of the acceleration, \ddot{x}_m , required to bring the ZMP to some particular, desired ZMP location, is always opposite the sign of the displacement:

$$\Delta x = x_{zmp} - x_m = -K_{pos} \ddot{x}_m \quad (2.10)$$

where $K_{pos} = \frac{z_m}{\ddot{z}_m + g} > 0$. We know this intuitively, of course. If you stand in equilibrium and wish to accelerate to the right, you must shift your center of pressure to push with

your left foot. Immediate regulation of the ZMP would only result in re-centering the mass between the two feet in the improbable event the initial *velocity* just happened to be *perfectly* matched such that the acceleration slows the mass down to a dead stop at the desired position.

Consider the particular case where we use inverse kinematics to constrain the mass to travel a constant height, $z_m = L$. Now, the relationship from the desired ZMP, x_{zmp} , to position of the mass, x_m , shown in Equation 2.9 becomes identical to the standard, linearized inverted pendulum dynamics. Using the Laplace operator ($s \equiv \partial/\partial t$, $s^2 \equiv \partial^2/\partial t^2$, etc.), this transfer function is:

$$\frac{X_m(s)}{X_{zmp}(s)} = \frac{1}{s^2 - \frac{L}{g}} \quad (2.11)$$

The two real-valued poles for this height-constrained system are $s = -\sqrt{L/g}$ (stable) and $s = +\sqrt{L/g}$ (unstable). One should also note that the solution for the evolution of the dynamics is equivalent to the dynamics of the pendulum-foot model (Figure 2-1) if it is linearized about equilibrium. This should explain to the reader why we earlier claimed that the effect of the torque, τ , as defined in Equation 2.7 is “negligible”, as mentioned in passing on Page 41. Near equilibrium, the natural dynamics of the inverted pendulum already maintain the ZMP at the center of the support base already, so the required torque command is approximately⁵ zero.

We can stabilize such dynamics using any of a number of classical approaches. In general both within this thesis and in the broader ZMP-based robotics literature, we will assume the *future desired ZMP trajectory is known* for some finite epoch ahead in time. We can take advantage of this to essentially “violate the causality” implicit in Equation 2.11, so that the mass moves in a way which takes advantage of the upcoming desired trajectory. The essential idea behind most of the published results for ZMP control is simply to plan and to follow a particular trajectory of ZMP to achieve some desired overall motion of the body of the robot. Some examples of legged locomotion using this general methodology are discussed below, briefly.

⁵The location of the actual pendulum pivot is not quite at the center of the foot in Fig. 2-1

ZMP-based approaches to control walking

A variety of robot gaits have been designed using a ZMP criterion. ASIMO, described in section 1.2.1 may be the most famous example to date [65], but ZMP methods are also used extensively throughout the rest of the humanoid robotics community. Of particular note, the work of Shuuji Kajita in employing Preview Control to follow a desired ZMP trajectory has produced successful dynamic walking gaits bipeds [86] and even some initial results for running bipeds [88]. Preview Control [155, 174, 91, 168] will be discussed in more detail throughout Chapter 3.

ZMP-planning techniques have also been used in designing more dynamic motions, including an intermittent, omnidirectional trot for the Titan VI quadruped [192] by planning motions in which the ZMP lies exactly on the support line of diagonally-opposed feet. We successfully demonstrate a similar strategy for designing double-support gaits for the LittleDog quadruped in Section 3.7.

Others have used general ZMP-inspired concepts toward analysis of the dynamics of passive walkers [123], i.e., by following the point of contact between an arc foot and the ground. Various researchers have also used or been directly inspired by ZMP principles in controlling many other robotic systems and models [152, 165, 126, 9, 68, 113, 30, 69, 61, 2], including ZMP-based running bipeds [87] or quadrupeds [126, 192, 103].

Planning and predicting the ground reaction forces which will act on a robot is clearly a useful technique in understanding what motions of the mass and inertia of a legged robot are attainable; this is also a key element in designing our own dynamic motions for the LittleDog robot in Chapter 3. However, the practicality of the zero-moment point as a *stability margin* is inherently limited, as it does not provide direct information about how easily a robot can be toppled by disturbances. That is, it does not truly quantify stability.

Often, an increased ZMP margin does not even correlate directly with (let alone meaningfully quantify) increased stability. Passive walkers can have point feet with no ZMP margin yet exhibit stable limit cycles; ZMP location merely correlates to a likelihood of “tipping”, which may be an intentional feature in well-designed, efficient walking strategies. Second, even if we chose to limit our desired behaviors to remain “fully actuated”⁶,

⁶i.e., in the sense of maintaining a non-zero ZMP margin, to avoid the underactuated toppling dynamic of

the geometric ZMP margin is an incomplete quantification of danger. For example, if the simple pendulum with a base from Figure 2-1 is highly asymmetric, the correspondence between the ZMP location and the propensity to tip may be highly asymmetric as well. The physical location of the ZMP does not provide information about the *magnitude* of the pressure force acted between the robot and ground, nor the height of the mass; thus, it does not cleanly quantify the magnitude of the disturbance required to instigate toppling. To better capture a notion of the size of perturbation which will cause tipping, a number of other metrics have been proposed and are outlined briefly in the following section.

2.1.4 Velocity- and energy-based metrics

As legged gaits become increasingly dynamic, the COM and ZMP stability margins are clearly insufficient metrics of stability. Stable limit cycle gaits exist for walkers which do not have a meaningful ZMP margin, for example, because their support polygon provides zero “margin” for control⁷ and/or because contact is only intermittent⁸. For robots that exploit underactuated motions, it is necessary to discuss stability in terms of better metrics. Several new metrics have been suggested within the last decade or so; some notable examples are reviewed briefly below.

Goswami introduced the Foot Rotator Index (FRI) [55] and also later suggested a new criterion for legged robots based on angular momentum [57]. As mentioned, ZMP-based legged motion planning generally involves selecting a desired trajectory of the ZMP over some future time horizon. Generally, the goal is to avoid tipping over the edge of the support polygon, so one can plan motions using a fully-actuated model. Once a robot begins to tip, one may calculate the theoretical location in which one would need to place the ZMP in order to prevent tipping, and this is the so-called FRI. The actual ZMP cannot exit the support polygon, of course, but the FRI indicates how far one would need to extend the support polygon in order to recapture a moment-controlling center of pressure to prevent tipping. The ZMP, FRI and an additional ground reference point called the Centroidal

many natural walking gaits.

⁷i.e., in the case of either a point contact or a line of support.

⁸e.g, for running and bounding gaits.

Moment Point (CMP) are discussed and compared in more detail in [141].

Hirose provides a good survey of stability margins and suggests his own metric: the Normalized Energy (NE) Stability Margin [67]. This margin is related to the Energy Stability Margin [119], which is the change in potential energy ($mg\Delta h$) required for a robot to topple over the edge of a support polygon. Hirose normalizes this value by the weight of the robot (mg), so that a heavier machine is not considered more stable. The goal is to create a stability margin which captures a more general measure of stability in terms of geometry and mass distribution; the NE Stability Margin is simply the height difference which will occur during the topple (Δh). Yoneda and Hirose had previously suggested the Tumble Stability Margin [192]. The concept here is to calculate the net moment required to topple a robot about each particular line of the support polygon and to normalize these values (as with NE Stability Margin) by the weight; this margin is the shortest resulting distance (after all lines of the support polygon are analyzed) and coincides with the SSM in the limit on level ground when there is no acceleration of the mass. Related stability metrics reviewed in [67] also look at the angle of inclination required for tipping.

Garcia et al. provide another survey of the various types of stability margins in use for legged robots [49], and they suggest a modified version of NE Stability Margin, which they call the Normalized *Dynamic* Energy Stability Margin [48]. Here, the goal is provide a metric to indicate as directly as possible the largest impulsive energy change the robot can safely take, normalized again by its weight, mg . The key aspect here (as indicated by the name) is that full dynamics (most notably, inertia and any torques or forces from all interactions with the environment, which may include manipulation with arms) are considered.

Although not strictly a “margin” of stability, Hirukawa introduces the universal stability criterion [69], which looks at the full contact wrench between the robot and its environment. This criterion allows one to evaluate a ZMP-type stability which includes arbitrary roughness of terrain and 3D contact points, and it checks if the friction cone has been violated. Friction considerations are often neglected in stability margin analyses but are worth mentioning briefly, as they were an important consideration in planning acceleration of the mass for dynamic lunging in Chapter 3.

In general, all the metrics listed so far in this section aim at ensuring that a robot does

not topple over. However, we reiterate once more, passive-dynamic principles allow for stable limit cycle walking which exploits (rather than avoids) such toppling dynamics. Wisse et al. boldly state that in walking, stability has just two requirements: “You will never fall forward if you put your swing leg fast enough in front of your stance leg. In order to prevent falling backward the next step, the swing leg shouldnt be too far in front.” [188] This concept is appealing and sensible, but it does not provide a quantifiable metric for stability. Passive-based walkers remain notoriously sensitive to perturbations and initial conditions, both in theory and practice.

Toward quantifying the stability of a variety of legged gaits more generally, Pratt and Tedrake introduce velocity-based metrics, with the aim of predicting the necessary stopping time for a robot to come to rest safely [145]. This is an intuitively-pleasing metric, since it gives a direct estimate of the requirements needed to safely arrest forward motion without falling. They provide several particular values of interest, along with both brute-force and approximate estimation methodologies for obtaining them.

As Full et al. note, defining “stability” for a legged system is a difficult task [47]. For a cyclical motion, a better definition should, they argue, look at the ability to return to a nominal gait or cycle after a perturbation. Wieber presents a related idea about stability essentially meaning *convergence* and proposes a Lyapunov Stability Margin [187].

We believe the ideal stability metric for walking should incorporate all factors which may ultimately lead a robot to fall down. It should provide an overall estimate of the instantaneous likelihood of falling and/or the expected length continuous walking to quantify stability. The particular metric we propose here is simply the expected average time of travel before falling, called the mean first-passage time (MFPT).

2.2 The mean first-passage time (MFPT) metric

The mean first-passage time (MFPT) is a term borrowed from the physics community. In engineering it is also often called the “mean time to failure” (MTTF). We believe this metric provides the most straight-forward and intuitive measure of stability for walking robots. Its greatest liability is the potential difficulty one might have in estimating its value. Chap-

ter 4 addresses this directly, however, demonstrating tools to analyze long term stability of walking systems. Below, we discuss the appropriateness of the MFPT metric briefly.

In general, legged gaits involve motions which orbit near some particular limit cycle in state space. For idealized systems, one can prove *local stability* through a Poincaré analysis, guaranteeing that small perturbations of the limit cycle in any direction in state space contract back to a fixed point [58]. One can also calculate the region(s) in state space from which the limit cycle dynamics are stable and will converge to the fixed point [154]. However, global stability is much more difficult to demonstrate, even for simplified models of walking. Global stability has been demonstrated for the “rimless wheel” model on a slope⁹ [171], but only if either the standing-still fixed point is considered meaningful (for a walking system) or if the slope of terrain is so steep that the wheel can never stand still. Also, even in the trivial case of the ideal rimless wheel model (with instantaneous, inelastic collisions), the inherent, one-way underactuation of the ground-foot contact may cause the model to exit its assumed dynamic regime entirely; both feet will lose contact with the ground, if it is rolling too fast (such that actual acceleration of the point mass in the idealized model would infeasibly exceed gravity). Legged locomotion is almost always susceptible to this underactuated contact between the foot and the ground – specifically, the foot can push at the ground but cannot pull¹⁰. Also, legged locomotion usually involves significant step-to-step variability¹¹. The terrain and perturbations to which a walker will be subjected are never known perfectly, and their effects on the underlying dynamic stability of a walking system are difficult to model analytically.

As a result, we believe well-designed, practical, legged robots are properly classified as systems demonstrating *stochastic* stability, as discussed in Chapter 4. Our goal should be to attain excellent performance most of the time with only occasional failure. We should not require *global stability*: guarantees of recovery from any possible perturbation are not realistic in realworld walking. Moreover, it is not sufficient to consider the effect of any single perturbation on the nominal limit cycle in quantifying what our controlled system can

⁹The rimless wheel is described in Section 4.3.1.

¹⁰*Almost* always. Geckos provide a nice counterexample [94].

¹¹Again, I’m sure the clever reader can come up with important exceptions, such as physical training on a treadmill, but *almost* always...

handle, because every step in locomotion adds new perturbations. It is critical to understand the *mixing* effects of all these perturbations to quantify the stochastic stability accurately.

2.3 Metastability: Beyond stability

There is generally a conflict between the goals of exceptional dynamic behavior and of guaranteeing stability of a walking system. As mentioned, walking systems are always underactuated, and it is impractical and undesirable to limit their operation to such carefully prescribed environment that absolute guarantees of stability can be made¹². Realistically, what we should expect is exceptional behavior *most* of the time, with an emphasis on *minimizing* the rate of failure.

Many systems exhibit long-living behaviors which are not technically “stable”. This is a well-defined concept, which is referred to as “metastability” [51]: literally, “beyond stability”. These stochastic dynamic systems exhibit behaviors which are impressively long-living, but which are also guaranteed to exit (“fail”) with probability one given enough time. Such systems cannot be classified as “stable”, but it is also misleading and incomplete to classify them as “unstable”. Physicists have long used the term *metastable* to capture this interesting phenomenon and have developed a number of tools for quantifying this behavior [60, 90, 127, 169]. Many other branches of science and engineering have also borrowed the terminology to describe dynamic systems in a wide variety of fields. Familiar metastable systems include crystalline structures (e.g., diamonds), flip-flop circuits, radioactive elements, oscillatory wave patterns in the brain, and ferromagnetic materials, such as spin glass or magnetic tape film (which explains why a taped recording sitting in storage inevitably fades over time).

The canonical example of metastability is a particle in a potential well subject to Brownian motion, as cartooned in Figure 2-3. Such systems have local attractors which tend to keep the dynamics within a particular neighborhood in state space. In the limit as these systems become deterministic (no noise), the local attractors are fixed points, and the sys-

¹²That is, one cannot give a global guarantee that *no* disturbance, no matter how unlikely, could possibly make them fall.




Figure 2-3: Cartoon of a particle subject to Brownian motion in a potential $U(x)$ with two metastable states, A and B .

tem is truly stable whenever the dynamics begin with an initial condition somewhere inside the basin of attraction of the fixed point. In contrast, stochasticity constantly pushes the dynamics about within this neighborhood, and for some systems and noise types, this turns a stable system into a metastable one. Over time, such systems deviate particularly far from a metastable attractor in state space (making “escape attempts”) and eventually they will exit (by which we mean entering a region where a different local attractor dominates).

2.3.1 Walking as a metastable limit cycle

Many researchers discuss the existence of strictly stable limit cycles for ideal passive dynamic or mildly-active passivity-based walkers [114, 58, 50, 56, 154, 83, 8]. As we have previously mentioned¹³, passive-based robots are notoriously sensitive to both initial conditions and perturbations. With additional complexity, such as passive knees, parameters (lengths and relative masses) of a passive walker must be chosen with extreme care for a stable limit cycle to exist at all [78].

In order to claim that a system is metastable, we need to show that it will converge to its long-living state (or limit cycle) very quickly, with respect to the time scale on which failures occur. That is, initial conditions should be forgotten “quickly”, with respect the time scale on which a particular sequence of successful steps occurs. Additionally, we need to define what an acceptable definition of “long-living” is, in the sense of a particular

¹³in Section 1.2.2 (p. 28) and elsewhere.

number of steps taken on average before falling. We do not attempt to define a strict limit on the required MFPT for metastable walking, but intuitively, we can probably agree it should be on the order of thousands of steps (at least), rather than dozens.

2.4 Discussion and conclusions

Below, we call to the reader’s attention to some additional considerations about metrics for walking and about the practicalities of estimating of time-to-failure on real-world system. We conclude this chapter by restating our belief that thinking about walking as an inherently stochastic process is appropriate and that the statistics of failure are correspondingly the best-correlated metrics for measuring stability.

2.4.1 Efficiency of walking

Another import metric for legged locomotion performance, beyond considering stability alone, is efficiency [99]. Humans are very efficient at walking compared to most bipedal robots, although walkers designed using passive dynamic principles have demonstrated comparable efficiency [33]. To date, there is certainly a trade-off between efficiency and stability, however. For example, as discussed in Section 1.2, there are two distinct solutions for bipedal locomotion which exhibit this trade-off: the dynamic walking approach in Section 1.2.2 is efficient but also highly sensitive to perturbations, while highly-actuated humanoids such as ASIMO (Section 1.2.1) maintain much better control authority. Additional control authority adds capability to remain more stable, but it also makes operation much less efficient. The author has personally “pushed around” both an ASIMO¹⁴ and several passive-based walkers. A slight tap can destabilize a dynamic walker, but ASIMO can withstand significant horizontal pressure and respond to stand stably (distributing weight evenly on its feet, to keep its ZMP well within the support polygon). Although this thesis focuses specifically on the issue of legged locomotion stability, incorporating both stability and efficiency into a practical robot is certainly the ultimate goal in legged robotics, and a

¹⁴Honda Research in Cambridge, MA has one.

cost function based on both of these criteria could clearly be used in optimizing a control policy for a robot.

2.4.2 The right answer?

For different regimes of locomotion, the Static Stability Margin or ZMP may be entirely appropriate and acceptable as a stability metric, while other motions will require limit cycle analyses. Universally, however, we always want robots which perform exceptionally most of the time with only occasional failures (falling), and a metric which can quantify this level of reliability therefore seems natural for encompassing the entire class of entities (machines and animals) undergoing legged locomotion. The strongest disadvantage of the MFPT metric is in knowing the best way to estimate this value. We are confident that this remains the “right way” to think about legged locomotion, however, and we are optimistic that the mathematical tools presenting ahead in this thesis for estimation of the mean first-passage time both on walking models and on real robots can and will prove their utility (and undergo significant refinement) over time.

Chapter 3

Kinodynamic Planning for LittleDog

In this chapter, we describe a design methodology for planning kinematically and dynamically feasible locomotion for a small, quadrupedal robot called “LittleDog”, pictured in Figure 3-1¹. The focus of this chapter is on *design of control* solutions for locomotion, while Chapter 6 discusses issues relating to their *evaluation and improvement*. The most significant challenge in our work has been the extension of traditional kinodynamic planning to achieve reliable double-support motions which are highly dynamic and underactuated. As in human and other animal locomotion, control of a robot in underactuated motion can produce significant gains in speed and in extremity of traversable rough terrain.

3.1 Introduction: Control solutions for locomotion

The task of locomotion is essentially an exceptional case of manipulation. Rather than controlling the degrees of freedom of particular, external objects in the environment, however, successfully locomoting oneself requires the generation of appropriate forces to control one’s *own* degrees of freedom over time. Locomotion requires movement of the entire robot within the environment, which makes it uniquely challenging as a “manipulation” task. In traditional manipulation tasks, one can often make the convenient and reasonable simplification that the robot is either bolted down or is massive enough to assume it has a fixed

¹LittleDog is a product of Boston Dynamics Inc. (BDI) which also developed the BigDog robot, described in Chapter 1.



Figure 3-1: The LittleDog robot
Alec Shkolnik (left) and Katie Byl are also pictured in the background (for scale).

reference frame with respect to the environment, and one can then employ high-impedance actuators to achieve an essentially fully-actuated robot. By contrast, a legged robot can not achieve a fully-actuated reference frame with respect to its environment: feet push but cannot pull, contact is intermittent, and lateral forces are limited by friction constraints. Both manipulation and locomotion require careful planning of kinematics and achievable dynamics; locomotion, however, occurs within a significantly more underactuated regime.

As robots are developed with more degrees of freedom, it becomes non-trivial to plan successful locomotion which maintains both kinematic and dynamic feasibility; resolving these two requirements simultaneously is known as *kinodynamic motion planning*. Kinodynamic motion planning for high-degree-of-freedom, humanoid robots has become a popular topic of research, particularly in recent years [107, 110]. A now-standard approach currently employed for solving this problem involves two parts. First, a series of kinematically-feasible poses are found, each of which is also statically stable, as defined in Section 2.1.2. Then, a “speed knob” is essentially “turned up”, increasing the speed of execution of these motions while insuring the zero-moment point (ZMP) (aka, center of pressure) remains within some safety margin of all edges of the support polygon, as described in Section 2.1.3. In other words, kinodynamic planning is traditionally employed to ensure that a robot operates in a fully-actuated regime, only.

Our primary aim in working with LittleDog has been to create motions which are both highly agile and highly reliable. Agility is clearly desirable in negotiating extreme terrain². High reliability is also an obvious goal, since falling down slows you down and also carries risk of physical self-damage. We obtain motions which are both repeatable and agile, such as the dynamic lunge depicted in Figure 3-2, by reasoning about ground reaction forces and the dynamic coupling between linear and angular accelerations. Then, we calculate inverse kinematic solutions for these physically realizable motions, using the methodology described in [157]. The calculated joint angles can be executed with good repeatability, due to the high impedance of the actuators. The resulting overall body motion is also quite repeatable, as described later in the chapter. Portions of this chapter, particularly the results for double-support motions in Sections 3.5.2 and 3.7 also appear in [24].

²For example, recall the agile mountain goat on page 18.

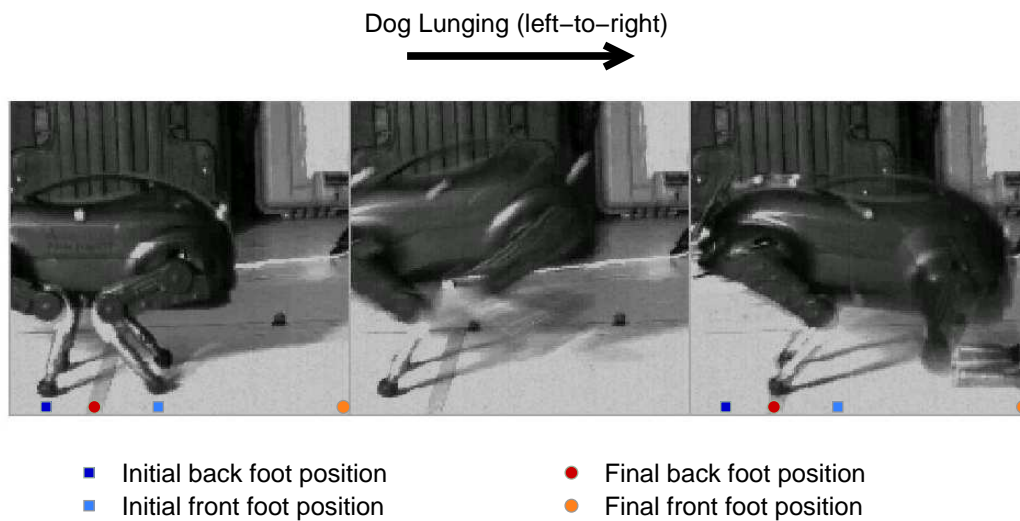


Figure 3-2: LittleDog executing a dynamic lunge

Frames were captured using a high-speed camera. Note that the back leg loses contact with the ground once the dog gains forward and upward momentum; it moves forward about 7 cm during the entire motion.

3.2 Toward dynamic maneuvers on rough terrain

Legged robots today move far too conservatively. A driving motivation in the study of legged machines is the promise that they will effectively navigate extreme terrain, making use of intermittent footholds and dynamic stability to go where wheeled vehicles cannot. Animals walk and run over incredible terrain with apparent ease [118], often leaping from one foothold to the next. In comparison, our machines move slowly and methodically, constrained by overly restrictive measures of dynamic stability. As described by Koditschek and Buehler, what we strive toward as roboticists are “dynamically dextrous robots”[97].

One major difference between animals and walking machines is the mechanical design. Many legged robots rely on large gear-reductions in order to obtain the necessary torque from electric motors with high-impedance transmissions. As a consequence, the joints in these robots have a high impedance, and the motions of these robots often appear rigid and stiff. In contrast, the impressively-dynamic capabilities of biological systems are often attributed to their use of compliant joints, which permit energy storage (and fast release)

and provide a level of mechanical robustness. While many researchers are interested in the mechanical design of low-impedance walking robots to achieve dynamic gaits [142], in this chapter, we consider solutions involving only control. In particular, we develop a motion-planning algorithm which allows a stiff, position-controlled quadruped robot to execute reliable dynamic maneuvers that are often associated with more compliant systems. By reasoning about the ground reaction forces and the dynamics of the passive degrees-of-freedom between the feet and the ground, we are able to design trajectories of the center of mass which maneuver the robot into and through a “bipedal” double-support phase. Our experimental results demonstrate the utility of these dynamic maneuvers in negotiating a variety of rough terrain in a motion capture environment.

3.2.1 Compliance versus stiffness

The success of many dynamic quadrupeds depends on a compliant (soft) leg design, to provide some level of natural (passive) dynamics, to moderate impacts and/or to facilitate force control. For example, there are now a number of compelling demonstrations of small, compliant quadrupeds which locomote quite well using only simple, oscillatory control (e.g., [96, 81]). Simplified control strategies have worked very well for a running quadruped with pneumatic legs [146], and significant refinement of these control ideas has more recently resulted in highly-dynamic legged locomotion and dynamic stabilization capabilities for BigDog [21]³.

Stiffness isn’t always an impediment, however. Although compliance has been instrumental in achieving dynamic gaits in a variety of robots, many of these more compliant quadrupeds have not simultaneously demonstrated the ability to achieve careful foot placements – a skill necessary for careful walking on particularly treacherous terrain. Quadruped robots capable of careful foot placement currently tend to move much more slowly (e.g., [48, 66]).

The robot used in the experiments presented here is LittleDog, a small quadruped built by Boston Dynamics, Inc. LittleDog was designed as a position-controlled robot capable of executing careful foot placements. Although there are many clear benefits of low

³BigDog was described in more detail in Section 1.2.4, page 30.

impedance for walking robots [142], a high-impedance robot such as LittleDog can still execute very dynamic motions, as demonstrated in our results.

3.2.2 Current state-of-the-art in kinodynamic planning

The phrase “kinodynamic planning problem” seems to originate with [29], which was later published in extended form as [40]. However, the basic problem was certainly considered before this, e.g., [133, 149, 156]. As the name suggests, kinodynamic planning involves simultaneous satisfaction of both kinematic and dynamic constraints on a robotic system. In its original formulation, the dynamic limitations considered were primarily actuator limits in velocity and/or acceleration, and the goal of the problem was one of determining a minimum-time solution connecting desired start and end states. In this context, the kinodynamic planning problem is an *optimal control problem* in its purest form. Early forms of the kinodynamic planning problem applied to robot *manipulation* (rather than to locomotion): for example, [41] extended kinodynamic planning from the problem of moving a point mass to one of a generalized, open-chain manipulator. In these early formulations, ground reaction forces were not generally considered; a robot arm could be assumed to travel with respect to a known, reliable coordinate system. In particular, many examples involved a fixed coordinate system, where the robot was mounted on a rigid, immovable base.

As research into humanoid robots grew throughout the 1980’s and 90’s, the *dynamic* constraints on kinodynamic planning grew to include requirements on the ZMP location, to achieve quasi-full actuation during locomotion, as discussed in Section 2.1.3. Although exact kinodynamic planning solutions have been developed for 1D [133] and 2D [28] problems, practical extensions to higher dimensionality require significant complexity. Higher-dimensionality solutions generally claim at best *approximate* optimality [156]. For a high-degree-of-freedom robot, such as a humanoid, kinodynamic planning is so challenging that a popular approach which is still used today uses randomized trees [108] to find *any* viable solution [110], rather than attempting to find an *optimal* solution.

In the current work, we extend kinodynamic planning to design motions with inten-

tional underactuation and coupled dynamics. For our quadrupedal robot, such underactuation is required to enable lunging and bounding type motions, which can in turn provide significant advantages both in attainable speed and in the extremity of traversable terrain⁴. As mentioned in Section 3.1, one approach in designing motions for high-degree-of-freedom robots involves first planning a series of statically-stable poses and then determining how rapidly they can be pieced together while maintaining a desirable ZMP stability margin. By contrast, our highly-coupled motions must be executed at a particular speed to achieve our desired underactuated motions. Previous work by others has also investigated the problem of appropriate time-scaling of trajectories in motion planning [74], emphasizing the idea that *slower* trajectories are not necessarily *safer* – and in fact may make a desired motion trajectory infeasible in some cases. This is precisely the case in our underactuated lunges, where accelerations which are either too large or too small will either topple the robot or fail to initiate a lunge.

3.3 The Learning Locomotion program

MIT is one of six research institutions participating in a DARPA (Defense Advanced Research Projects Agency) program entitled “Learning Locomotion”. The goal of this project is to generate computer algorithms which autonomously control the LittleDog robot to cross known but challenging terrain quickly and reliably. The format of the Learning Locomotion program allows for direct comparison of *control* strategies and algorithms in a head-to-head format, since all teams must compete using identical hardware (i.e., actuation and sensing). More technical details on robot operation and desired performance metrics are given below in Sections 3.3.1 and 3.3.2.

3.3.1 LittleDog hardware and environment

The LittleDog quadruped has a mass of about 2.4 kg and is actuated at 12 leg joints. Planning for desired joint trajectories is done on an off-board computer, which transmits desired position and velocity values for all joints to the robot at 100 Hz. A low-level PD control

⁴e.g., wide gaps, stepped terrain, or “hurdling” style obstacles

loop on the robot (running at 400 Hz) regulates the joints. All joints are actuated through high gear-ratio (84:1) transmissions, giving them high impedance to external disturbances. Round feet on LittleDog provide rolling contacts which can not produce ankle torque. The ground contact angle at the support legs becomes unactuated during double-support motions. The robot operates in a motion capture (mocap) environment to estimate the 6 DOF of the body and to detect upcoming, pre-scanned terrain boards.

The work presented in this chapter focuses primarily on the dynamic motions planned for Phase 2 of the Learning Locomotion project, between September 2007 and April 2008. Various contributions in writing the code to control the dog and in debugging and testing were done by fellow teammates on the MIT LittleDog project. They include: Russ Tedrake and Nick Roy (principal investigators), Alec Shkolnik, Sam Prentice, Steve Proulx, Khash Rohanimanesh, Michael Price, Olivier Chatot, Emma Brunskill, John Roberts and Lauren White.

3.3.2 Metrics for LittleDog

DARPA enforces strict go/no-go metrics for many programs, and the metrics for the Learning Locomotion program are listed in Table 3.1. Speeds and heights were selected assuming a nominal leg length of approximately 12 cm; this is a typical value for the ground clearance under the belly of the robot during walking. Overall performance measures for each team during DAPRA testing were based on the best two out of three “blind” test trials on each of seven particular terrain class types. Teams were given examples of each terrain class on which to practice both at their own sites and during periodic (monthly) off-site “spiral” tests. Final performance for each team was based solely on metric scores achieved during go/no-go testing on the blind terrain boards. Failure to complete a run resulted in a trial speed of zero. The original intent of the metrics was to achieve an average speed at or above the metric on each of the seven classes of terrain. Our overall speed average during final (off-site) go/no-go testing was 6.0 cm/s, and the MIT team achieved metric speeds on 6 out of the 7 different terrain types.

Table 3.1: DARPA metrics for the Learning Locomotion project
 Height values give estimates of the maximum change in height per nominal leg length (12 cm) to be expected in trials; speed values are prescribed DARPA metrics.

Attribute	[units]	Phase I	Phase II	Phase III
Speed	[cm/s]	1.2	4.2	7.2
Height	[cm]	4.8	7.8	10.8

3.4 Locomotion planning on rough terrain

Successful long-living (i.e., metastable) locomotion on rough terrain requires both high-level planning (of the path taken through the terrain and of the motion *types* selected along the way) and low-level planning (of the dynamics and kinematics which can achieve the general motion types we wish to perform). The work presented throughout this chapter emphasizes methodologies for generating reliable *low-level* solutions to motion planning. While both low-level and high-level planning are certainly essential to both animal and robot locomotion, we believe focusing on the lower level is of more imminent interest (and is also significantly more challenging) than high-level planning for several reasons. First, we provide evidence later in this thesis⁵ that near-optimal solutions to metastable locomotion can sometimes be found using a relatively short look-ahead on the terrain. Second, we would often intentionally like to leave our high-level planning options as open as possible, to be able to command a robot to perform particular, high-level tasks, such as accompanying humans or surveying terrain. Good low-level solutions must exist to allow a legged robot to operate as a simple “vehicle”, even if significant high-level path planning can simply be done by a human operator. We note that if these low-level tasks are designed well and have quantifiable performance metrics, then the task of a high-level planner *should* (ideally) become a relatively straight-forward one of making selections which results in the best estimated performance metric. For all of these reasons, exploring what is achievable in low-level control is a critical step in planning locomotion on rough terrain.

Throughout this thesis, low-level motion planning consists of two, fundamental tasks: (1) the design of the motion itself, and (2) the evaluation of its performance, by which we

⁵in particular, in Section 5.8, p. 154 of Chapter 5.

mean both estimates of the expected end state and estimates of the probability of success and/or distribution of end states which may occur. This chapter focuses on the design of reliable motions. This task of *evaluation* (i.e., quantification) of the performance expected is a significant challenge for a high-degree-of-freedom robot such as LittleDog, traversing complex terrain. We discuss the problems of evaluation and optimization of low-level motion control for LittleDog briefly in Chapter 6.

3.4.1 Path planning

Although our focus in this chapter is on the task of planning low-level (step-to-step) plans once a global path has already been dictated, we will consider issues in global path planning briefly in this particular section. Implicit in the idea that initial conditions are forgotten rapidly for metastable locomotion is a guarantee that “greedy” short-sighted decisions will not lead to a large, unforeseen penalty several steps later. There are many problems, such as chess, where a long horizon look-ahead is required to estimate the comparative values of short term actions. It is beyond the scope of this thesis to provide the exact conditions which guarantee that such long-horizon planning is or is not required, and (regardless of such guarantees) it is of interest to develop low-level solutions where some user-generated, global plan has simply been dictated. However, we will still provide two general reasons to believe that particular foot selection for legged locomotion along a prescribed global path can generally be accomplished with a relatively short time-horizon look-ahead.

One piece of evidence for the near-optimality of short-horizon foothold planning is that many animals can navigate quite robustly on complex and stochastic terrain when only a step or two of terrain knowledge is available. This is often the case when hiking, for instance, where paths may turn suddenly or the ground profile may be obscured by leaves and underbrush. Also, dogs and horses trained for highly-agile locomotion must respond within a footstep or two to commit to particular obstacles, because they simply do not know what path will be taken until they are signalled. In dog agility, courses also include obstacles, such as tunnels, which significantly hide upcoming terrain, and in competitions, dogs are required to run an agility course without prior knowledge of the specific (global)

course.

A second reason to believe short-sighted planning may often work well for legged robots comes from analyzing results for a controlled compass gait walker on rough terrain, which will be presented in detail in second half of Chapter 5. Here, we demonstrate that both an infinite-step look-ahead policy and an optimal one-step look-ahead policy result in continuous walking on known, wrapping terrain. In that work, the vertical magnitude of the terrain features was iteratively adjusted until the optimal policy with infinite look-ahead could just barely succeed. The success of a one-step policy on the same terrain was correspondingly both surprising and highly suggestive to the author: although additional evidence is required, it seems both intuitive and reasonable to suspect that the nature of legged locomotion often requires only a short look-ahead. The success of short-horizon planning is in part due to the impulse-generating nature of “taking a step”: energy can be imparted or dissipated during contact with the ground at each step, providing a natural discretization for motion planning and control.

Above the level of individual foothold planning, some global path for the robot must also be determined. In many cases, there may not be much choice in a global path, either because obvious “paths” (e.g., manmade) preexist or because the robot is required to traverse a particular path or to go between specified way-points.

To achieve the DARPA metrics for the Learning Locomotion program, teams are provided with significant pre-planning time, to process terrain and to compute a global path or strategy for traversal to the goal. Global path planning for LittleDog is a high-dimensional task; the robot has 18 DOF and must plan continuous motion for approximately 45 steps to get from start to goal on most terrain. Searching for “optimal” solutions is arguably impossible within the 90-second planning period provided. Instead of finding an optimal path, we instead look for a “reasonable” solution, using a depth-first searching algorithm. Several flavors of algorithm have been employed throughout Phases 1 and 2 of the project. Qualitatively, the most computationally intensive tasks in planning are collision detection and inverse kinematics. We use a heuristic cost function to estimate the “easiness” of finding locally collision-free and kinematically-feasible poses to identify an approximate global corridor to the goal; this reduces the search space for particular body pose and foothold

selections.

For the most difficult terrain⁶, we use the longest look-ahead, planning entirely from start to goal. On terrains where particular, extreme obstacles can be most-easily identified, we employ the most dynamic behavior, obtain the fastest speeds, and use a much shorter look-ahead. We also use extensive “re-planning” throughout on such terrain, as necessary. In Phase 2, such terrains consisted of discrete obstacles (e.g., a gap or Jersey barrier, as shown in Figure 3-17) surrounded by flat terrain, so that footholds are only planned with care at the initiation of a lunge, to be sure of appropriately clearing the obstacle. In future work on more generalized rough terrain which includes intermittent obstacle-like features, we are hopeful that dynamic motions can be planned to connect a wider variety of start and end positions, so that any required look-ahead is not essential for fast traversal.

Figure 3-3 shows an example of a multi-step path planning process developed by the author for Phase 1 of our LittleDog DARPA work. It involves a 3-step process: 1) planning an approximate path for the robot body, such that body collisions with the terrain are minimized, 2) filtering the initial path to look for swing foot trajectories which minimize the chance of terrain collision, and 3) creating “railroad tracks” to optimize the stance width. These railroad tracks are placed to either side of a nominal center of body (COB) path and are then pinched inward where appropriate to avoid dangerous or collision-prone footholds. Footholds are subsequently selected which approximately follow these tracks.

The reader may begin to suspect, from the preceding paragraphs, that avoidance of terrain collisions has been a significant issue on particular terrain boards. This is absolutely so and continues to be problematic. Because of the lack of compliance of the position-controlled robot limbs, terrain collisions can cause nearly impulsive disturbances, knocking it backward and imparting destabilizing yaw and roll. This may cause immediate failure during a trial, if a robot falls to the ground sideways, or it may deflect the robot so far off course that significant replanning is required to get back on the planned path. Upcoming terrain collisions are computationally expensive to predict. This issue was addressed toward the end of Phase 1 through a support vector machine (SVM) approach [42] with mixed results, but collision avoidance continues to be a significant issue in planning on

⁶That is to say, by definition, “on the terrain on which we performed the worst.”

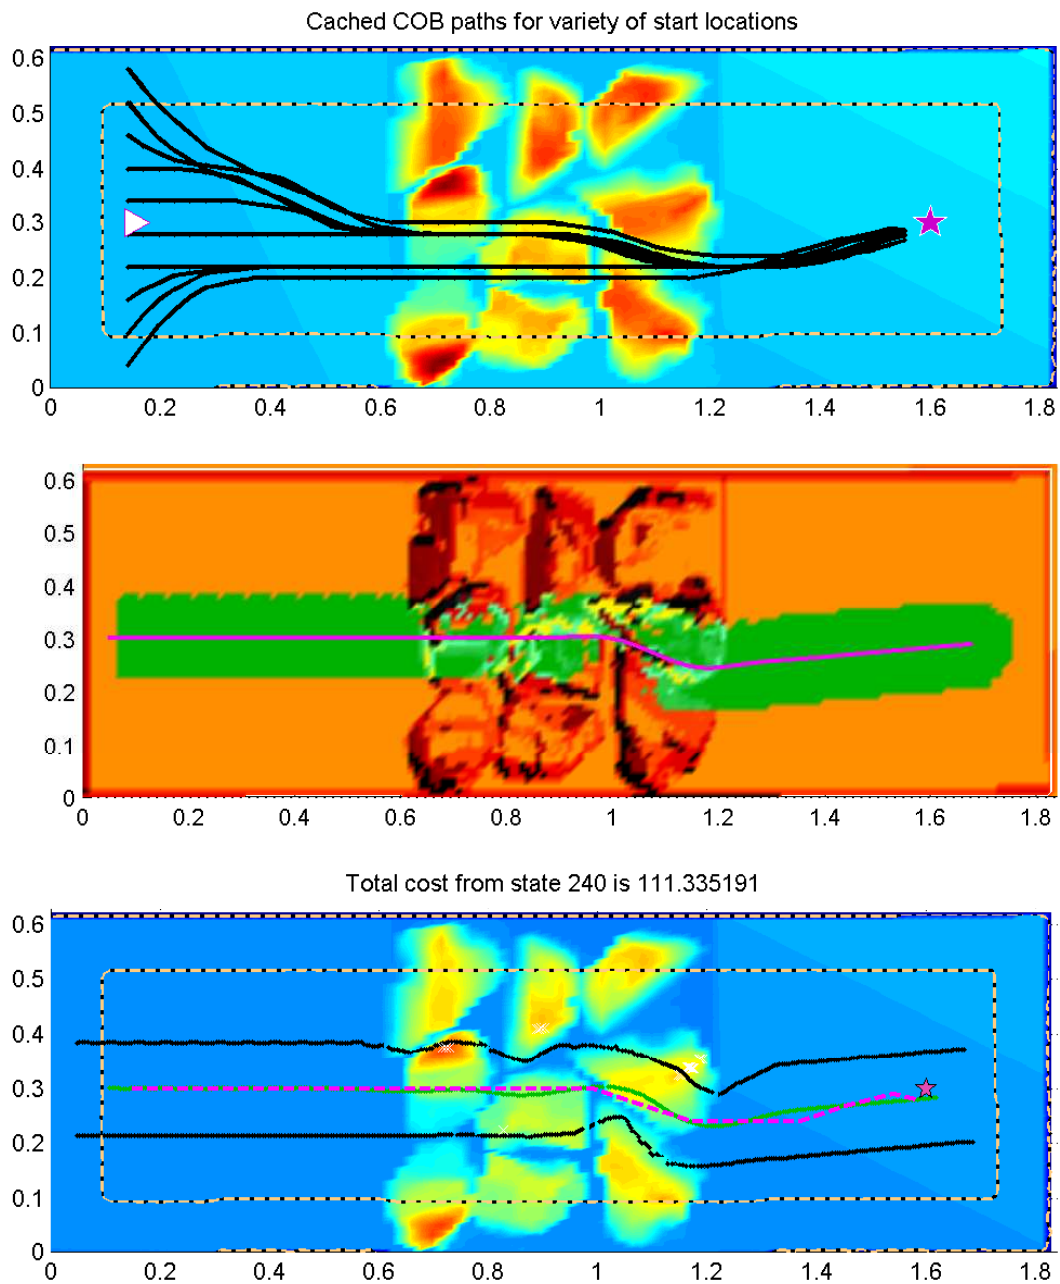


Figure 3-3: Global path planning on rough terrain

Above, the Phase 1 path planner is used to analyze Phase 2 terrain. This global path planner involves a 3-step process, described in the text. A variety of other approaches were also implemented and tested throughout Phases 1 and 2 of Learning Locomotion.

rough terrain with significant vertical features. We note that the utility of global path planning for LittleDog emerges primarily because of the computational expense of collision detection and inverse kinematics, and we strongly suspect such planning is not inherently prerequisite, given the dynamic nature of walking itself.

Although we do not provide extensive detail on the various path planning strategies which were used in Phase 2 in this thesis, we will note that they involved a significant, collaborative effort. Sam Prentice, Alec Shkolnik, John Roberts and Russ Tedrake all contributing significantly toward successful path planning in Phase 2.

3.4.2 Motion planning strategy

To negotiate rough terrain, the planner for LittleDog can select among motion types which range from a slow, statically-stable walk to more dynamic gaits and lunges. Dynamic motions allow the robot to move more quickly, but they also require more care in planning. To achieve fast locomotion over a variety of terrain, we operate in two dynamic regimes: 1) fast walking and 2) double-support motions. Double-support motions are our focus here, but our development of both motion types are based on the same, central idea of attaining desired ground reaction forces to achieve desired trajectories of the body of the robot.

Our ability to execute so-called “open-loop” joint trajectories depends inherently on having both (1) a dynamically feasible, co-compatible set of motions and (2) a good reference clock. The ability to regulate the timing of a desired reference trajectory is such an important aspect of many successful open-loop strategies that a large class of open-loop control may more accurately be considered as *time-feedback* control as addressed recently in [109].

3.5 Crawl gaits

In this section, we discuss gaits in which “full actuation” is assumed during locomotion. During such gaits, we will be concerned with the location of the center of pressure (COP), otherwise called the zero-moment point (ZMP), as previously discussed in Section 2.1.3 on page 38. In the limit as accelerations (linear and angular) of the body are small, the

ZMP will be nearly collocated with the projection of the center of mass; traditional motion planning often begins with this limiting assumption and proceeds by designing overall kinematic poses for a robot, as mentioned in Section 3.2.2. This simplification is convenient but also limits the overall speed and performance capabilities of a legged robot. It is more elegant and practical to employ a technique which works over a range of motion speeds. This section summarizes our methodology for obtaining such speed-variable walking gaits, planned using the assumption of fully-actuated motion; this will then be extended toward underactuated motion planning in Sections 3.6 and 3.7.

It is important to note (briefly) at the outstart of this section that our crawl gaits with LittleDog must be planned to avoid violation of the fully-actuated assumption in several ways. For example, (1) feet can only push, not pull, where they contact the ground, (2) lateral accelerations are limited by “friction cone” constraints, (3) inverse kinematics must be planned to ensure the feet maintain their (rolling) contact with the ground over time, and (4) disturbances such as unplanned swing leg collisions can impart significant, impulsive disturbances which can begin to topple the robot. However, if the terrain is appropriately well-characterized and the motions are intentionally planning within particular, dynamic constraints, using a fully-actuated model yields reliable results in this slow-walking regime.

In a crawling gait, there are always at least three feet in contact with the ground at any given time, and one can always define a support polygon from the convex hull of these support points [38]. Keeping the center of pressure of the robot strictly within this support polygon at all time ensures that the robot will never begin to topple, and so crawling gaits are often employed in an attempt to move conservatively (with minimal risk of falling over). Such a gait is also sometimes referred to as creeping [116].

When the magnitude of the accelerations of the body remain low, it is practical to consider crawling gaits as a series of connected, statically-stable snapshots in time, because the center of pressure is essentially collocated with the projection of the center of mass downward. Faster walking is considered “dynamic” when the accelerations become significant and this assumption is no longer valid. We discuss this regime of “ZMP walking” for LittleDog in more detail just ahead in Section 3.5.2. Referring to the case of level walking, as depicted in Figure 3-4, we can mathematically represent the (x, y) location of the center




Figure 3-4: Center of pressure during level-ground walking

The center of pressure (COP) is not collocated with the center of mass (COM) in this example. Assuming the height and angular rotations of the body remain constant, there is an instantaneous acceleration acting on the body in the direction from the COP to the COM in the x-y plane.

of pressure⁷ as obeying:

$$(x_{com} - x_{cop})(\ddot{z}_{com} + g) = (z_{com} - z_{cop})\ddot{x}_{com} \quad (3.1)$$

and

$$(y_{com} - y_{cop})(\ddot{z}_{com} + g) = (z_{com} - z_{cop})\ddot{y}_{com} \quad (3.2)$$

where g is the (positively-signed) *magnitude* of gravity: 9.8 m/s^2 . Note that these relationships will be reduced as given in 3.5 when motions are planar and on level ground. Following from 3.1 and 3.2, approximate collocation of the COM projection with the COP means that the lateral distance between the COM and COP is small with respect to the scale of the robot. Taking the height of robot from the ground, $z_{com} - z_{cop}$, as a representative scale for non-dimensionalization, the static assumption is only valid when both lateral accelerations are small, i.e.:

$$\left| \frac{\ddot{x}_{com}}{(\ddot{z}_{com} + g)} \right| \ll 1 \quad (3.3)$$

⁷See also Eqns. 2.8 and 2.9.

and

$$\left| \frac{\ddot{y}_{com}}{(\ddot{z}_{com} + g)} \right| \ll 1 \quad (3.4)$$

3.5.1 Statically-stable walking

In some situations, slow, deliberate walking is practical. Examples include negotiating extreme terrain, for instance, where footholds may be intermittent or where body or leg collisions must be carefully avoided. Figure 3-5 shows an example of deliberate, careful walking. For this speed of walking, the projection of the center of mass onto the ground is very close to the actual center of pressure, and these gaits are often (correspondingly) referred to as “statically stable”.



Figure 3-5: LittleDog walking on pegs

3.5.2 Dynamic, ZMP-based walking

Fast walking is achieved by planning the trajectory of the zero-moment point (ZMP) [178, 179] for one or two upcoming steps at a time. Our method is inspired by the ZMP control for bipedal walking presented in [86], which uses a preview control method originally described in [91]. Calculation of the forward problem, mapping body motions to resulting ZMP, is straight-forward. Equation 3.5 below gives the solution for the instantaneous ZMP

location for the planar robot depicted in Figure 3-6.

$$x_{zmp} = \frac{x_m(\ddot{z}_m + g) - z_m\ddot{x}_m}{\ddot{z}_m + g} \quad (3.5)$$

This forward solution was already presented in Chapter 2 (Eq. 2.9). To ensure (3.5) is valid in this pre-lunge orientation, we assume all four legs are in contact with the ground and that motions are planned to maintain the ZMP strictly within the support polygon defined by the point contacts of the feet at the ground. In the planar model in Figure 3-6, we also assume that the body is not allowed to rotate, so that rotational acceleration terms are not included in (3.5). In the cart-table model presented in [86], the resulting point-mass representation of the robot is also constrained to travel at a constant height, i.e., so that z_m is constant and $\ddot{z}_m = 0$. The solution to the forward problem for this model, taking a particular body trajectory and calculating the corresponding zmp, is then extremely simple:

$$x_{zmp} = x_m - \frac{z_c\ddot{x}_m}{g} \quad (3.6)$$

where z_c is the constant height along which the mass is constrained to travel.

Calculation of the inverse problem, deriving the required body motions to approximate a desired zmp trajectory over time, is more difficult. The preview control method outlined in [86] finds approximate solutions to the inverse problem by 1) augmenting the second-order system to include acceleration as a state variable (resulting in a third-order system), so that x_{zmp} in (3.5) can be represented as a state output, 2) assuming z remains constant over time, and 3) solving a linear-quadratic regulator (LQR) problem to calculate the $x_m(t)$ trajectory which results in the lowest weighted quadratic cost in error in ZMP location and (to a lesser extent) control effort. Because rotations and changes in height (z_m) are assumed to be negligible, we can decouple motions to plan in the x-z and y-z planes, independently. Finally, joint angles are planned using inverse kinematics appropriate to the robot to attain the desired trajectories in x and y . High-impedance actuation insures high fidelity in commanded joint trajectories, and this in turn results in ground reaction forces close to the pre-planned values.




Figure 3-6: Planar ZMP model for fully-actuated motion

The image at left is the “cart-table model”, taken (with the author’s permission) from [86]. Here, the robot is represented as a point mass (illustrated by the cart) which is constrained to move only in the horizontal direction (along a table top). The table has a base of support, which corresponds to the support polygon of a robot. Maintaining the ZMP strictly within this support area ensures the robot will never begin to topple. At right is our corresponding planar model for LittleDog. If the robot body remains at constant height and does not rotate, its dynamics are entirely captured by the cart-table model..



Figure 3-7: Rocky terrain, negotiated with a dynamic gait

3.5.3 Design methodology

This section outlines our method for generating walking gaits which maintain a positive ZMP margin during locomotion, to avoid unplanned, underactuated rotations of the body.

ZMP and COB trajectory planning

For ZMP-based walking, we begin planning by determining a set of sequential, static poses, much as we would for the much slower motions in Section 3.5.1. Although our fast walking motions may use the same footholds used in slow, statically-stable walking, the body trajectory must now be adapted to maintain a stable ZMP trajectory. Using preview control, we specify a desired ZMP trajectory which connects points within sequential support polygons. The speed at which the desired ZMP travels and the (approximately constant) height of the center of mass determine the resulting COM trajectory.

We note here that we extend this same strategy to plan double-support pacing and trot-like walking gaits. Although there is no longer a support polygon in either of these gaits, we can still plan to attempt to stabilize the center of pressure along an infinitesimally narrow line of support during locomotion. Figures 3-8 and 3-9 illustrate corresponding x-y trajectories for pacing and trot-walking, assuming a center of mass height of 15 cm. The trajectory-generation machinery we use is very similar for walking both with and without a support polygon. The primary difference is that we have much more freedom in where the ZMP travels over time in the case where three feet provide a polygon of support. The reader should also note that the ZMP can move *instantaneously*, so the desired trajectory can include step changes in location, if desired.

Six-DOF motion planning for body on rough terrain

Once the center of mass position is calculated, we must still determine the three rotational degrees of freedom for the robot over time. To do so, we simply connect the planned static poses (from Section 3.5.1) using appropriately scaled half-cosine waveforms. In practice, the resulting overall motions are satisfactorily smooth.

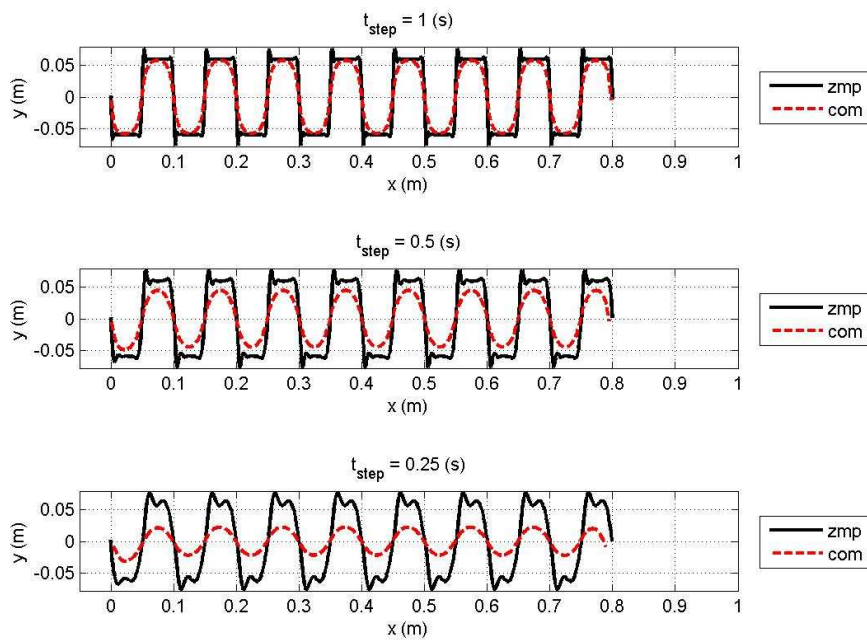


Figure 3-8: COM vs. ZMP during pacing

Swing leg trajectories

Swing leg trajectory design is an important consideration for two primary reasons. 1) Limits in joint velocities should be respected, to avoid operating any motors at saturation, and 2) we wish to avoid collisions, both with the ground and (proprioceptively) with other body parts of the robot. A variety of swing leg trajectories were developed to achieve smooth but fast motions on different terrain. On rough terrain, the trajectory follows the convex hull of the terrain, and speed is adjusted to respect the velocity limits of the motors.

On extremely rough terrain, the achievable swing leg speed often determines the limit to the overall velocity of the robot. The danger of feet sliding, due to violation of the friction cone, is not nearly as significant as the danger of swing foot collisions, which can produce impulsive disturbances on the entire robot. In practice, such collisions account for a significant fraction of the total failures (robot falling) when walking on rocky terrain.

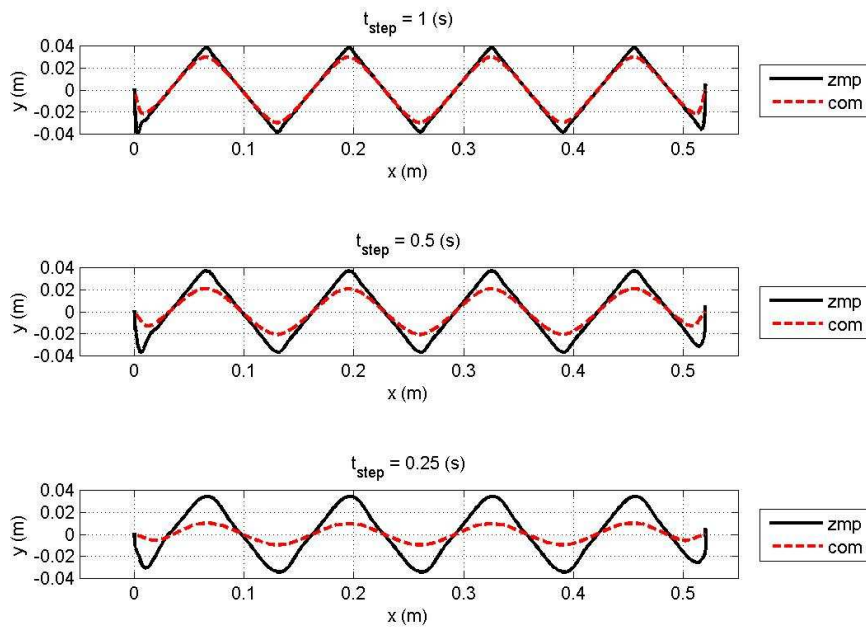


Figure 3-9: COM vs. ZMP during trot-walking

3.5.4 Results and discussion

Our results for open-loop ZMP walking represent a deliberate balance between the speed of locomotion and control calculation (which is excellent), and the accuracy of the final COM position attained after walking (which is good, but inherently suffers from cumulative errors). The incorporation of a low-level slip recovery strategy and/or of force feedback in the control may provide significant improvements in stability in future work with LittleDog during Phase 3 of Learning Locomotion.

3.6 Double-support walking

Including a brief, planned double-support phase allows us to increase the speed at which we can execute the same, planned sequence of poses⁸. Although the preview control method is an elegant and straight-forward way to plan COM trajectories here and has inspired

⁸The same strategy described in Section 3.5.1 is used to plan end-poses.

our own approach, we intentionally use a more aggressive motion planning strategy for this gait, to maximize speed while maintaining stability. Specifically, the preview control solution results in slower initial accelerations than necessary, and it is actually possible to “push-off” with significantly more force with the rearmost leg to accelerate forward faster at each step.

In practice, we connect all six degrees of freedom using a half-cosine waveform. Diagonally opposing front and rear legs swing to new positions during this motion. The front leg lifts off the ground at the start of the COM motion; the rear leg lifts off halfway through, when the forward acceleration of the body is planned to be exactly zero (i.e., midway through the cosine half-wave). For the remainder of the motion, the body has a net deceleration, causing it to pitch forward somewhat about the line of support of the two stance feet. At the end of each this motion, the robot intentionally pauses for a short interval (on the order of 0.05 seconds) with all four feet on the ground to obtain an updated estimate of the six degrees of freedom of the robot body from motion capture.

The period of this motion is set to be as fast as possible while still maintaining a small ZMP margin as the rear leg pushes off. We then calculate the expected pitching moment of the robot, assuming the mass moves as planned in x and y . The time to travel from one pose to another is generally about 0.5 seconds, and the robot is underactuated for approximately the second half (0.25 sec) of the motion. The maximum distance the center of mass travels in one step is about 0.1 meters. During the 0.25 seconds of underactuated motion, the predicted accumulation in (nose-down) pitch is about 7 or 8 degrees, which is within a degree or two of what we observe during experimental runs on the actual robot, as illustrated in Figure 3-11.

3.7 Double-support lunging and pacing motions

This section presents a methodology for planning repeatable *highly-dynamic double-support motions* (i.e., bipedal phases) for LittleDog. Planning such underactuated motions is inherently difficult: the ground contact angle of the foot is unactuated (no ankle torque), and the dynamics are nonlinear and strongly coupled. To approach the problem, we use a simpli-

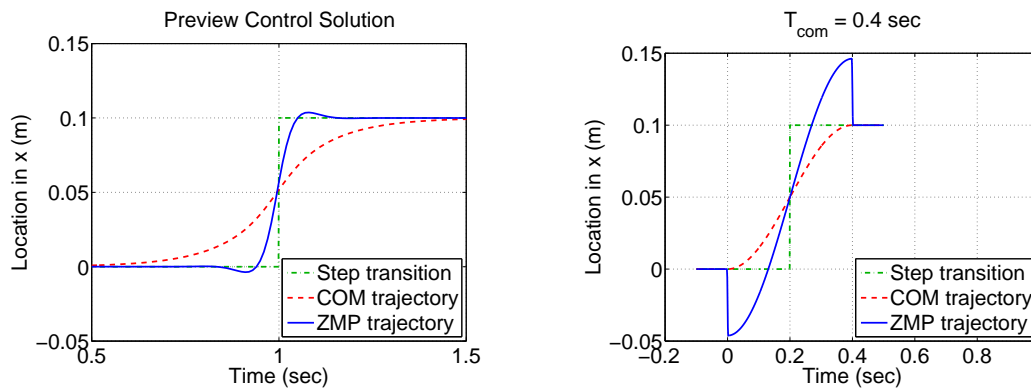


Figure 3-10: The **inverse** (left) and **forward** (right) problems ZMP trajectory. The preview control solution at left minimizes errors in ZMP location, but it also requires 1.0 second to move the COB, compared with 0.4 seconds for the half-cosine solution for COB motion at right. For faster speed, we use the half-cosine trajectory shown at right rather than the preview control solution, which results in an instantaneous jump in ZMP at the start of motion. This jump is planned to remain small enough so that the ZMP never exits the support polygon during triple-support phases of ZMP-based walking gaits.


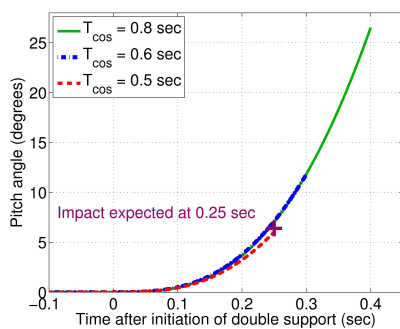


Figure 3-11: Predicted and observed pitch during underactuation in trot-walk. For a half-cosine wave transition in body motion of 0.1 meters in 0.5 seconds, we predict an accumulated pitch of about 7° (at left) and actually observe a pitch of about 8° during actual trot-walking (as captured in the image at right).

fied model of the dynamics and assume all motion is restricted to the plane perpendicular to the axis connecting the two support feet. We apply the same basic equations used in zero-moment point (ZMP) analysis to this 2D model to calculate the ground reaction forces required to obtain desired rotations and translations of the body. For this second-order physical solution, good solutions can be obtained simply by planning a step input in acceleration in the x direction for a particular time window, as described in Section 3.7.2 ahead. The inverse kinematic solutions for the corresponding joint angle trajectories are commanded using a low-level PD controller. Once again, high gear-ratio joints ensure low errors in desired joint angle. The resulting motions of the robot are both close to the planned behavior and highly repeatable.

3.7.1 Modeling double-support motions

Double-support motions include lunging, where the back (or front) legs support the robot as it rears (or bucks) up, and pacing, where the legs move in sequential left and right side pairings to support the robot as it moves. The two images in Figure 3-12 show a planar model of the dog rearing up in a lunge. In the real robot, there is an unactuated, rolling contact between the ground and each of the spherical feet (radius ≈ 1 cm) which we model as a fixed point contact. Unlike double-support *walking*, in which we plan motions with near-zero moment about the point of ground contact, in these highly-dynamic motions, we intentionally plan for a finite moment to rotate the body about the support point in our planar model. Although our goal here is different, our method still relies on an approach similar to the preview control method for the ZMP. Our goal is to plan feasible ground reaction forces *to achieve a desired rotation* of the body while maintaining feasible kinematic joint angles. The forces in x are limited by friction:

$$|F_x| \leq \mu F_z \quad (3.7)$$

where $F_z = m(g + \ddot{z})$, $F_x = m\ddot{x}_m$, and μ is the coefficient of friction. We assume $\mu \approx 0.3$ on our terrain and that $|\ddot{z}| \ll g$; correspondingly, we plan our linear accelerations of the robot such that $|\ddot{x}_m| \leq 3$ [m/s²].




Figure 3-12: Planar model of LittleDog in a double-support lunge.

At top, the planar model includes 2-segment legs. The legs are modeled as massless, so only the location of the end point of each legs is important. At bottom, the geometric locations of the end points are more easily represented by modeling each leg as a single, variable-length leg segment which can rotate about the hip over time.

In our simplified model, all motion is strictly planar. The body has a fixed inertia and mass, while the legs are modeled as massless. When there is a single, unactuated contact point for the planar robot, as depicted in Figure 3-12, we can equivalently represent the actuated joint angles (shown in the first illustration) with a length from the hip to the ground, L_1 , and an inter-leg angle, θ_1 (as shown at right). Using L_1 and θ_1 as the actuated degrees of freedom, the kinematic equations for the location of the mass are:

$$x_m = L_1 \cos \beta + L_a \cos \alpha - L_b \sin \alpha \quad (3.8)$$

$$z_m = L_1 \sin \beta + L_a \sin \alpha + L_b \cos \alpha \quad (3.9)$$

where $\beta = \pi + \alpha - \theta_1$. Taking derivatives of (3.8) and (3.9):

$$\dot{x}_m = -L_1 \sin \beta \dot{\beta} - L_a \sin \alpha \dot{\alpha} - L_b \cos \alpha \dot{\alpha} + \dot{L}_1 \cos \beta \quad (3.10)$$

$$\dot{z}_m = L_1 \cos \beta \dot{\beta} + L_a \cos \alpha \dot{\alpha} - L_b \sin \alpha \dot{\alpha} + \dot{L}_1 \sin \beta \quad (3.11)$$

and

$$\begin{aligned} \ddot{x}_m = & -L_1 \cos \beta \dot{\beta}^2 + (-L_a \cos \alpha + L_b \sin \alpha) \dot{\alpha}^2 - L_1 \sin \beta \ddot{\beta} + \\ & -(L_a \sin \alpha + L_b \cos \alpha) \ddot{\alpha} - \dot{L}_1 \sin \beta \dot{\beta} + \ddot{L}_1 \cos \beta \end{aligned} \quad (3.12)$$

$$\begin{aligned} \ddot{z}_m = & -L_1 \sin \beta \dot{\beta}^2 - (L_a \sin \alpha + L_b \cos \alpha) \dot{\alpha}^2 + L_1 \cos \beta \ddot{\beta} + \\ & +(L_a \cos \alpha - L_b \sin \alpha) \ddot{\alpha} - \dot{L}_1 \cos \beta \dot{\beta} + \ddot{L}_1 \sin \beta \end{aligned} \quad (3.13)$$

We can use (3.12) and (3.13) along with the torque balance below, (3.14), to solve for the (underactuated) rotational acceleration as a function of the actuated degrees of freedom and their time derivatives.

$$\ddot{\alpha} = \frac{m}{J_y} (\ddot{x}_m z_m - (\ddot{z}_m + g)x_m) \quad (3.14)$$

where J_y is the moment of inertia in the x - z plane. Note that when the left-hand side of (3.14) is zero (no net moment on the body), this relationship reduces to the ZMP relationship in Equation 3.5, where x_{zmp} is now constrained to the single point of ground contact (defined as $x = 0$). Expanding (3.14), we find:

$$\begin{aligned} \ddot{\alpha} = & \{ z_m [-L_1 \cos \beta \dot{\beta}^2 + (-L_a \cos \alpha + L_b \sin \alpha) \dot{\alpha}^2 + \\ & L_1 \sin \beta \ddot{\theta}_1 + \ddot{L}_1 \cos \beta - \dot{L}_1 \sin \beta \dot{\beta}_1] + \\ & x_m [L_1 \sin \beta \dot{\beta}^2 + (L_a \sin \alpha + L_b \cos \alpha) \dot{\alpha}^2 + \\ & L_1 \cos \beta \ddot{\theta}_1 - \ddot{L}_1 \sin \beta + \dot{L}_1 \cos \beta \dot{\beta}_1 + g] \} \div \\ & \{ x_m [L_1 \cos \beta + L_a \cos \alpha - L_b \sin \alpha] + \\ & z_m [L_1 \sin \beta + L_a \sin \alpha + L_b \cos \alpha] + (J_y/m) \} \end{aligned} \quad (3.15)$$

Equation 3.15 gives the forward solution to the underactuated dynamics given known values of the actuated degrees of freedom, L_1 and θ_1 , and we can use this relationship to simulate the expected motions for a given plan of actuation. However, (3.15) does not lend itself to an efficient calculation of the actions required to, for instance, achieve a particular maximum angle α of rotation or a desired x displacement for the front feet. To achieve goals such as this, we consider a simple family of attainable ground reaction forces and iterate in a few steps to obtain approximate solutions which obey the required coupling

equations. This method is described in more detail in Section 3.7.2 below.

3.7.2 Design of double-support trajectories

To achieve a double-support phase, the most important element in our method is achieving a forward acceleration in x . The achievable magnitude of the force is limited by friction constraints to approximately $|\ddot{x}_m| < 3 [m/s^2]$, as previously discussed. To relate the ground forces to body motions, we now consider the other two equations of motion (along with (3.14)) which completely describe the dynamics of the system:

$$\ddot{x}_m = \frac{1}{m} F_x \quad (3.16)$$

$$\ddot{z}_m = \frac{1}{m} F_z \quad (3.17)$$

In an initial iteration, we will assume z_m is constant, and correspondingly $\ddot{z}_m = 0$. Equation 3.14 then relates the motion in x to the desired rotations, α , with no coupling to the z motion. This approximation surprisingly captures the gross dynamics in α quite well, particularly during the initial half of the motion, when x_m is small and when z_m has not changed significantly.

Figure 3-13 shows simulation results for our simple, planar model of the robot. The leftmost column shows the x position over time and its derivatives. The acceleration \ddot{x}_m is chosen first. As shown in the upper left, the time-optimal solution for fast rotation becomes a trivial bang-bang command in x acceleration, limited in magnitude to remain less than 3 (m/s²). The magnitude of the x acceleration essentially sets the slope of the change in α , shown in the lower right plot. The duration of the positive acceleration, \ddot{x}_m , affects the impact time. Dashed lines in middle and right columns show a simplified approximation to the solution, which neglects z motion. The more exact (solid line) solution can be found by enforcing some kinematic assumptions on the relationship between x , z and α . For example, one can assume that the bent-leg length in Figure 3-12, L_1 , remains constant. Doing so, we then simulate the dynamics using the same \ddot{x}_m commands and solve at each dt for kinematically compatible solutions for z and α . The final predicted trajectories for

x_m , z_m and α are shown as solid lines in Figure 3-13. Inverse kinematic solutions to obtain joint angles from these trajectories are then calculated as described in [157].

The selection of an initial condition in x is an important consideration which we have not yet addressed. To increase the length of time over which some \ddot{x}_m can be maintained, we wish to have $\dot{x}_m < 0$ (negative velocity) and x_m relatively small (for greater rate of rotation). In the example shown in Figure 3-13, $x_o = 0.02$ [m] and $\dot{x}_o = -0.15$ [m/s]. We select these values before the solution process described above, using practical considerations of kinematic feasibility and of the static stability of the initial stance to guide selection.

Once desired joint trajectories have been planned, we can simulate the resulting forward dynamics which would ideally occur using the forward solution for $\ddot{\alpha}$ in (3.15). In particular, we can examine the effects of adding disturbances (e.g., external forces or errors in commanded joint trajectory) on the resulting motions of the robot. Our simulations and experimental results both show good repeatability for our resulting kinodynamic motion plans, as discussed below.

3.7.3 Result summary for 3 double-support motions

There are three pairwise combinations of legs which can be used for double-support locomotion on a quadruped: diagonally-opposed legs (trot-walk); the left pair versus the right (pacing); and the front pair versus the back (lunging). Discussions of individual design details and results for each case are given below in Sections 3.7.4, 3.7.5, and 3.7.6, respectively.

The primary concern in executing double-support motions is that the pitch angle, α , is regulated to clear nearby terrain while avoiding extreme pitch which would topple the robot. Trot-walking is used on a wide variety of rough terrain, although the overall speed and time of double-support are both reduced as terrain becomes more extreme. Lunging motions are used on more extreme regions of terrain where the robot must negotiate upcoming obstacles, such as gaps or steps. Finally, although a *continuous* forward-going pacing gait provides no speed advantage over the trot walk, we do frequently employ *intermittent*




Figure 3-13: Simulated dynamics for a planar dynamic lunge. State variables and their derivatives over time for simple, second-order “brick” dynamics. Given a particular initial condition, the global pitch angle, α , depends much more heavily on acceleration in x over time than on changes in z .

pacing-class motions to reposition the overall stance of all four legs rapidly. Our lunging motions generally result in significant displacements (primarily in the forward direction). Pacing motions produce smaller lateral displacements in the center of body, but they allow for greater flexibility in direction of translation and can efficiently reorient the rotation angle in the x - y plane: the robot can turn in place, step sideways, step forward, or combine these displacements during a single pacing maneuver. In practice, our lunging motions are pre-calculated for few particular, dramatic types of terrain negotiation, while pacing steps are generated on-the-fly for generalized, rapid repositioning of all four feet.

3.7.4 Diagonal trot-walk

On flat ground, the speed of the diagonal trot-walk is 12.5 cm/s, which is approximately three times the Phase 2 metric (of 4.2 cm/s). Figure 3-14 shows LittleDog during the double-support phase in level-ground trot-walking.



Figure 3-14: Snapshot of double-support during dynamic, “diagonal” walk

The speed of the swing leg trajectories and the time of double-support overlap between diagonal feet are both carefully set as functions of the height variations of terrain and the lateral distance traveled of the center of mass. Figure 3-15 shows the four Phase 2 terrain types on which we employ this trot-walking gait.



Figure 3-15: Four terrain types using a diagonal trot-walk

Our gait on four of the seven Phase 2 terrain types employs a time-scaled diagonal trot-walk. Starting with the top, left image, our top speeds on these four terrain types are as follows: 10.3 cm/s on gray modular terrain, 8.4 cm/s on brown modular terrain, 7.7 cm/s on sloped boards, and 5.6 on molded rocks.

3.7.5 Dynamic lunge

Our results for lunging agree reasonably well with the expected motions predicted by our simple, planar model. Figure 3-16 shows the pitch angle α during repeated trials of the dynamic lunge. The theoretical pitch was obtained by simulating the desired joint trajectories in the legs and solving the dynamics over time using the coupling relationship in (3.15). Causes of variability include unmodeled dynamic effects such as backlash effects, kinematic asymmetries in the robot, variable friction, and packet loss in communication.




Figure 3-16: Comparison of experimental results and theory for a dynamic lunge

Lunging is a practical way of negotiating significant terrain obstacles rapidly. We use our dynamic lunge on the remaining three of the seven Phase 2 terrains. Figure 3-17 shows two examples: a hurdle-style vertical obstacle, called a “Jersey barrier”, and a negative obstacle (i.e., gap). Figure 3-18 shows frames from the remaining terrain, which involves a step-climbing sequence. Here, the overall speed of traversal is limited by the time required to carefully retract the rear legs during climbing and to position all four feet before each lunge.

We have also negotiated more extreme terrain in lab. Figure 3-19 shows frames from video taken of LittleDog using a dramatic lunging to climb over a barrier which is too tall to allow the legs to step over it directly. This lunge is highly repeatable in successfully



Figure 3-17: LittleDog lunges across Jersey barrier (left) and gap (right)

The dynamic lunge described in Section 3.7.5 can be used to allow both front feet to cross a discrete terrain obstacle quickly and efficiently. Our speed in traversing a one-meter long course with a single barrier or gap is approximately twice as fast when this lunge is employed than when using only a classic crawl gait (i.e., approx. 10 cm/s vs 5 cm/s).

landing the front elbows of LittleDog on the top of the barrier. Note the displacement in the position of the rear legs between frame 3 (first row, third column) and frame 6 (second row, third column). This displacement occurs as the rear legs are unloaded during the end of the lunge and the friction force is not large enough to prevent slip. This is an advantage in practice, as it allows for greater forward travel during the lunging maneuver, but the sliding does introduce additional uncertainty in the final, resting pose of the robot once it has come to rest. In actual trials, we simply realign the robot's position after dramatic lunges to account for this, and pacing steps aid in doing such repositioning with impressive speed and accuracy.

3.7.6 Pacing motions

Pacing motions involve moving the left and right pairs of legs sequentially. We use this class of motion 1) to locomote forward, backward, left, or right, 2) to turn in place, 3) to quickly reposition all four feet – for example, in getting ready for the dynamic lunge maneuver described above – in about half the time required by static-crawl motions, and 4) to combine various of the motion types above. It's usefulness comes more from its extreme versatility than from its speed. Our top speed in forward locomotion with pacing is signifi-




Figure 3-18: Climbing steps through sequential dynamic lunging
Our overall speed is around 5.2 cm/sec in climbing these steps with the dynamic lunge.




Figure 3-19: LittleDog employing a dynamic lunge to clear a tall obstacle. This terrain showcases the utility of the dynamic lunge. Terrain includes a gap followed by a brick-shaped obstacle which is at the Phase 3 metric height of 10.8 cm. At this height, straddling the obstacle would make “walking” very difficult, as the legs must be almost fully extended to clear the belly of the robot from the terrain. With the added inclusion of the gap, simply walking across this terrain becomes potentially infeasible.

cantly slower than for the dynamic double-support “trot-walk”⁹ described in Section 3.7.4, however, so we use this motion type primarily for turning and fast repositioning of feet.

To demonstrate the repeatability of pacing motions, we choreographed a 10-second “dance” of motions. The motions sequentially do the following: (a) narrow the stance width, (b) turn approximately 90° in three rapid motions, (c) increase the stance width, and (d) step sideways. At the end of this pacing sequence, the robot performs a complete forward lunge, pushing with the rear legs to lift the front legs, and it also narrows its stance again during this motion. Figure 3-20 shows the resulting footholds during a particular trial (top) with images of the actual robot at the end and start positions immediately below. The roll angles during pacing motions shown in this figure are quantitatively more repeatable than the pitch angles achieved during the more dramatic lunges for which data is given in Figure 3-16, as one might expect since the displacements in COM with respect to the stance legs remain much smaller for pacing. Data for sequential test trials of this series of pacing motions are shown in Figure 3-20. An overhead view of the four-leg stances achieved during one particular trial is shown at top of this figure. The left, front foot in each stance is labeled with a letter corresponding to the end of a labeled region of roll and pitch data in one of the two, lower plots. An arrow in front of this foot indicates the current “forward” direction of the dog. Data show good repeatability. Close-ups of foothold selection accuracy during pacing on flat terrain are shown in Figure 3-21.

Pacing motions are inherently “easier” to execute than either the diagonal or bounding style double-support motions. This is somewhat nonintuitive but can be explained for two reasons. First, both lunging and pacing involve controlled pitching which ensures the body will “land” after rotation onto both previously non-support feet. In contrast, the pitching of the body is maintained near a delicate balance during the diagonal gait. In practice, the body will actually begin to tip during the diagonal gait, and the motion is inherently sensitive in whether it tips toward the front or back non-support foot. The support feet can also be lifted further from dangerous terrain features during both pacing and lunging. Second, in comparing pacing to lunging, the magnitude of the inertia to be rotated is smaller for pacing, which means it requires smaller magnitude accelerations to manipulate the roll

⁹On flat terrain, pacing is about 8 cm/sec, versus 12.5 cm/sec for a trot-walk.


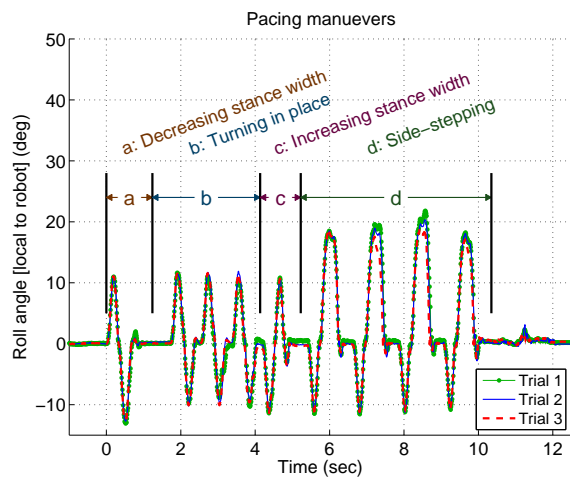

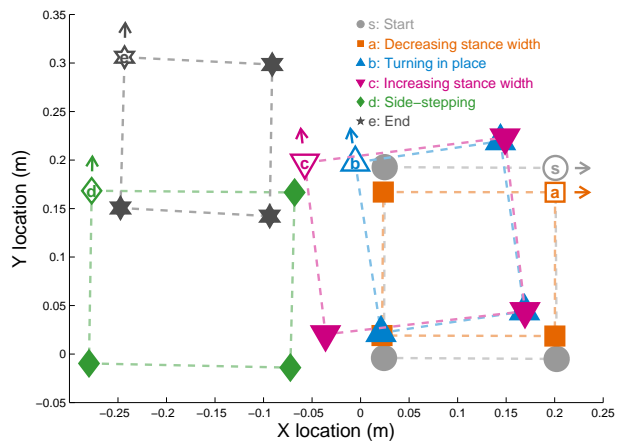


Figure 3-20: Repeatability of long, open-loop playback of various double-support motions

of the body. In summary, the motions during pacing require smaller accelerations and are less prone to destabilization than either the diagonal trot-walk or the dynamic lunge.

Finally, we note that sensitivity to initial conditions (see Figure 3-22) and to external disturbance are both less significant for this high impedance strategy, where joint angles are commanded accurately over time, than for a system where force commands at the feet are instead regulated directly over time. For this simulation, both control strategies (joint angles versus forces) are clocked solely on time and not on the state variables during execution. The stiffness of the robot has a modulating effect on the overall geometry which in turn better regulates the unactuated degrees of freedom in the robot than does a direct force command.



Figure 3-21: Demonstration of accuracy in positioning pacing steps. The rear left foot should align with the black circle marks and the front left foot should align with the red triangles. Spacing between consecutive circles (and triangles) is 4 cm. Pacing steps of up to 10 cm can be achieved to create continuous, forward locomotion.




Figure 3-22: Sensitivity of pitch angle, α , to initial conditions (x_o)
Simulation data comparing expected results when planned force trajectories ($F_x(t)$, $F_z(t)$) are commanded [top] versus when a particular inter-leg angle (θ_1) is position controlled [bottom]. Position control of the joint results in less sensitivity to both initial conditions and external disturbances.

3.7.7 Conclusions

We demonstrate a method for designing highly-repeatable double-support motions with a minimal feedback requirement on the unactuated states of the robots. This has practical use, as real-time measurements of six degrees of freedom (x , y , z , roll, pitch and yaw) of the body of a robot are not as easily obtained as the joint angles. It is also of interest because it allows a robot with high impedance at all actuated joints to obtain the types of gaits more commonly implemented to date in robots with significant passive compliance.

Our method involves 1) calculation of the acceleration in x required to attain a desired

rotational motion, e.g., to achieve some α_{max} (assuming $\dot{z}_m = 0$), 2) a second refined ODE solution using this $\ddot{x}_m(t)$ and kinematic coupling assumptions to relate x , z and α over time, 3) calculation of inverse kinematic solutions to obtain the actuated degrees of freedom dictated by the planned trajectory from 2), and 4) playback of these joint commands to obtain double-support lunging and pacing motions on the real robot. (As mentioned, joint trajectories are commanded using a PD control loop at 400 Hz.) Despite the lack of additional feedback (on forces or global body motion), the commanded joint angles produce repeatable results which are not highly sensitive to the types of variability (initial conditions; backlash loading; battery voltage; wireless packet loss; etc.) experienced in our laboratory.

Obstacles which can be negotiating during lunging include gaps, barriers and steps, and pacing is used as an effective way to reposition all four legs rapidly.

3.8 Future work

Our ultimate goal in this work is to obtain well-controlled motions for a quadruped robot on variable terrain (smooth, rough and intermittent) as quickly and reliably as possible. This chapter presents a methodology for obtaining a family of low-level motions which can then be chosen by a high-level planner, as appropriate for the upcoming terrain. We provide simple, practical control ideas to allow the same, particular robot to transition rapidly from slow, careful walking (on intermittent or other precarious footholds) to fast, dynamic motions. We have successfully obtained such performance on runs up to 1.5 meters in length in a motion capture environment on known terrain boards. Demonstrating this level of flexibility in a field environment is an eventual goal which we believe would be a unique accomplishment in legged robotics.

Below are details on some desired extensions of our approach which are of current research interest in the Robot Locomotion Group at MIT.

3.8.1 Piecewise motion planning on rough terrain

One important challenge in achieving fast, reliable motions on rough terrain is in mating different motions together effectively. Our method for doing this involves a “funneling” approach, similar to the idea described in [23]. Ideal motions are those which can take a large set of initial conditions and funnel them into a small set of end conditions. Using a lunge to negotiate a discrete obstacle (like a step or gap) currently requires beginning within a particular set of initial conditions, so that the front feet successfully get across the obstacle. Then, more carefully planned crawling footsteps can be used to mate together highly dynamic motions, as needed to cross terrain.

3.8.2 Toward dynamic bounding across rough terrain

One of the ultimate goals for dynamic legged robots is to achieve the kind of agility and speed demonstrated by animals such as the leaping mountain goat, pictured in Figure 1-1 (page 18). Obtaining reliable motions such as this of course depends highly on the actuation and sensing capabilities of the robot, which largely determine achievable control. LittleDog, for example, lacks natural compliance to strategically store energy and has important actuator limitations (most particularly, in angular velocity at the hip). However, the double-support lunge, described in Section 3.7.5, provides significantly faster locomotion (by about a factor of two) over terrain with a single discrete obstacle, such as a gap or a Jersey barrier, as depicted in Figure 3-17, than does a ZMP-based walk. Below, we describe several ways to extend usage of the dynamic lunge on a broader class of rough terrain.

Figure 3-23 shows one way in which the dynamic lunge can facilitate locomotion across extreme rough terrain. Here, the lunge is used to get both front feet onto an elevated “step” in terrain quickly. The rear legs can then be brought up more carefully in a climbing style of motion. The initial jump here is identical to the jump used to climb the steps shown in Figure 3-18, since it also begins on level terrain. In this usage, the dynamic lunge improves both speed and repeatability of traversal on such terrain.

There are two other additional ways in which we hope to modify the dynamic lunge to extend the dynamic capabilities of this stiff robot on more general types of rough terrain.




Figure 3-23: LittleDog climbs onto rough terrain

First, we can modify the jump to work for a variety of initial conditions, so that a jump can be initiated from rough terrain. Second, we can add an additional rocking motion to retract the legs, to create a continuous “bounding” gait. Each of these two extensions of the basic dynamic lunge is discussed below briefly. Note that conceptually, both ideas can also be combined, to create continuous episodes of highly dynamic locomotion across rough and uneven terrain.

So far in this Section, we have analyzed the dynamic lunge from a particular initial condition, where all four feet begin on flat, even ground. This is adequate to clear a “clean” obstacle such as the Jersey barrier or gap shown in Figure 3-17. However, we can also extend the range of initial conditions to lunge when the feet are initially on rough terrain. Figures 3-24 and 3-25 illustrate such jumps. In Figure 3-24, the feet rest on shallow steps, so that the rear feet initially differ in height by 4 cm from one another and the front feet are both at a third height, while in Figure 3-25, the robot begins on more general “rough terrain”. Theoretically, if the center of pressure remains at the mid-point between the two rear feet during the underactuated lunge, we can still plan the kinematics of the dog to “mimic” the planar-model motion observed for a lunge from a symmetric initial condition.

We observe good repeatability in initial tests for this generalization of the dynamic lunge. Figure 3-26 compares the pitch of the body during three dynamic lunges: one

initiated from completely flat terrain and the others initiated on the terrain shown in Figures 3-24 and 3-25. The overall shape of the pitch trajectory and its peak value are all very similar in all three cases, as hoped. Additional work should still be done to estimate the range of initial conditions from which a repeatable lunge can be performed, to plan kinematics more precisely (to account for significant spring-compression during loading), and to plan for the front feet to clear terrain successfully and (ideally) to land in an orientation which matches the upcoming terrain profile and with some velocity-matching (to minimize impact forces).




Figure 3-24: Dynamic lunge executed with feet at varying initial heights. Various viewpoints are given, showing orientation of the robot before (top, left and right), during (bottom-left) and after (bottom-right) a lunge. The pitch trajectory for this lunge is shown in middle of Figure 3-26

A second, natural extension of the dynamic lunging motion is the creation of a continuous bounding gait. From initial testing, we predict such motion should result in locomotion speeds at least as fast as the trot-walk described in Section 3.6 (12.5 cm/s) on flat ground. The potential advantages of bounding come on rough terrain, however, where the natural



Figure 3-25: Dynamic lunge initiated on rough rocks

Initial joint angles are automatically selected such that the body is level, with no pitch or roll initially. Two particular initial conditions are depicted here. Top two images show different viewpoints for one test jump. Lower two images show the robot before and during another jump. Resulting motion still follows the expected planar dynamic model, where body position is (as before) measured with respect to the mid-point between the rear legs to model planar motions. Pitch data for this lunge are shown at right in Figure 3-26.




Figure 3-26: Pitch data for dynamic lunge on rough terrain

Data above show the pitch angle of the robot during a dynamic lunge from three different terrain types. Left: The standard lunge, executed from flat terrain with all feet at the same initial height. Middle: Robot begins with rear legs on flat steps at different heights than one another (and different from the front feet), as depicted in Figure 3-24. Right: Robot initiates lunge from rocky terrain, where each foot is at a different height and on a different slope at the initial condition, as shown in Figure 3-25.

leg clearance provided by bounding may allow for much faster traversal than the fast walk over rugged terrain.

Robust, continuous bounding is a significant challenge and will likely require a different control architecture, with faster replanning capabilities within our low-level C code. These challenges will be discussed more in Chapter 6, which also provides initial data for open-loop bounding.

3.9 Summary and conclusions

In this chapter, we present a reasoned approach for planning highly-repeatable dynamic motions for a quadruped robot which lacks both passive compliance and sensory feedback bandwidth – two typical approaches for achieving dynamic gaits. Instead of using such natural or control-based compliant dynamics, we reason about a simplified physical model of the robot body to design planned ground forces and body motions which are compatible for underactuated, double-support phases of motion. If saturation limits (primarily in velocity) of the robot are carefully avoided, the high impedance of the robot ensures that planned joint trajectories are executed with high fidelity, and our planned underactuated

trajectories demonstrate good repeatability over each short (0.3 sec) double-support phase. Interspersing double-support phases with three- or four-legged support periodically corrects for any small deviations (e.g., in pitch) to keep the overall trajectory of the robot near a higher-level set of pre-planned, nominal gait poses. This straight-forward strategy allows us to negotiate significant examples of rough terrain (gaps, hurdles, etc.) at over twice the speed of a traditional “crawl” gait. Equally significant and perhaps less intuitively, we obtain *more reliable* results by using this careful reasoning about the ground reaction forces and body motions than in using a crawl gait. Although, a crawl gait is often assumed to be a “conservative” strategy on rough terrain, violations of the static assumption and swing-leg collisions with terrain become more probable as speed increases – both resulting in unexpected moments which may topple the robot. Our approach provides a practical solution to the dual goals of increasing both speed *and* reliability of locomotion while also enabling efficient negotiation of significant terrain obstacles.

Chapter 4

Metastable Legged Locomotion

In Chapter 2, we argued that traditional stability margins fail to quantify performance accurately and proposed (in Section 2.2) that the expected time-to-failure, or *mean first-passage time* (MFPT) provides an appropriate metric for quantifying the stochastic stability of walking. The concept of walking as a metastable process was then introduced in Section 2.3 (on page 49). In the present chapter, we discuss the quantification of performance for such metastable systems in more detail and demonstrate the use of statistical, mathematical tools to do so. Some of the results presented here can also be found in [27].

4.1 Approach

Simplified models of limit-cycle walking on flat terrain have provided important insights into the nature of legged locomotion. Real walking robots (and humans), however, do not exhibit true limit cycle dynamics because terrain, even in a carefully designed laboratory setting, is inevitably non-flat. Walking systems on stochastically rough terrain may not satisfy strict conditions for limit-cycle stability but can still demonstrate impressively long-living periods of continuous walking. Here, we examine the dynamics of passive rimless-wheel and compass-gait walking on randomly generated rough terrain and employ tools for analyzing stochastic processes to describe the ‘stochastic stability’ of these gaits. This analysis generalizes our understanding of walking stability and provides statistical tools which can be adapted for experimental limit cycle analyses on real walking systems.

4.2 Metastable walking

Systems which are guaranteed to eventually fail cannot be classified as stable. Traditional stability margins and guarantees are no longer applicable in this regime, and we seek better descriptors. Many fields have already approached this problem by embracing the concept of metastability, which literally means “beyond stability”.

The science of legged locomotion is plagued with complexity. Many of the fundamental results for legged robots have come from detailed analytical and computational investigations of simplified models (e.g., [32, 50, 58, 97, 114]). These analyses reveal the limit cycle nature of ideal walking systems and employ Poincaré map analysis to assess the stability of these limit cycles. However, the very simplifications which have made these models tractable for analysis can limit their utility.

Experimental analyses of real machines based on these simple models [171] have revealed that real machines differ from these idealized dynamics in a number of important ways. Certainly, the dynamics of impact and contact with the ground are more subtle than what is captured by the idealized models. But perhaps more fundamental is the inevitable stochasticity in the real system. More than just measurement noise, robots that walk are inherently prone to the stochastic influences of their environment by interacting with slightly different terrain on each footstep. Even in a carefully designed laboratory setting, and especially for passive and minimally-actuated walking machines, the effects of this stochasticity can have a major impact on the system dynamics. In practice, it is very difficult (and technically incorrect) to apply deterministic limit cycle stability analyses from the simple systems to our experimental walking machines - the real machines do not have true limit cycle dynamics.

In this chapter, we extend the analysis of simplified walking models toward real machines by adding an element of stochasticity into the walking model. Although we have considered a number of sources of uncertainty, we will focus our presentation on a particularly compact model - where the geometry of the ground is drawn from a random distribution. Even with very mild deviations in the ground from a nominal slope angle, the resulting trajectories of the machine are different on every step and end with the robot even-

tually falling down (with probability one as $t \rightarrow \infty$). As the deviations from the nominal slope increase, the analysis also suggests a way to model and even control a walking robot on moderately rough, unmodeled terrain.

4.2.1 Suggested reading on metastability

Here, we briefly point the reader to some references on metastability which provide additional background. These include a small, representative set of the literature on metastable system dynamics [60, 79, 127, 52], tools for estimating stochastic stability in discrete cases (Markov chains) [169, 16, 15, 17, 84, 105, 181], and issues of model order reduction [76, 177, 11]. Additionally, two recommended texts on stochastic processes are [51, 90].

4.2.2 Metastable limit cycle analysis

The dynamics of walking systems are continuous, but they are punctuated by discrete impact events, such as when a foot comes into contact with the ground. These impacts provide a natural time-discretization of a gait onto a Poincaré map. Therefore, we will consider walking systems governed by the discrete, closed-loop return-map dynamics:

$$\mathbf{x}[n + 1] = \mathbf{f}(\mathbf{x}[n], \gamma[n]), \quad (4.1)$$

where $\mathbf{x}[n]$ denotes the state of the robot at step n and $\gamma[n]$ represents the slope of the ground, which is a random variable drawn independently from a distribution P_γ at each n . This model for stochastically rough terrain dramatically simplifies our presentation here, but it does require us to restrict our analysis to strictly forward walking¹. These state evolution equations represent a discrete-time, continuous-state Markov process (or infinite Markov chain). For computational purposes, we will also discretize the state space into a finite set of states, x_i . Defining the state distribution vector, $\mathbf{p}[n]$, as

$$p_i[n] = \Pr(\mathbf{X}[n] = x_i), \quad (4.2)$$

¹Including backward steps is straightforward, but it requires the model to include spatio-temporal correlations in the slope angle.

we can describe the state distribution (master) equation in the matrix-form

$$\mathbf{p}[n + 1] = \mathbf{p}[n]\mathbf{T}, \quad T_{ij} = \Pr(\mathbf{X}[n + 1] = x_j \mid \mathbf{X}[n] = x_i). \quad (4.3)$$

\mathbf{T} is the state-transition matrix; it is a stochastic matrix (each row must sum to one). The n -step dynamics are revealed by the Chapman-Kolmogorov equation,

$$\mathbf{p}[n] = \mathbf{p}[0]\mathbf{T}^n.$$

We obtain the transition matrix numerically by integrating the governing differential equation forward from each mesh point and then using barycentric interpolation [128] to assign individual transition probabilities.

For walking, we will designate one special state, x_1 , as an absorbing state representing all configurations in which the robot has fallen down. Transitions to this state can come from many regions of the state space; there are no transitions away from this state. Assuming that it is possible to get to this absorbing state (possibly in multiple steps) from any state, then this absorbing Markov chain will have a unique stationary distribution, where the entire probability mass is in the absorbing failure state.

Figure 4-1 shows a toy example of a Markov chain. For the legged systems we will analyze in this chapter, the states in a Markov chain will include a set of discretized points in continuous state space plus one additional, “fallen” state. The fallen state is modeled as an absorbing failure state, which must transition back to itself with probability 1. An absorbing state will correspondingly have a one in the diagonal at its row (and column) of the transition matrix. In the remainder of this chapter, we will assume this fallen state is the first state in the transition matrix: i.e., $T(1, 1) = 1$.

The dynamics of convergence to the absorbing state can be investigated using an eigenvalue analysis. Without loss of generality, let us order the eigenvalues, λ_i , in order of decreasing magnitude, and label the corresponding (left) eigenvectors, \mathbf{v}_i , and the characteristic time constants, $\tau_i = \frac{-1}{\log(\lambda_i)}$. The transition matrix for an absorbing Markov chain will have $\lambda_1 = 1$, with \mathbf{v}_1 representing the stationary distribution of the absorbing state. The magnitude of the remaining eigenvalues ($0 \leq |\lambda_i| < 1, \forall i > 1$) describe the transient




Figure 4-1: A toy example of a Markov chain and corresponding transition matrix. At each, labeled state, arrows leading away are labeled with the probability that the dynamics will result in a transition to the state to which the arrow points, conditioned on starting from that particular state. The transition matrix at right contains the identical information. Each row must sum to unity, because all possible transitions are enumerated in the Markov chain and transition matrix.

dynamics and convergence rate (or mixing time) to this stationary distribution. Our analyses of the walking models we investigate here will reveal a general phenomenon: λ_2 is very close to 1, and $\tau_2 \gg \tau_3$. This is an example of metastability: initial conditions in eigenmodes 3 and higher are forgotten quickly, and \mathbf{v}_2 describes the long-living (metastable) neighborhood of the dynamics. In metastable systems, it becomes useful to define the *metastable distribution*, ϕ , as the stationary distribution conditioned on having not entered the absorbing state:

$$\phi_i = \lim_{n \rightarrow \infty} \Pr(\mathbf{X}[n] = x_i \mid \mathbf{X}[n] \neq x_1).$$

This is easily computed by zeroing the first element of \mathbf{v}_2 and normalizing the resulting vector to sum to one.

Individual trajectories in the vicinity of a metastable attractor are characterized by random fluctuations around the attractor, with occasional “escape attempts”, in which the system has entered a region of relatively low influence from the attractor. For walking systems this is equivalent to noisy, random fluctuations around the nominal limit cycle, with occasional transitions to the absorbing (fallen) state. The existence of these escape attempts suggests a natural quantification of the relative stability of metastable attractors in terms of first-passage times. The *mean* first-passage time (MFPT) to the fallen absorbing state is the

time we expect our robot to walk on average before falling down.

Let us define the mean first-passage time vector, \mathbf{m} , where m_i is the expected time to transition from the state x_i into the absorbing state. Fortunately, the mean first-passage time is particularly easy to compute, as it obeys the relation:

$$m_i = \begin{cases} 0 & i = 1 \\ 1 + \sum_{j>1} T_{ij} m_j & \text{otherwise} \end{cases}$$

That is, we assume at least one successful step is (by definition) taken from every non-fallen state, and correspondingly, the expected first-passage time must be one step more than the expected value of the first-passage time after this single step is taken. In matrix form, this yields the following one-shot calculation of the vector of MFPT's for all non-fallen states:

$$\mathbf{m} = \begin{bmatrix} 0 \\ (\mathbf{I} - \hat{\mathbf{T}})^{-1} \mathbf{1} \end{bmatrix}, \quad (4.4)$$

where $\hat{\mathbf{T}}$ is \mathbf{T} with the first row and first column removed. \mathbf{m} quantifies the relative stability of each point in state space. One interesting characteristic of metastable systems is that the mean first-passage time around an attractor tends to be very flat; most system trajectories rapidly converge to the same metastable distribution (forgetting initial conditions) before escaping to the absorbing state. Therefore, it is also meaningful to define a system-wide mean first-passage time, M , by computing the expected first-passage time over the entire metastable distribution,

$$M = \sum_i m_i \phi_i. \quad (4.5)$$

When $\tau_2 \gg \tau_3$, we have $M \approx \tau_2$, and when $\lambda_2 \approx 1$, we find that

$$M \approx \tau_2 = \frac{-1}{\log(\lambda_2)} \approx \frac{1}{1 - \lambda_2}.$$

4.3 Numerical modeling results

This section uses two simple, classic walking models to demonstrate use of the methodology presented in Section 4.2.2 and to illustrate some of the important characteristics typical for metastable walking systems more generally. The two systems presented here are the rimless wheel and the passive compass gait walker, each of which is illustrated in Figure 4-2.




Figure 4-2: The rimless wheel (RW) (left) and compass gait (CG) walker (right) models

4.3.1 Rimless wheel

The rimless wheel (RW) model consists of a set of N massless, equally-spaced spokes about a point mass. Potential energy is gained as it rolls downhill, while conservation of angular momentum results in a loss of kinetic energy at each impact. For the right combination of constant slope and initial conditions, a particular RW will converge to a steady limit cycle behavior, rolling forever and approaching a particular velocity at any (Poincaré) “snapshot” in its motion (e.g., when the mass is vertically above a leg and $\theta = 0$). The motions of the rimless wheel on a constant slope have been studied in depth [32, 171].

In this section, we will examine the dynamics of the RW when the slope varies stochastically at each new impact. To do this, we discretize the continuous set of velocities, using a set of 250 values of ω , from 0.01 to 2.5 (rad/s). We also include an absorbing failure state,

which is defined here to include all cases where the wheel did not have sufficient velocity to roll past its apex on a particular step. Our wheel model has $N = 8$ spokes ($\alpha = \frac{\pi}{4}$). At each ground collision, we assume that slope between ground contact points of the previous and new stance leg is drawn from an approximately² Gaussian distribution with a mean of $\bar{\gamma} = 8^\circ$. For clarity, we will study only wheels which begin at $\theta = 0$ with some initial, downhill velocity, ω_o , and we consider a wheel to have failed on a particular step if it does not reach an apex in travel, $\theta = 0$ with $\omega > 0$. (ω is defined as positive going downhill, i.e., clockwise in Fig. 4-2.) Note that the dynamic evolution of angular velocity over time will not depend on the choice of a particular magnitude of the point mass, and we will use spokes of unit length, $l = 1$ meter, throughout.

On a constant slope of $\gamma = 8^\circ$, any wheel which starts with $\omega_o > 0$ has a deterministic evolution over time and is guaranteed to converge to a fixed point of $\omega = 1.2097$ (rad/s). The return map defining the step-to-step transitions from ω_n to ω_{n+1} is given as:

$$\omega_{n+1} = \sqrt{\cos^2 \alpha \left(\omega_n^2 + 2 \left| \frac{g}{L} \right| (1 - \cos \beta_1) \right) - 2 \left| \frac{g}{L} \right| (1 - \cos \beta_2)}$$

where $\beta_1 = \frac{\alpha}{2} + \gamma$ and $\beta_2 = \frac{\alpha}{2} - \gamma$, with $\gamma > 0$ as the downhill slope. A plot of this return function is shown in Figure 4-3.

When the slope between successive ground contacts is drawn from a stochastic distribution, the function given in Figure 4-3 is now replaced by a probabilistic description of the transitions, as illustrated in Figure 4-4. Given the current state is some particular ω_n , there is a corresponding probability density function (PDF) to describe what the next state, ω_{n+1} , will be. Figure 4-5 shows this set of PDF's clearly; this 3D plot illustrates the same probabilistic return map shown from overhead in Figure 4-4. For our discretized model, each height value in Figure 4-5 is in fact the magnitude of the value of the transition matrix, T_{ij} , where i is the state we are coming from (ω_n , on the x-axis) and j is the state we are going to (ω_{n+1} , on the y-axis). Figures 4-4 and 4-5 provide a graphical representation of the transition matrix describing this metastable dynamic system.

To generate the discrete transition matrix, we calculate $\omega_{n+1} = \omega_{n+1}(\omega_n, \gamma)$ for each of

²To avoid simulating pathological cases, the distribution is always truncated to remain within $\pm 10^\circ$, or roughly 6σ , of the mean.

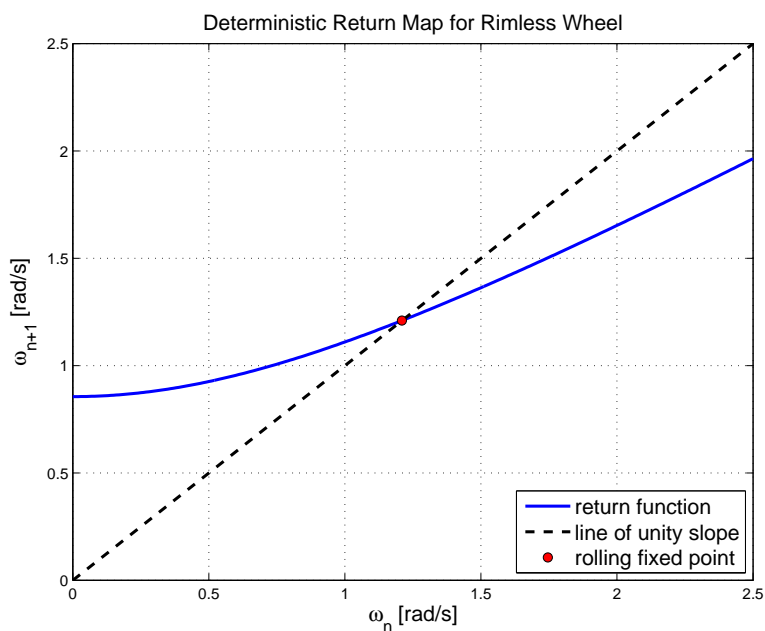


Figure 4-3: Return map and fixed point for an 8-spoke rimless wheel on flat, downhill terrain with a constant slope of 8 degrees

The velocity, ω , is the speed of rotation when a spoke is exactly vertical (i.e., when the potential energy between two impacting “steps” is greatest and the velocity is therefore lowest). The wheel correspondingly fails to complete a step whenever $\omega < 0$.




Figure 4-4: Return distribution and metastable “neighborhood” for an 8-spoke rimless wheel on downhill terrain with a mean step-to-step slope of 8 degrees and standard deviation of 1.5 degrees

There is now a *probability density function* describing the transition from ω_n to ω_{n+1} .




Figure 4-5: 3D view of the return distribution for the rimless wheel system
This surface plot is a smoothed rendering of the step-to-step transition matrix, T , with the probability density functions for some particular states (ω_n) overlaid for greater clarity. For the deterministic rimless wheel, the corresponding picture would show a narrow line of infinite height, representing the deterministic state-to-state return function.

a discrete set of 601 possible γ values, in the range of ± 10 degrees from the mean. Each new state is then represented in the mesh using barycentric weighting interpolation [128], which (we note) inherently adds a small level of additional (unintended) noise to the modeled dynamics. In Figures 4-4 and 4-5, the noise in terrain slope has a standard deviation of $\sigma = 1.5^\circ$. Using MATLAB to take the 3 largest eigenvalues of the transpose of the transition matrix for this case, we find that the largest eigenvalue, λ_2 , is within 10^{-14} of $\lambda_1 = 1$, which is within the mathematical accuracy expected. This eigenvalue corresponds to the absorbing failure state, and the corresponding eigenvector sums to 1, with all values except the failure state having essentially zero weight³ in this vector (since all wheels will eventually be at this state, as $t \rightarrow \infty$). All other eigenvectors sum to zero (within numerical limits⁴), since they must die away as $t \rightarrow \infty$. The second-largest eigenvalue is $\lambda_2 = 0.999998446$. Using the methods presented in Section 4.2.2, this corresponds to a system-wide MFPT of about $1/0.000001554 = 643,600$ steps. Each initial condition has a particular MFPT, $m(\omega)$, which is obtained from Eq. 4.4 and plotted in Figure 4-6.

³All states except the failure state had a magnitude less than 10^{-10} , numerically.

⁴Our numerical limits are determined by the resolution in MATLAB, which is on the order of $1e-15$.

Note that the value of the mean first-passage time is nearly flat throughout a large portion of state space. This is characteristic for metastable systems, which justifies the notion of a “system-wide” MFPT, $M \approx 1/(1 - \lambda_2)$, quantifying the overall stochastic stability of a particular dynamic system. For this particular case, there are no regions in state space (except the failure state) with MFPT significantly lower than the system-wide value, which is *not* typical more generally; the passive compass gait walker discussed next in Section 4.3.2 is highly sensitive to initial conditions, for instance, although it too has regions of state space which share a nearly uniform MFPT value.




Figure 4-6: Mean first-passage time as a function of the initial condition, ω_o . Data are for a rimless wheel on stochastic terrain with mean slope of 8 deg and σ of 1.5 deg. Points show the approximation obtained through eigenanalysis of the discretized system, and a smoothed line is overlaid. Note that MFPT is largely constant over a large portion of state space.

The eigenvector associated with λ_2 yields the PDF of the metastable dynamic process – the relative probability of being at each location in state space, given initial conditions have been forgotten and the walker has not yet failed. Figure 4-7 shows the resulting probability distribution functions for the rimless wheel for each of several levels of noise. Pictorially, each system-wide PDF for a metastable system is analogous to the fixed point for a stable, deterministic system. In the deterministic case, the probability of being exactly at the fixed point approach unity as $t \rightarrow \infty$.

The third-largest eigenvalue of the transition matrix, λ_3 , quantifies the characteristic time scale in which initial conditions are forgotten, as the dynamics evolve toward the metastable distribution (or toward failure). For the case presented here ($\sigma = 1.5^\circ$), $\lambda_3 \approx 0.50009$, which means almost half of the PDF of the initial condition composed of this eigenvector is lost (“forgotten”) with each, successive step; an even larger fraction-per-step

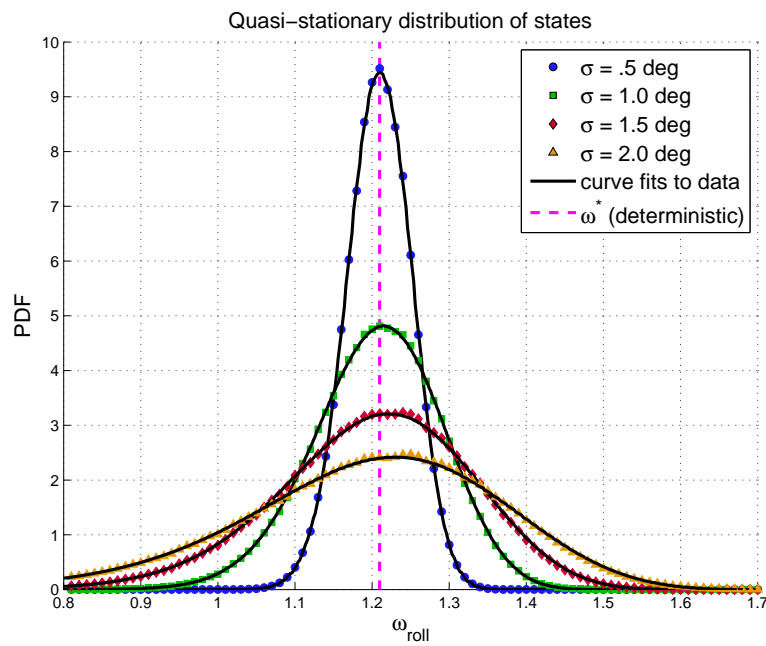


Figure 4-7: Quasi-stationary probability density functions for the stochastic rimless wheel for each of several values of terrain noise, σ

Each distribution is estimated by renormalizing the eigenvector associated with the second-largest eigenvalue of the transpose of the transition matrix. Note that meshing inherently adds noise to the dynamic system; smoothed lines are drawn on top of the raw data (shown as points) from the scaled eigenvectors.

is lost for all remaining eigenvectors (with even smaller values of λ). Within a few steps, initial conditions for any wheel beginning in our range of analysis $0 < \omega_o \leq 2.5$ [rad/s] have therefore predominantly evolved into the metastable PDF (or have failed). If we multiply the metastable PDF, $\phi(\omega)$, by the transition matrix, we obtain the conditional probability, $Pr(\omega_{n+1}|\omega_n)$ of having just transitioned from ω_n to ω_{n+1} , which is shown both as a 3D plot in Figure 4-8 and as a set of contour lines overlaid in Figure 4-4.




Figure 4-8: 3D view of the metastable “neighborhood” of state-to-state transitions
 If a rimless wheel starts from some arbitrary initial condition and has not fallen after several steps, this contour map represents the joint probability density function of being in state ω_n now and transitioning to ω_{n+1} . The contour lines drawn are identical to those overlaid in Figure 4-4. They delineate neighborhoods of likely (ω_n, ω_{n+1}) pairings, analogous to the unique fixed point of the deterministic case.

This particular system has a beautiful simplicity which allows us to extract some additional insight from the conditional probability in Figure 4-8. Because of the definition of ω_n as being the velocity when the mass is at its apex in a given step, the value of $\omega_{n+1} = 0$ represents the boundary to the absorbing failure state in this example. If we visualize the contours of the conditional probability as they extend toward $\omega_{n+1} = 0$ in Figure 4-4, we see that most failures do not occur because we transition from a very slow state (ω_n close to zero) to failure but are more typically due to sudden transitions from more dominant states in the metastable distribution to failure.

Finally, when this methodology is used to analyze the rimless wheel for each of a variety of noise levels (σ), the dependence of system-wide MFPT on σ falls as shown in Figure 4-9.

For very low levels of noise, MATLAB does not find a meaningful solution⁵. As the level of noise increases, the MFPT decreases smoothly but precipitously. (Note that the y-axis is plotted on a logarithmic scale.) The stochastic stability of each particular system can be quantified and compared by calculating this estimate of MFPT which comes from λ_2 of the transition matrix.




Figure 4-9: Mean first-passage time (MFPT) for the rimless wheel, as a function of terrain variation, σ

Estimates above 10^{14} correspond to eigenvalues on the order of “ $1 - 10^{-14}$ ” and are beyond the calculation capabilities of MATLAB.

4.3.2 Passive compass gait walker

The second metastable dynamic system we analyze in this paper is a passive compass gait (CG) walker. This system consists of two, massless legs with concentrated masses at the intersection of the legs (“the hip”) and partway along each leg, and it has been studied in detail by several authors, e.g., [50, 58, 162]. Referring to Figure 4-2, the parameters used for our metastable passive walker are $m = 5$, $m_h = 1.5$, $a = .7$, and $b = .3$. Given an appropriate combination of initial conditions, physical parameters and constant terrain slope, this ideal model will walk downhill forever.

⁵The numerical resolution on MATLAB cannot capture MFPT values above $1e14$ or so.

When each step-to-step terrain slope is instead selected from a stochastic distribution (near-Gaussian, as in Section 4.3.1), evolution of the dynamics becomes stochastic, too, and we can analyze the stochastic stability by creating a step-to-step transition matrix, as described in detail for the rimless wheel. The resulting system-wide MFPT as a function of terrain noise, $M(\sigma)$, is shown in Figure 4-10. Note that it is similar in shape to the dependence shown in Figure 4-9.




Figure 4-10: Mean first-passage time for a passive compass gait, as a function of terrain variation

Results for analysis of a compass gait walker using a discretized (meshed) approximation of the transitions. Average slope is 4 degrees, with the standard deviation in slope shown on the x-axis. Noise is a truncated Gaussian distribution, limited to between 0 and 8 degrees for all cases.

To analyze this system, our discretized mesh is defined using the state immediately after each leg-ground collision. The state is defined completely by the two leg angles and their velocities. On a constant slope, these four states are reduced to three states, because a particular combination of slope and interleg angle will exactly define the orientation of both the stance and swing leg during impact. We use three states to define our mesh, although we allow for variations in slope from the nominal (mean) value. To approximate the transition matrix, we simulate the deterministic dynamics a short distance forward or backward in time to find a state where the slope of the line connecting the “feet” of the legs is equiv-

alent to our desired, nominal slope. Because the dynamics between collisions are entirely deterministic, these two states are mathematically equivalent for the stochastic analysis. If such a state does not exist for a particular collision (which occurs only very rarely), we treat this as a member of the absorbing failure state. The advantage of this approximation is that reduction of the dimensionality from 4 states to 3 provides a significant improvement in accuracy. Specifically, it allows us to mesh finely enough to capture near-infinite MFPT for low-noise systems, while using four states did not. Instead of using either the absolute stance leg angle, X_1 , or the relative swing leg angle, X_2 , we mesh using the inter-leg angle, $\alpha = \pi - X_2$. The three states we use are then: (1) absolute angular velocity of the stance leg, X_3 , (2) relative velocity of the swing leg, X_4 , and (3) the inter-leg angle, α .

Figure 4-11 shows a slice of the deterministic basin of attraction for this compass gait on a constant slope (top), along with regions in state space with nearly-constant MFPT (bottom two) for two different magnitude of noise (σ) in terrain. Each slice is taken at the same inter-leg angle, $\alpha \approx 25.2^\circ$. In the deterministic case, the basin of attraction defines the set of all states with infinite first-passage time: all walkers beginning with an initial condition in this set will converge toward the fixed point with probability 1. For stochastic systems which result in metastable dynamics, there is an analogous region which defines initial conditions having MFPT very close to the system-wide value, M . Interestingly, the deterministic and stochastic basin shapes are quite similar here; we expect this may often be the case for systems such as this with discrete jumps in state space.

The image at the top of Figure 4-12 shows the deterministic basin of attraction for this CG walker more clearly. This plot was generated by sampling carefully over the state space and simulating the dynamics. By contrast, the plot at the top of Figure 4-11 intentionally uses the same mesh discretization used for the stochastic system, to provide a better head-to-head comparison of the change in shape due to the addition of terrain noise (as opposed to the noise of the discretization itself). The second image in Figure 4-12 shows the deterministic basin of attraction for a different set of physical parameters ($m = 0.5m_h$; $a = .6$ m; $b = .4$ m) on the same, constant slope of 4° . This basin looks qualitatively more delicate and the resulting performance of this walker on stochastic terrain is in fact much worse (i.e., MFPT of about 20 steps when $\sigma = 0.5^\circ$, where we find $M = 180,000$ for the

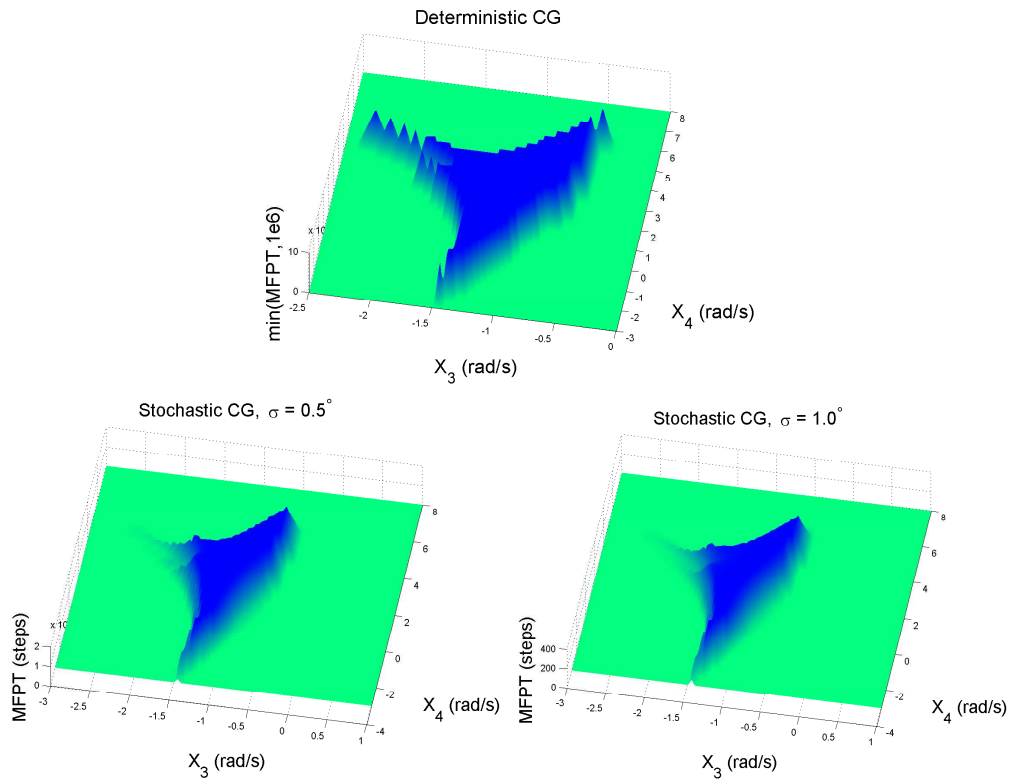


Figure 4-11: Basin of attraction versus a map of MFPT

Basin of attraction (top) for deterministic CG walker and maps of MFPT for CG walkers on stochastic terrain ($\sigma = 0.5^\circ$, bottom left; $\sigma = 1.0^\circ$, bottom right). To aid in visual comparison, all 3 plots use the same meshing discretization. Note that in the stochastic cases, there is an apparent “MFPT basin” which is essentially a low-pass filtered version of the deterministic basin of attraction. The shape of the region with constant MFPT does not change significantly, even when the magnitude of the MFPT itself varies greatly (about 180,000 steps for the left-hand plot; 390 for the one on the right). This region of nearly uniform MFPT represents a boundary in state space where a walker is likely to converge *very* rapidly toward a neighborhood near the fixed point (for a deterministic system) or to the metastable neighborhood (for stochastic systems).

other walker).

Just as in the case of the rimless wheel, the fixed point (for our deterministic compass gait system) is now replaced (in the stochastic case) by a probability density function, defining the likelihood of being in any particular state (conditioned on not having fallen) as $t \rightarrow \infty$. Figure 4-13 shows 2D contour plot sections of the PDF obtained from the eigenanalysis of the stochastic compass gait. The outermost contour defines a boundary containing 0.999 of the probability distribution in state space. The distribution spreads over more of state space as the level of noise increases, in a manner analogous to the widening of the distribution with noise seen in Figure 4-7.

Finally, we note that the relationship in state space between the PDF of the metastable dynamics, shown in Figure 4-13, and the region of nearly-uniform mean first-passage time, M , shown at the bottom of Figure 4-11, hints at where successful “escape attempts” are most likely to occur over time. Figure 4-14 overlays these two regions across a different dimensional slice of the 3D space for the $\sigma = .5^\circ$ and $\sigma = 1.0^\circ$ cases. As the tails of the metastable PDF (shown in yellow) approach the boundary of the uniform-MFPT basin (shown in blue), there is a higher probability of failing on any given step during the metastable process, resulting in turn in a less stochastically stable system (i.e., with a lower system-wide value of M).

4.4 Discussion

This section briefly discusses some additional issues in the use of the stochastic methods presented toward both the design of controllers for and evaluation of the performance of walking systems. We highlight the potential impacts on control design of using this methodology, discuss implementation issues, and also provides a few further observations on the properties of metastable system which result in multiple attractors (e.g., period- n gaits).




Figure 4-12: Basins of attraction (blue regions) and fixed point for two compass gait walkers, each on terrain of constant slope (4 deg)

For the walker at top, $m = \frac{10}{3}m_h$, $a = .7$ m, and $b = .3$ m; at bottom, $m = \frac{1}{2}m_h$, $a = .6$ m and $b = 0.4$ m. Each image shows a slice of the 3D basin, taken at the inter-leg angle of the fixed point for each walker. The fixed point is at $X_3 = -.89$ (rad/s), $X_4 = 2.89$ (rad/s), $\alpha = 25.2^\circ$ for the first walker and at $X_3 = -1.14$ (rad/s), $X_4 = 1.26$ (rad/s), $\alpha = 33.4^\circ$ for the lower one.

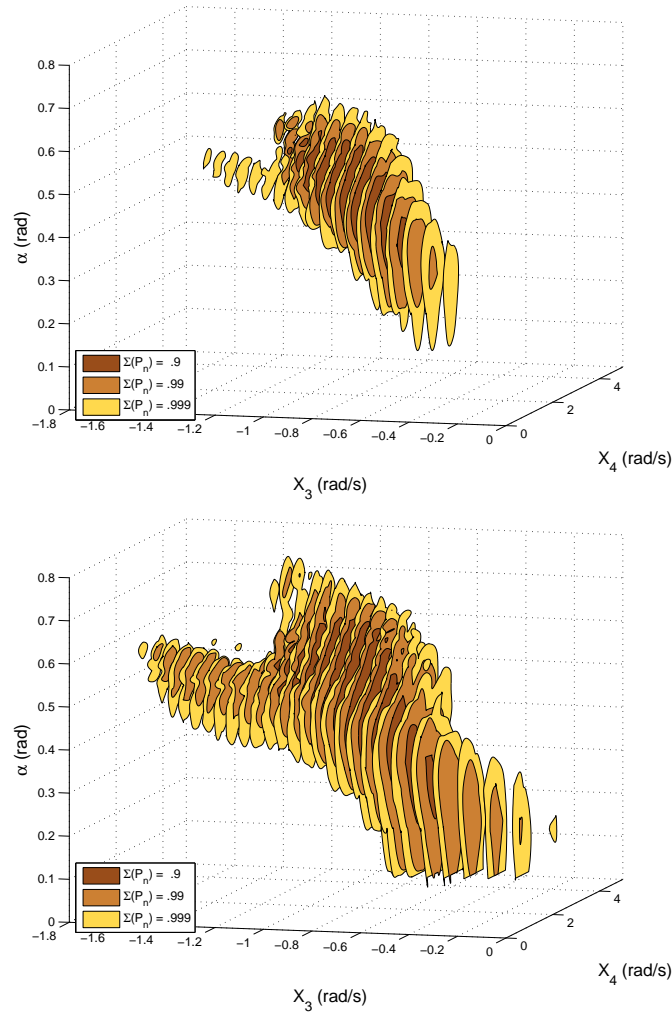


Figure 4-13: Metastable neighborhoods for a passive compass gait walker

On stochastic terrain, there is no fixed point for the compass gait walker. Instead, there are metastable “neighborhoods” of state space which are visited most often. As time goes to infinity, if a walker has not fallen, it will most likely be within this region. The contours shown here are analogous to the PDF magnitude contours in Figure 4-7; they are drawn to enclose regions capturing 90%, 99%, and 99.9% of walkers at any snapshot during metastable walking. Top picture corresponds to σ of 0.5 degrees in terrain slope; σ of 1.0 degrees (bottom) results in larger excursions in state space, as expected.




Figure 4-14: Metastable system: Contours of the stochastic “basin of attraction” Basins are shown where MFPT is $0.5M$, $0.9M$ and $0.99M$ (blue) versus contours where the integral of the PDF accounts for $.9$, $.99$, and $.999$ of the total metastable distribution (yellow). The metastable dynamics tend to keep the system well inside the “yellow” neighborhood. When the tails of this region extend more significantly out of the blue region, the system dynamics are less stochastically stable, with a lower value of M . The axis into the page represents the swing leg relative velocity, X_4 , and a slice is taken at $X_4 = 2.33$ [rad/s]. Terrain variation for the top plot is $\sigma = 0.5$ degrees, which corresponds to $M \approx 180,000$ steps. The system at bottom has twice the magnitude of noise and is less stable; terrain variation of $\sigma = 1.0$ degrees results $M \approx 390$ steps or so.

4.4.1 Impacts on control design

One of the primary goals of a controller is to enhance the dynamic stability of a system. For walking systems, we propose throughout this paper that this should be defined as increasing the *stochastic stability*. We would like time-to-failure to be long, and we would like a system to converge toward the metastable distribution from a large set of initial conditions. The tools provided here can be used in optimizing controllers with either or both of these two aims in mind.

As an example, consider an active compass gait walker, with a torque source at the hip but with the ankles still unactuated at the ground contact. Putting this walker on a *repeating* terrain, as depicted in Figure 4-15, allows us to mesh across a region of the state space of possible poses. By designing a low-level PD controller on inter-leg angle, we can discretize the action space on a single once-per-step policy decision. The optimal high-level policy (to select desired inter-leg angle) for the system can now be solved via value iteration. Our results for such a control methodology allow this underactuated compass gait model to walk continuously over impressively rough terrain, as discussed further in Section 5.5.

4.4.2 Implementation issues

Practical use of the statistic tools outlined in this chapter depend on thoughtful implementation. Issues which deserve attention here include the following: scaling, interpolation, and induction algorithms for estimating state transitions.

Scaling is a major concern for two reasons. First, including more dimensions to define the states of a systems results in exponentially more elements in the resulting transition matrix. Any practical storage and manipulation of the matrix (e.g., using MATLAB) becomes impractical for even moderately large degree-of-freedom systems. In our experiments with the compass gait walker described in Chapter 5, for instance, we use a clever trick to “fast-forward” the dynamics of each continuous step trajectory to a particular plane in state space and correspondingly reduce the number of degrees of freedom required to described each post-impact state from 4 to 3. This reduction in dimensionality may sound modest, but we were never able to successfully refine a mesh for either the passive or active compass

gait models without it. As the compass gait is a fairly simple system, we expect three degrees of freedom is likely to be a practical limit for dynamic models of any reasonable sophistication.

As a result, model order reduction is essential as systems become more complicated. A promising trick is to create a hierarchy in a controller, using a low-level control loop which is driven by a high-level representation of some subset of the total state space. For instance, recall that in Chapter 3 we described some dynamic gaits for a quadruped which has 18 degrees of freedom: 12 actuated joint angles, plus the 6 degrees of freedom describing the location and orientation of the body (which are not directly actuated). For a dynamic lunging maneuver, we collapsed these 18 degrees of freedom into 2 DOF: an actuated hip angle for the stance leg and an underactuated pitch angle of the robot body. Our results showed good agreement between the highly-simplified model and the actual robots. Our message here is that although model order reduction can often be a significant challenge requiring thoughtful planning, we do not believe it is an absolute road block to practical implementation.

Estimating the transition probabilities is potentially challenging for a real robot, as well. The basic, well-established principles of Bayesian inference should provide our starting point [39]. However, if we simply try to obtain our estimates through Monte Carlo style trials, we will often require an impractically large number of experiments to get good estimates of rare events as discussed further in Section 5.4.2. We should expect practical estimation of transition probabilities will also require intuition about the physics of the system to enhance our estimations of total transition probabilities wherever practical. Bayesian estimation is an area of great interest recently in some fields of robotics – most notably in Simultaneous Localization and Mapping (SLAM) [158, 180]. However, the issue of estimating state transition matrices for the *dynamics* of a robot systems is not treated widely in the literature. Bayesian estimation of the state transition functions for robot manipulation tasks is addressed by several authors [54, 112].

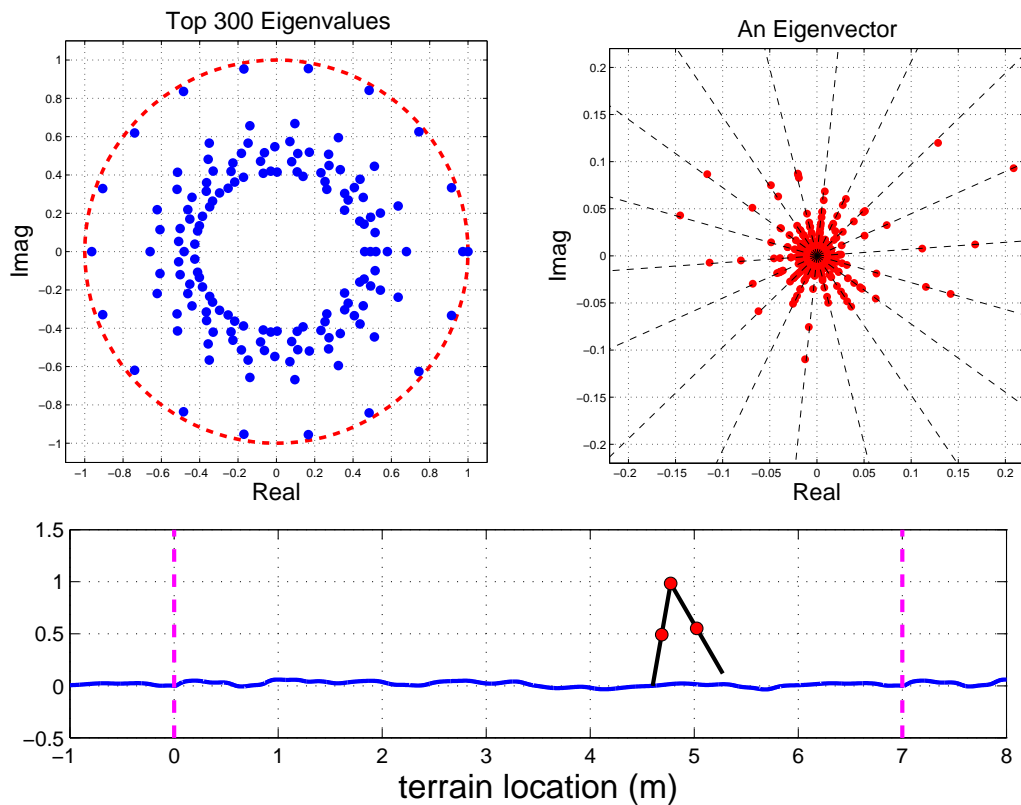


Figure 4-15: Controlled compass gait walker, with torque at the hip
 To solve for an optimal policy using value iteration, the terrain wraps every 7 meters. The optimization maximizes the MFPT from any given state. An eigenanalysis reveals a complex set of eigenvalues (top), spaced evenly about (but strictly inside of) the unit circle. Corresponding eigenvectors are also complex.

4.4.3 Multiple stable limit cycles

Metastable dynamic systems sometimes have an inherent periodicity. We expect this may be the case on a slightly steeper slope, for instance, where compass gait models experience period-doubling bifurcations [58]. Another case where periodicity arises is for wrapping terrain, such as the terrain for the controlled walker in Figure 4-15. Wrapping is a realistic model for many in-laboratory walking robots, as they are often confined to walk on a boom – repeatedly covering the same terrain again and again. A controlled version of the compass gait will be discussed in detail in the next chapter. The important detail to note now is that optimal policies for walking on repeating terrains may also have periodicity, selecting the same particular footholds every n^{th} time around the terrain.

In our simulations for a CG walker on wrapping terrain, we observe that a repeating, n -step cycle results in multiple eigenvalues, λ_2 through λ_{n+1} , all with magnitude just under unity. They are complex eigenvalues, as are the corresponding eigenvectors. In cyclical dynamic systems, this group of eigenvalues together replace the role of the second-largest eigenvalue, λ_2 , in systems which lack periodicity. The top left image in Figure 4-15 shows such a set of eigenvalues, all lying just within the unit circle. The next-smallest set of eigenvalues are all significantly smaller in this example. The complex eigenvalues and eigenvectors mathematically capture an inherent periodicity, in which the probability density function changes over time in a cyclical manner.

4.5 Conclusions

The goal of this chapter has been to motivate the use of stochastic analysis in studying and (ultimately) enhancing the stability of walking systems. Robots that walk are inherently more prone to the stochastic influences of their environment than traditional (e.g., factory) robots. Locomotory systems capable of interacting with the real world must deal with significant uncertainty and must perform well with limited energy budgets and despite limited control authority.

The stochastic dynamics of walking on rough terrain fit nicely into the well-developed study of metastability. The simplified models studied here elucidate the essential picture of

metastable limit-cycle dynamics which make occasional escapes to the fallen-down state. We anticipate that metrics for stochastic stability, such as the mean first-passage time, can be used as potent metrics for quantifying both the relative stability of different states across state-space and the overall system-wide stability of real walking systems.

Chapter 5

Compass Gait Model on Rough Terrain

In this chapter, we investigate the stochastic stability of both passive and active models on the Compass Gait (CG) model on rough terrain. In both regimes, we mesh state space and model the dynamics as a discrete Markov Chain. For metastable walking systems, an eigenanalysis on the (sparse) transition matrix which represents the system dynamics reveals both a characteristic, system-wide mean first-passage time and the “metastable neighborhood” of the regions visited most often in state space.

For the actuated versions of the CG model, we use the value iteration algorithm [166] to solve for approximate optimal control of a CG biped model on rough terrain. The resulting control policy is considered only *approximately* optimal because we do not attempt to solve for a continuous policy over time. Instead, a hierarchical control strategy is used, with a low-level controller regulating the hip angle trajectory, and a higher-level policy selecting a low-level option once per step.

5.1 Introduction

In Chapter 4, we introduced a simple model of walking called the “compass gait” biped. This model provides a very appealing simplification of human walking while still capturing much that is fundamental. Researchers have long noted that experimental data from human subjects show little muscle activity in the swing leg except at the start and end of a stride [117, 125, 100], for example. The compass gait captures the remaining “passive

dynamics” of human walking: an inverted pendulum stance leg coupled with a regular (downward) pendulum swing leg.

However, as we observed in Chapter 4, the purely passive compass gait walker with no regulating control is highly susceptible to falling on even mildly-varying terrain. For legged robotic locomotion to become practical outside of a controlled, laboratory setting, robots will need robust, dynamic control strategies which can cope with the high degree of stochasticity expected in day-to-day operation. Such robots will encounter variations in terrain profile and in characteristics such as contact friction and coefficient of restitution, and they will be subject to intermittent external forces.

Passive dynamic principles have inspired a promising avenue of research in legged robotics. Purely passive walkers, which require no actuation or active control, exhibit stable limit cycles walking down a gentle slope. Unfortunately, the inverted pendulum dynamics of these passive walkers also exhibit notoriously fragile basins of attraction; minor perturbations, such as subtle variations in terrain, tend to cause them to fall. Although there has been much recent progress in the study and design of robots based on passive dynamics principles [114, 56, 102, 33, 184], there are no good examples of an underactuated robot or model based on passive dynamics which can walk on significantly rough terrain.

It seems natural to posit, however, that such dynamics can be made significantly more robust through perspicacious design of control. In this paper, we present highlights of our results in the study of a compass gait (CG) walker model with torque actuation at the hip and an axially-directed source of impulse at each toe. The focus of this work is to assess the capabilities of this machine in negotiating difficult terrain; no attempt has been made to consider energy efficiency. Much of the material presented in this chapter also appears in [26].

5.2 Related work for CG models on rough terrain

Dynamic walking approaches to legged locomotion are inherently sensitive to perturbations, as discussed in Section 1.2.2. To date, little has been done either to address optimal control for passive-based walkers on rough terrain or to quantify their stability for such

walking. Some relevant approaches which could likely be adapted toward this goal are cited briefly below.

[44] use a manifold approach to control an underactuated compass gait walker on mildly rough terrain. This approach has shown promise, but it is currently limited to terrain with a gradually-changing slope. It does not address more typical rough terrain, where foothold selection is important, and/or variations in slopes are significant. By selecting appropriate manifolds, a simple biped might be able to negotiate more extreme terrain, but a methodology to do so (and to quantify resulting performance) remains an unsolved (and non-trivial) problem to date.

The hybrid zero dynamic (HZD) approach developed for underactuated bipedal walking in [185] also seems promising as a good low-level control strategy for negotiating rough terrain. The HZD approach addresses constant-slope walking with isolated perturbations [183] (i.e., to enlarge the basin of attraction of the walker), but it has not yet been applied toward rough terrain walking (and its mixing effects). We suggest the HZD approach could be developed for rough terrain walking by employing a hierarchical strategy analogous to the one we present in Section 5.5, where a family of low-level trajectories are designed from which one can select a high-level sequence of one-step control policies.

A third approach, presented in [162], stabilizes a fully-actuated compass gait, where a torque source is assumed at the stance foot contact with the ground. The fully-actuated version of the compass gait can indeed negotiate rough terrain with variations in slope. We find this approach highly undesirable, however, because we believe maintaining an underactuated contact where the foot meets the ground is an essential feature in walking.

Finally, we note [163] address quantification of stability on rough terrain by examining models of *human* walking which are based on passive dynamic models. This work acknowledges that guarantees of global stability are impossible for real-world walking systems, since it is always possible to find a perturbation large enough to cause an underactuated walker to fall. The authors of [163] investigate both short-term and long-term stability of a passive compass gait model on stochastic terrain with the eventual goal of making predictions on the risk of falling for humans. To date, this research has not provided such quantifiable estimates, however, and it certainly does not address the issues of control and

optimization of stability (passive or active).

5.3 Stability versus agility for compass gait walking

Before presenting results for the passive and active compass gait walkers, we will briefly discuss the considerations taken in selecting the particular parameters used for the passive and controlled experiments – which were in fact distinctly different. In Section 4.3.2, we discussed metastability of the compass gait walker, and we noted that changing the mass distribution along the legs can (not surprisingly) have a tremendous impact on the resulting stochastic stability of a passive walker. Referring to the basins of attraction for each walker shown in Figure 4-12¹ of Chapter 4, the basin at the top corresponded to the walker with a higher mean first-passage time on all the rough terrain cases examined. In this chapter, we will revisit the passive stability of the less stable passive walker from Chapter 4 (i.e., the one whose basin of attraction is given at the bottom of Fig. 4-12), and we will introduce a new parametrization (mass and length distributions) for the compass gait model we use in our optimal control simulations.

Our work on *active* control of the compass gait uses a model which has low passive stability, comparable to the worse of the two walkers presented in Chapter 4, because we have observed a trade-off between passive stability and active control authority. Although perturbations have less effect on the passively-robust walker, input from additional control also has less effect on the inherent dynamics. This trade-off is not surprisingly, as it is commonly seen in a variety of dynamic systems. In aircraft, for example, a well-built glider is designed to have greater passive stability than a high-performance jet; each design exploits an inherent trade-off between passive stability and active control capabilities for agile motion. Similarly, we have intentionally selected a walker with dynamics which are sensitive to our actuated force inputs. The actively-controller walker is qualitatively “bulkier” than its passively-stable counterpart, with more coupling between the stance and swing legs (through the commanded torque at the hip), and the bulkier walker has (as mentioned) fragile stability properties in the passive (uncontrolled) case.

¹on page 119

5.4 Passive biped model on rough terrain

This section has two aims. First, it quantifies the passive walking stability of the passive compass gait model we use throughout the control analysis in this chapter. This provides a baseline for comparison with the performance of the control strategies presenting in Section 5.5. The second purpose of this section is to expand somewhat on the passive compass gait results presented in Section 4.3.2, particularly in highlighting the inherent advantages of the Markov chain modeling approach (as opposed to Monte Carlo or other simulation techniques) in estimating mean first-passage times.

5.4.1 Passive compass gait model




Figure 5-1: The compass gait biped model.

The model presented in this chapter is not the so-called “simplest model” [50, 154, 102] for bipedal walking, since it does include a distributed mass (at the hip and legs) rather than concentrating all mass at the hip, but it does assume several important idealizations. Collisions are assumed to be perfectly inelastic and instantaneously, at which point the stance and swing legs trade identities. We also assume that the stance leg is effectively pinned to the ground during walking. That is, we assume the stance leg never violates the

friction constraint at the ground contact. This assumption is reasonable, since the standard deviation in slope height for any terrain tested for which the purely passive compass gait could successfully take more than one or two steps did not exceed more than about six degrees away from being flat terrain.

Having a distributed mass along legs makes the motions during walking more “interesting”; that is, the dynamics between the stance and swing legs are strongly coupled. This is particularly important toward understanding control of the *active* models presented. For our controlled models, a high-level policy selects a nominal interleg angle, to be regulated by a PD controller at the lower level. However, the motion of the swing leg and stance leg are coupled, so that the size of the step commanded affects not just the eventual foothold taken but also the angular velocities of the legs at impact, as well as whether or not the step is even dynamically feasible at all.

5.4.2 Methods

We explored the stability of the passive CG walker through two methodologies: Monte Carlo simulations and discretization of the dynamics into a Markov process. In both cases, terrain slope variations provide the only randomness in the system. The Markov modeling approach, which was described in Section 4.3.2 (on page 114), has several advantages, which are highlighted below.

Monte Carlo simulation

The most straight-forward approach to estimate the mean first-passage time (MFPT) for a model is, perhaps, Monte Carlo simulation [121]. This method uses computational brute force in cases where statistical interactions lead to analytically complicated dynamics and was originally developed for use on the ENIAC computer, to model dynamics of and estimate neutron multiplication rates for fission chain reactions [120].

The Monte Carlo method is useful for the compass gait dynamics, because there is currently no closed-form, analytic solution for the step-to-step transitions of this model. To estimate the MFPT, we simply start up the model from a particular initial condition

repeatedly and then track the long-term statistics of failure events.

The advantage of this approach is that relatively low effort is required to set up our simulations. The disadvantage is excessively long simulation time. For one thing, recall that our eventual goal is to obtain long-living, metastable walking cycles. In the limit as we analyze better dynamic systems, failure events become increasingly rare, and it takes a correspondingly long time to observe a statistically sufficient number of failures. Secondly, note that our real goal is not to characterize failures from *one, particular* initial state but rather for initial conditions which span some volume in the state space of interest.

There are some advantage to Monte Carlo simulation, however. One is that it allows for arbitrarily-accurate spatial resolution, if we are interested in analyzing a small subset of state space. For example, Figure 5-2 shows a 2D slice of the 3-dimensional set of post-collision states for the passive walker. Simulating to estimate MFPT over this plot took over a day of computation, but the resolution is much finer than we could obtain when using a discretized Markov chain model, due to matrix size limitations in MATLAB. Figures 4-11 (page 117) and 4-14 (page 121) show examples of the resolution possible when using a discretized Markov model across all of state space.

A second advantage of Monte Carlo simulation is that it can provide a useful, intuitive description of the statistical variation in the dynamic system. In Chapter 6, we will discuss stabilization of a bounding motion for LittleDog. A significant failure mode for this gait involves the robot tipping forward too far as it supports itself on the two front legs to retract the rear ones. This occurs when the robot pitches past a critical “tipping point”. Monte Carlo simulations can provide a good online estimate of the variation in pitch angle, which in turn can provide an estimate of the overall failure rate. It would take far longer to observe enough data to accurately model the dynamics over a complete mesh (by many orders of magnitude), and so Monte Carlo simulation can be used to capture important details of particular dynamics in cases where filling out a full mesh is not practical.

Markov chain model

Although Monte Carlo simulation is appropriate in some cases, modeling the dynamics of our walking systems with a *discrete Markov chain* generally has several advantages





Figure 5-2: Deterministic basin of attraction (top) and stochastic map (below) of mean first-passage time

The 2D stochastic map slice was obtained through Monte Carlo simulation. Dynamics were simulated for many trials from each, particular initial condition. This plot can only be obtained practically for a walker with weak stability properties (the blue regions corresponds to about 18 steps on terrain with mean slope 4° downhill and $\sigma = 0.5^\circ$); otherwise, simulation time is prohibitive. Top image also appears in Figure 4-12.

Table 5.1: MFPT for three different passive CG models
Slope at each step is drawn from a Gaussian with mean of 4° and $\sigma = 0.5^\circ$.

ID	m (kg)	m_h (kg)	a	b	MFPT (steps)
A	2	2	0.5	0.5	9.3
B	5	10	0.6	0.4	18.1
C	5	1.5	0.7	0.3	185,000

over this approach. First, while the Monte Carlo method requires starting at each initial condition and simulating many steps (until failure) for each of many trials, our simulation requirements are much more efficient for a Markov chain. From each discrete state in the Markov chain mesh, we just simulate *one* step for each of a discrete set of “noise” conditions (for example, for a set of upcoming slope angles). The act of taking a step during walking also provides a natural discretization in time which is well-suited to the Markov chain model.

Another big advantage of the Markov chain is that it allows for the one-shot calculation of state-dependent MFPT and for the system-wide eigenanalysis of the dynamics². Finally, the discretized model lends itself to dynamic programming, which will allow us to solve efficiently for optimal control policies later in this chapter.

5.4.3 Discussion of passive performance

Given appropriately chosen parameters for the mass distribution of the walker and the stochasticity of the terrain variations, the results both from Monte Carlo simulation and from representing the dynamics as a discrete Markov chain show that the passive compass gait walking on rough terrain demonstrates typical metastable behavior. That is, initial conditions are forgotten relatively quickly, there is a region (the “stochastic basin of attraction”) across state space where MFPT is essentially uniform, and there is a second region (near the fixed point of the deterministic system) where the dynamics tend to remain over time. We will observe these same characteristics in successful (metastable) versions of the actively-controlled dynamic system.

²Both were described earlier in Section 4.2.2, on page 103.

Table 5.1 compares the mean first-passage time for each of the three passive walkers described throughout this thesis on a particular, stochastic terrain. Data for walkers A and B are obtained from Monte Carlo trials, while the result for walker C came from an eigenanalysis. We would expect to wait an average of 1.4 hours to observe a single failure event for this walker if we ran Monte Carlo computer simulations, so the eigenanalysis is truly much more practical here! Note that model B is the walker analyzed to produce Figure 5-2, and both B and C were examined in Section 4.3.2 of the previous chapter.

The walker on which we will optimize control throughout the remainder of the present chapter is model A. Qualitatively, the goals of active control are to increase stability by both enlarging the stochastic basin of attraction (of states with uniform MFPT) and speeding up the mixing time in which initial conditions are forgotten.

5.5 Optimal control of the biped on rough terrain

5.5.1 Actuated compass gait model




Figure 5-3: Actuated compass gait model with torque source at hip. Simulations for the actuated compass gait use the following parameters: $m = m_h = 2\text{kg}$, $a = b = 0.5\text{m}$

The basic compass gait model, introduced by [56], is used in all simulations and depicted in Figure 5-3. In addition to the modeling details previously given for the passive compass gait (in Section 5.4.1), the actuated model now also includes a pure torque source at the “hip” linkage where the legs meet. The stance leg is oriented at an absolute angle θ_1 and contacts the ground at a zero-torque, frictionless joint, called the “toe”. The swing leg position is measured relative to the stance leg at angle θ_2 . Parameters for the model used for the controlled walker are given in the caption of Figure 5-3, and the direction of walking for all simulations is from the left to the right.

To address the practical concerns of defining and detecting ground collision, we assume that the toe of the swing leg is immediately retracted once a ground collision occurs and remains so until the stance leg reaches a position within 10° of its desired interleg angle, at which time it extends instantaneously. If, after extension of the toe, we discover the tip is below our ground height, we consider the walker to have failed to have taken a successful step. For successful steps, the identities of the stance and swing legs are swapped when the (extended) swing leg toe finally collides with the ground.

In addition to the hip torque actuation, we also include an ideal impulse actuation at each toe which we assume can be delivered at a prescribed magnitude, axially at the stance leg contact immediately preceding each collision between the current swing leg and the ground. Inclusion of an impulse source is inspired by the work of [100], who illustrate that this is an efficient method for imparting energy to the system to make up for energy lost in ground collisions³.

5.5.2 Method for approximate optimal control

To solve for the approximate optimal control solution for the compass gait walker on rough terrain, we discretize the dynamics, always using a post-collision state to define the state. We then use the value iteration algorithm [166] to find an optimal step-to-step feedback policy for the discretized system. We tested two low-level control strategies for defining the state-to-state dynamics. The first method was a Proportional-Derivative (PD) controller

³Note that [100] use the simplified version of the compass gait, with all mass concentrated at the hip and infinitesimal mass at each toe. However, the effect is similar in the model with distributed mass.

regulating the inter-leg angle, α . The second method was an impulsive toe-off, applied just before each ground collision. We tested each method alone and also tested both together. More details on implementation of the simulations are given below.

Value iteration algorithm

Our goal is to find a policy which will maximize the number of steps taken before falling. To determine the full capabilities of the walker, one would wish to discretize over time and to select a torque magnitude at each dt . We chose instead to discretize on a step-to-step basis, to reduce the size of the state space (by one dimension) in order to keep the problem computationally tractable. We were pleased to find that such a strategy still produces continuous walking on significantly rough terrain in our simulations.

The post-collision state of walker is represented in most plots using four meshing state variables: X_{m1} : x location of the stance leg; X_{m2} : the distance Δx from the stance leg ground contact to the swing leg contact; X_{m3} : the angular velocity of the stance leg, $\dot{\theta}_{st}$; and X_{m4} : the angular rate of change of the inter-leg angle, $\dot{\alpha}$. Details on the discretization used are given in Appendix B. Note that the inter-leg angle is increasing when the swing leg rotates “forward”. Counter-clockwise is positive for all angle measurements, as illustrated by Figure 5-3.

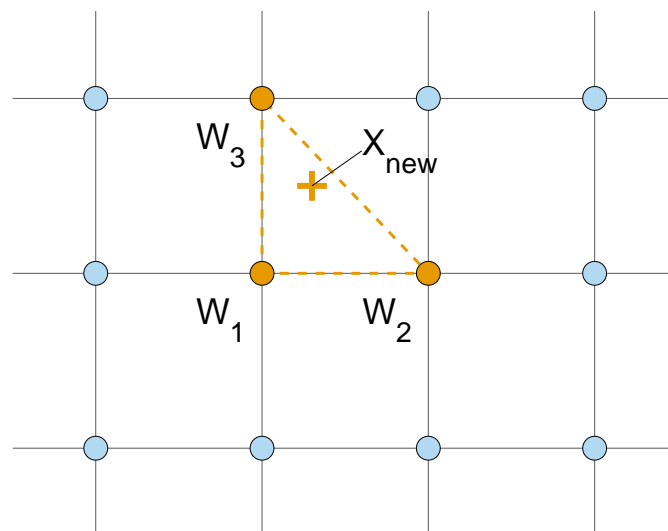


Figure 5-4: Illustration of meshing approximation

The value function and feedback policy between the discretization points must be approximated, since values are only stored at the particular mesh nodes. Figure 5-4 illustrates the effect of meshing; each new post-collision state, X_{new} , which we observe after a simulated step, must be approximated by a weighted sum of some subset of nodes which exist in our pre-determined mesh. The minimal number of nodes to create a meshed volume in an N -dimensional mesh is $N + 1$; for our 4-dimension mesh of post-collision states, this means each exact new state is modeled as if transitions occurs to 5 neighboring mesh nodes. The individual transition probabilities are set through barycentric interpolation [129] (and of course must sum to unity). Any post-collision states falling outside the defined mesh range were automatically binned into the absorbing failure (fallen) state.

In solving for the optimal policy, we begin by assigning a value $C_{last}(k) = 0$ for each state (k) in the mesh. Through iteration, we update these values to reflect which states can achieve the greatest estimated number of steps before falling. Equation 5.1 shows the update for the estimated cost of starting in state i and performing control action option a .

$$C_{new}(i|a) = \sum_{k=1}^5 W_k [\gamma C_{last}(k) + C_{onestep}(k)] \quad (5.1)$$

After this update is done for every possible *control action*, we can then update C_{last} for each state, so that it is the lowest of all C_{new} values (corresponding to the best action found so far):

$$C_{last}(i) = \min_a C_{new}(i|a) \quad (5.2)$$

and our optimal n -step policy is the set of actions at each step (after n iteration steps) which gives this lowest cost:

$$\pi(i) = \arg \min_a C_{new}(i|a) \quad (5.3)$$

We include a discount factor, $\gamma = 0.9$, in Equation 5.1 to ensure that the value function converges to a finite value everywhere as the number of iterations goes to infinity. The one-step cost function determines exactly what it is that we are optimizing, as described below.

Cost function to optimize metastability

To maximize the expected value of the *number of steps taken* from each state in the mesh, the appropriate one-step cost is simply:

$$C_{onestep}(i) = \begin{cases} 0 & i \in \text{fallen} \\ -1 & \text{otherwise} \end{cases} \quad (5.4)$$

In words, falling down yields no cost (or reward), while successfully making *one* step is rewarded with the value *minus one*, to indicate a single additional step has been taken. Technically, we are not maximizing the mean first-passage *time* in a literal sense here; we are instead maximizing the *number of steps taken*. When we refer to the MFPT in this chapter (and elsewhere), we are generally measuring “time” in number of steps taken, both for convenience and because our discretization (using post-collision states in the mesh) makes the notion of measuring time as discrete steps taken quite natural.

If we *did* wish to maximize the time traveled instead, the one-step cost would depend on the trajectory taken from one step to the next. We could perform this bookkeeping, if desired, and then the non-fallen one-step cost would equal the *negative* of the time taken to perform a particular (successful) step.

Another intuitive option is to maximize the *distance* traveled. To do this, we would assign the *negative* of the *distance traveled* in one step. For any post-collision step, this is simply the (negative) distance between the stance and swing toes (and does not depend on the trajectory taken to get to this state). Note that relative distance between stance and swing toe is one of the four state variables used to define the post-collision state, so implementing this one-step cost is quite trivial in practice.

Hierarchical controller design

Our simulations use two simple low-level control options: PD control of a desired inter-leg angle and impulsive stance-foot toe-off when the ground collision occurs for each new step. We tested three control strategies, using each of the two low-level controls alone and combining the two methods in creating a third technique. The high-level control actions at

each step in the value iteration correspondingly consist of one or both of the following: (1) a desired inter-leg angle, α_{des} , for the PD controller and (2) the magnitude of the impulse.

The primary purpose of the PD controller is to regulate of the step length of the walker, which in turn selects upcoming foot placement on the terrain. However, the dynamics are coupled, so that the controller also affects the entire dynamic state of the walker. There is currently no closed-form description for the step-to-step transitions resulting from this coupling, which is precisely why simulation of the dynamics provides a practical methodology for investigating the system.

The main goal in employing the impulsive toe-off action is to compensate for the energy which is lost at each ground collision. This in turn allows the walker to take larger steps than would otherwise be possible, since more energy is of course lost for larger step angles [32].

PD control of inter-leg angle

A low-level PD controller regulates the inter-leg angle, α , which is defined as:

$$\alpha = \theta_{sw} - \theta_{st} \quad (5.5)$$

Our PD controller was designed by hand, to obtain swing leg motions which do not overpower the stance leg dynamics entirely but which are still effective in regulating desired step length approximately. The PD controller is active throughout the course of any particular step. Equation 5.6 gives the commanded controller torque at the hip:

$$\tau = K_p(\alpha_{des} - \alpha) + K_d(0 - \dot{\alpha}) \quad (5.6)$$

where $K_p = 100$ and $K_d = 10$.

Impulsive toe-off at impact

In this control strategy, a particular magnitude impulse is applied axially from the stance leg downward, toward ground. Toward eventual implementation on real robot, we note again here that it is well-known that applying the impulse immediately *before* collision is more

efficient in imparting energy to the compass gait walker [100]. We assume the impulse is imparted instantaneously in our simulations; correspondingly, no additional dynamic simulations are necessary to implement value iteration with this when this control actuation is included.

A high-level feedback policy is evaluated at each the post-collision state, but we desire that a commanded impulse affect the current step. To model the transition with an impulse, the post-collision state is simply "rewound" to its corresponding pre-collision state, using the well-known equations for conservation of angular momentum during the inelastic collision [56]. Then, the contribution of the impulse (mv) is added as appropriate at the toe of the stance leg, and the equations for angular momentum (now going from pre- to post-collision) are applied once again. Note that the velocity of the swing leg must be tested after the impulse is applied to ensure that the velocity of the swing leg toe is still directed toward the ground. A large enough impulse could theoretically send the entire walker airborne! However, this has not been a practical concern for the impulse magnitudes we have tested in our simulations.

Our initial tests using value iteration allowed for the selection of one of 21 values ranging from 0 to 2 (kg-m/s). In practice, however, the largest magnitude of impulse was almost always selected. To simplify the controller and to reduce the run-time for value iteration (by a factor of 21), we eventually set the impulse value to a constant magnitude of 2 at every step.

5.6 Controlled CG on wrapping rough terrain

One significant goal in this chapter is to demonstrate the inherent capabilities of a compass gait biped with a torque-free pin joint where the stance leg meets the ground, walking on rough terrain. To explore the dynamic capabilities of this simple walker, we wish to allow perfect knowledge of upcoming terrain. To allow for perfect knowledge of the terrain for an arbitrary number of steps in the future, we first created a model with wrapping terrain. Our results demonstrate more capable walking for such a minimalist biped than can be found in the current state of the art. The results provide some evidence that this simple model

captures a large part of the dynamics not only for constant-slope walking but also for more realistic locomotion, where the environment adds noise.

Our results are presented in detail before. The meshing discretization used in these simulations is described in more detail in Appendix B. Figure 5-5 illustrates example terrain (at a somewhat exaggerated scale) and describes the four meshing states used.




Figure 5-5: Compass gait on wrapping terrain

At left, the actuation for this model includes a torque at the hip and a constant magnitude impulse which occurs at each step. At right, the four states which define each post-collision elements in the discrete mesh are given.

5.6.1 Wrapping rough terrain models

We attempted to solve the optimal control problem for each controller on a variety of both continuous and discontinuous terrains. Examples are shown in Figures 5-6, 5-7, 5-10 and 5-12. Step-to-step transitions were always computing using a continuous version of a particular terrain profile. When solving the value iteration problem for a discontinuous terrain, regions where gaps exist were simply defined as “dead” zones; any state involving a gap region as its next foothold transitioned immediately into the failed (fallen) state, by definition. Our tests of the hip-actuated compass gait walker on discontinuous terrain were largely inspired by the impressive results from [71] in using a simple, intuitive strategy to control the Raibert hopper on intermittent terrain. We note that many of the same basic principles which have proved effective for simple hopping robots have likely guided the

practical design of impressive, robust control for the dynamic BigDog quadruped (created by Boston Dynamics) on rough terrain [21]. We hope that the future for legged bipedal robots can follow a similarly impressive path, using lessons learned from the compass gait model as a basis for successful controllers on advanced robots based on passive dynamic principles.

To test the performance limits of each control strategy analyzed, each terrain was scaled to make its features more dramatic until value iteration failed to converge on a stable walking solution. Each terrain consisted of a particular profile which repeats every 7 meters, as shown in Figure 5-6. This allows value iteration to converge on a fixed policy using a finite representation for the terrain. A repeating terrain may also be thought to represent the terrain of a compass gait walker circling endlessly on a fixed boom; we intend to experiment with the control strategies described here on a real robot mounted to a such a boom, as described later in Section 7.2⁴.




Figure 5-6: Intermittent-foothold terrain. Terrain profile repeats every 7 meters, allowing the numerical mesh to “wrap”.

5.6.2 Results on wrapping terrain

This section highlights our simulation results for each of the three control strategies described in Section 5.5.2 on various wrapping terrains described in Section 5.6. Summarized briefly, the PD controller is effective in regulating foot placement on terrain; however, it allows for disappointingly limited variations in step width or height. The toe-impulse action is *intended* to enable the walker to negotiate more extreme terrain by pumping additional energy into the system efficiently, but performance of the impulse control alone

⁴See Figure 7-3 on page 178.

does not regulate the upcoming step length effectively, and so the resulting walker is unfortunately quite fragile. Combining a constant magnitude toe-off with the PD controller provides a significantly better design than either component demonstrates alone, allowing for significantly greater variations in both step length and height during continuous walking, through the combined actions of energy addition (through toe-off) and step length regulation (through hip torque).

Performance using only PD control

Figures 5-7 and 5-8 show examples of terrain which our simulation successfully negotiates using only the PD controller. The performance is good, but it is importantly limited by the inefficiency of adding energy through hip torque, alone. A single step which is too wide or too high results in rapid failure, as the compass gait simply cannot complete a subsequent step after a large loss in energy.




Figure 5-7: Examples of terrain which were successfully negotiated using PD control alone. Terrain at top is flat except for a set of bottomless no-go gaps. The middle and bottom examples consist of piecewise slopes. Each was successfully solved both as a continuous terrain, shown as a solid black line, and as an intermittent terrain, represented by the solid regions. Figure 5-8 shows detail from the bottom plot.

Several features can be noted about both the performance capabilities of the walker




Figure 5-8: A close-up of the bottommost terrain in Figure 5-7.

and the nature of the resulting optimal control policy. Qualitatively, the walker is able to perform much larger steps when the terrain is sloped downhill, as one might naturally expect. Also, the more extreme (“difficult”) the terrain was, the more liable the optimal policy was to select the same, repeated set of footholds over a particular portion of the terrain, as illustrated in Figure 5-9.

The terrains shown throughout the paper are what the author considers to be the most impressive examples for which successful optimal policies could be found; on more extreme terrain, there were no solutions which resulted in continuous walking. That noted, it is interesting to compare the features of the three types of intermittent terrain in Figures 5-7. On the flat (topmost) terrain, gap width was about 11 cm on average, up to a maximum of 18 cm; for the piecewise sloped terrain (bottom), the gaps needed to be somewhat smaller: 8 cm on average up to a maximum of about 10 cm. The downhill terrain (shown in the middle) has an average grade of 4° and correspondingly permits dramatically larger steps: 30 cm on average, with the larger step being about 47 cm – or nearly half the leg length.

The overall variation in terrain height is not significantly different for the downhill terrain with large steps as compared with the piecewise sloped terrain. Excluding the part of the terrain with intermittent gaps, the piecewise terrain has a standard deviation (σ) of 1.07 cm, and a max-to-min height difference of 5.15 cm. Subtracting out the net slope of 4° , the “downhill” terrain height has a standard deviation of 1.26 cm and max-to-min height different of 6.15 cm. Step width on this terrain does vary significantly (from about 28 cm to 63 cm), but it cannot maintain a gait with continuously small footsteps without tripping forward and falling after a few steps. Experimentally, we find torque-only actuation is often insufficient to completely regulate energy.

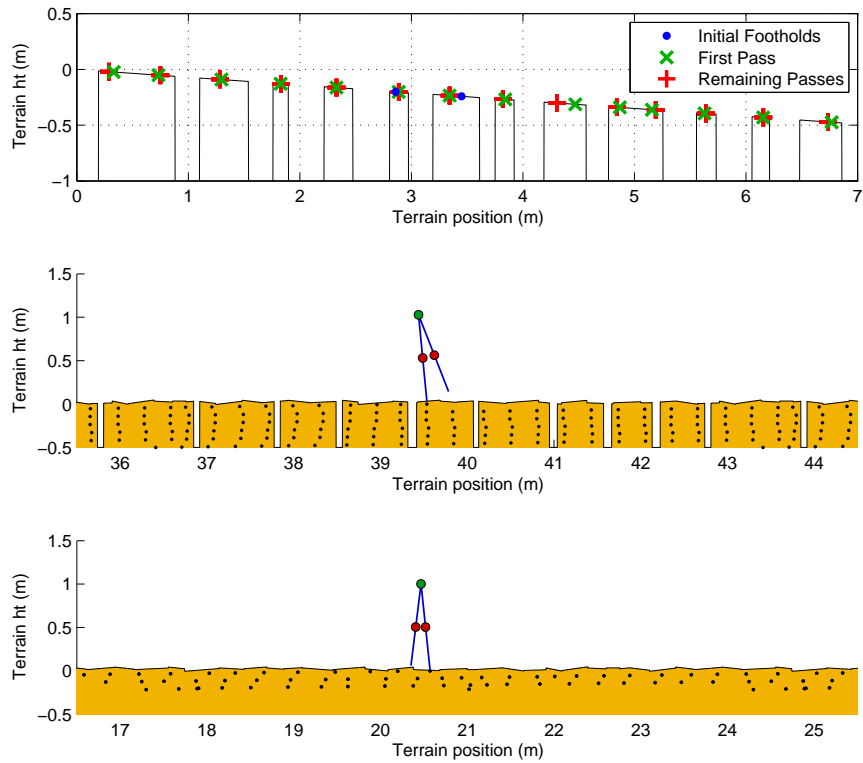


Figure 5-9: Foothold patterns from optimal control.

The optimal footholds on extreme terrain quickly converge to fixed pattern (top). On easier terrain, no fixed pattern emerges (bottom). In intermediate terrain, we see an intermediate level of organization in the pattern of footholds (middle). Footholds taken are plotted as dots in the lower two figures, with the earliest steps plotted lowest and with dots repeating as the terrain repeats, every 7 meters.

Performance using only toe-off control

It was hoped that an impulsive toe-off controller might produce better results than the PD controller. However, the toe-off control alone results in surprisingly fragile walking cycles which require fairly precise initial conditions and allow for almost no variations in terrain height. Figure 5-10 shows the footholds taken using a constant impulse value of $mv = 2$ (kg-m/s). The terrain is completely flat for the first steps, until $x = 4$ m. We intentionally illustrate a terrain where the walker can successfully take a single step after a change in terrain. It takes one step with a resulting downward slope of -0.7° and then trips up while attempting a subsequent step at 0.35° . The performance is quite poor, even if one allows for the selection of the magnitude of the impulse. Without more direct regulation of the upcoming step length, the walker quickly falls on even trivial terrain.




Figure 5-10: Footsteps taken using impulsive toe-off control, only. Terrain is completely flat between 0 and 4 meters, varies mildly (always within ± 1 cm) between 4 and 7 meters, and repeats exactly at 7-meter intervals.

Figure 5-11 shows the dynamic states (angles and angular velocities) of the walker during the same trial for which footsteps are depicted in Figure 5-10. These data are presented primarily to allow quantitative and qualitative comparison with the dynamic states during the trials using both PD and impulse control, shown in Figure 5-13. With the torque actuation from the PD controller, the steps are generally both much larger (X_2 is larger) and

faster (there are more steps per second, and X_3 and X_4 are both larger when hip torque is employed). Note that two of the dynamic states plotted in these figures differ from the meshing states defined in Section 5.5.2: $X_1 = \pi/2 + \theta_{st}$ and $X_2 = \pi + \alpha$, while $X_2 = X_{m2}$ and $X_3 = X_{m3}$.




Figure 5-11: Dynamic states using impulsive toe-off control, only.

Performance using both PD control and toe-off

Our most impressive performance by far is obtained by combining the low-level PD controller with the an impulse of magnitude 2 (kg-m/s) preceding every ground collision. The optimal control policy converges to a successful policy for continuous walking on the terrain shown in Figure 5-12, having a standard deviation of 4.7 cm and a max-to-min difference in height of 19.1 cm. The deviations in the width and height of the actual footsteps taken using the optimal policy are shown in Figure 5-14; the height of the footholds shown have a standard deviation of 3.3 cm, and the SD of the inter-leg angle at each step is about 5.9°.

5.7 Controlled CG on stochastically rough terrain

In addition to the wrapping terrain models described in Section 5.6, we also solved for the approximate optimal policy for the same, controlled walker on stochastically rough terrain. In the stochastic case, the height difference between successive footholds is drawn from a Gaussian distribution. Figure 5-16 illustrates the nature of the resulting terrain. This




Figure 5-12: Footsteps taken during a 60-second trial using the optimal control policy from value iteration (top) and using a one-step time horizon (bottom).




Figure 5-13: Dynamic states during a using the optimal control policy from value iteration (top) and using the one-step control strategy (bottom).




Figure 5-14: Comparison of step width and step height for the optimal control solution vs. for one-step control. Note that most of the steps selected using the one-step approach are near a nominal value of 28° , which results in a step width near .48 meters.




Figure 5-15: Smoothness of cost function over state space.




Figure 5-16: Compass gait on stochastic terrain

At left, the actuation for this model includes a torque at the hip and a constant magnitude impulse which occurs at each step. At right, the four states which define each post-collision elements in the discrete mesh are given.

model for the terrain is somewhat artificial, since height would typically be a function of step length on real terrain; by contrast, the next upcoming step in our simulation occurs at a particular pre-selected height, regardless of step length. However, this compact terrain model can still produce arbitrarily extreme terrain (by increasing the standard deviation of the step heights, as desired).

Based on the results obtained on wrapping terrain, we included both the low-level PD controller and the constant-magnitude impulsive toe-off in all simulations on stochastic terrain. As in the wrapping terrain case, details on the discretization of meshing are given in Appendix B.

In our initial simulations on stochastic terrain, the post-collision state (which again used four state variables, as for the wrapping terrain) captured the current orientations and velocities of the two legs; however, it did *not* capture any information about the specific, upcoming terrain. The resulting solution maximizes the mean first-passage time for a walker which is blind to the *specifics* of the terrain but does know the *statistics* of what may happen. Optimal control solutions for this “blind” walker were quite poor when compared with the performance obtained on wrapping terrain with statistically similar variations in terrain height.

On wrapping terrain, the upcoming profile of the terrain was inherently imbedded within the simulation itself. Our meshing state included the absolute position of the stance foot on the wrapping terrain, and as each new step was simulated, the z value of the terrain was always a deterministic function of the particular x location of the swing leg toe during its trajectory. This means that on wrapping terrain, the walker is employing a potentially “infinite” look-ahead; the number of steps we look ahead is determined by the number of steps performed in the value iteration⁵.

Intuitively, including some information about the immediately-upcoming terrain should improve the performance of our underactuated biped significantly. For example, it would allow the biped walker to execute a shorter step (which loses less energy) if going uphill or a longer step going downhill. We tested this hypothesis by implementing a one-step lookahead. This required enhancing the 4-variable representation of the post-collision state of the walker with an additional, fifth state: the Δz value of the next, immediate step.

This one-step lookahead improved the mean first-passage on stochastic terrain dramatically. Figure 5-17 compares the MFPT for the one-step and no lookahead cases. For example, on terrain where the value of Δz is drawn from a Gaussian with zero mean and a standard deviation of 1 cm, the MFPT is about 76 with no lookahead, versus 12,000 with a knowledge of just the next, immediate step.

Given this result, it is natural to enquire how a two-step or (more generally) an n -step lookahead would compare on a similar plot. Unfortunately, obtaining such results would require multiplying the size of the total mesh by a factor of 19 for each additional lookahead, because our representation of the upcoming terrain is discretized to be drawn from one of 19 particular values. To compare the one-step lookahead with a longer lookahead, we will instead return (in Section 5.8, ahead) to the case of wrapping terrain, where arbitrary lookahead can be obtained without adjusting the size of our state space (nor our mesh).

⁵We used 100 steps in value iteration on both wrapping and stochastic terrains, although in practice the policy usually converged within a few steps.




Figure 5-17: MFPT for one-step vs. no lookahead on stochastic terrain

5.8 One-step policy on wrapping terrain

For our passive walking examples (in Chapter 4 and in Section 5.4 of the present chapter), we noted that initial conditions are forgotten very quickly when walking on rough terrain. Correspondingly, we posited that the ability to negotiate such terrain may only require a short lookahead. Assuming the walker has not fallen down, a quick mixing time (to forget initial conditions) implies the consequences of decisions made now should not have a significant impact on stability of its state, given a few, successful steps are taken.

To explore this hypothesis, we began with an estimate of the value function for each post-collision state from a model of *stochastic* terrain which had approximately the same standard deviation in step-to-step terrain height as our new *wrapping* terrain. Then, we ran the value iteration algorithm again on wrapping terrain – but stopped iterating after only a single pass. This approach works surprisingly well, resulting in continuous walking on terrain which (as stated earlier) was intentionally scaled until the optimal solution (with essentially infinite lookahead) was only barely able to achieve continuous walking.

Interestingly, although this policy uses only a one-step lookahead, simulations of the biped on wrapping terrain show that a repeating pattern of footholds are automatically

selected over time. Figure 5-18 shows this convergence of states over time. As in the case for optimal solution, we notice selected footsteps are more apt to repeat themselves on more extreme examples of wrapping terrain; on easier terrain, the foothold selection appears more haphazard. We suspect this occurs because particular regions of the terrain have a locally-optimal “goal” foothold, which acts to regulate an entire emergent pattern on wrapping terrain.

Finally, we should emphasize, briefly, that resetting the value function with the optimal values found for control on stochastic terrain is a critical step: if the value function is initiated as zero everywhere (for example), then a single pass of value iteration only results in a value of either 0 or -1 at every mesh state, for each possible action. This means the only information we have is whether or not a single step will be successful. Not surprisingly, the resulting policy does not work well in our simulations. The walker has no notion of the stability of its internal dynamic state when the value function is initiated with such a “blank slate”; using the solution from stochastic terrain to preset the value function intrinsically provides some information about the capability to take future steps.

The fact that the one-step control policy is nearly as effective as the optimal policy on our wrapping terrain simulations provides some evidence that use of a limited lookahead (e.g., 1 to 3 steps) results in near-optimal performance. We find this particularly compelling, as it correspondingly implies that near-optimal results should also be possible on real-world legged robots having only a short-sighted knowledge of upcoming terrain. We anticipated this result after analyzing our results on passive CG walkers, where we noted a surprisingly fast mixing dynamics of the metastable limit cycle generated by a passive compass gait walker on statistically rough terrain [25, 27], and further work is currently planned to test this hypothesis on a real compass-gait robot, as discussed in Section 7.2.

5.9 Discussion

Before concluding, we wish to bring several issues to the attention of the reader in this section, concerning both the passive and active compass gait simulations.




Figure 5-18: One-step policy shows convergence on wrapping terrain. Even when only a one-step lookahead is used to fix a policy on wrapping terrain, the states converge over time to a particular, repeating pattern. At top (left) are the four states over time (in color) with a time-shifted version of the data overlaid (in thin, black lines); a close-up (top, right) shows they have converged after about one cycle on wrapping terrain. At bottom is a logarithmic plot showing the deviation in state from n to state $n + 18$, i.e., after one complete 18-step “lap” about the wrapping terrain; the plot illustrates an exponential convergence of the controlled system to a fixed trajectory through state space.

5.9.1 Policy interpolation

First, we note that care must be taken in interpolating the discontinuous policy for walking on rough (most particularly, on intermittent) terrain to select control actions. We discovered a persistent and catastrophic interpolation error which occurred from time to time when evaluating the optimal policy during simulations. There are regions on many of the terrains (such as gaps) which are clearly infeasible for a step, and as a result, there are transition points where one must decide to take either a small step or a large step to avoid such regions. We found there were occasionally but consistently states where a simple interpolation of the optimal policy resulted in a medium-sized step which lands in a no-go zone. Such a situation is illustrated in Figure 5-19. Although [129] provide an excellent reference on issues of stochasticity and discontinuity of cost function in implementing the solution to a dynamic system via value iteration, their solutions can be expensive in practice and do not directly address our particular interpolation issues, which are associated with sharp discontinuities in our policy.

In solving this interpolation problem, we noted that although the *policy* was often discontinuous, the *value function* generally remained smooth. For example, there would often be several feasible step sizes the walker could take from a particular location in state space. As you move in state space, a transition occurs where it becomes “safer” (in terms of maximizing mean first-passage time) to take a large step instead of a smaller one, or vice versa. The expected MFPT remains essentially flat over a large swath of state space, although the selection of optimal footstep size changes discretely.

Our solution was to simulate a single step with each of many possible step length. Unsuccessful steps were identified and eliminated as control options. We then sorted the remaining control actions to select the optimal discrete action with the lowest cost. When two neighboring actions to this discrete action were also feasible, we used a quadratic fit of the set of three discrete actions and their function value to interpolate to estimate the true minimal cost action over continuous action space. Essentially, this changed our problem from one of directly interpolating the *policy* (which was inherently discontinuous) to one of indirectly interpolating over the *cost function* which results when one selects a particular




Figure 5-19: Policy interpolation

Sometimes, taking either a small step (top) or a large one (bottom) is acceptable, while a medium-sized step (center) would be deadly. In practice, there are regions in state space where the optimal action changes abruptly from a choice of a small step to a large one. Special care must be taken in interpolating the optimal policy near such discontinuities to avoid taking a deadly step.

(possibly non-optimal) action and assumes optimal actions will be selected from then on.

5.9.2 Efficiency versus stability

A second issue we bring to the reader's attention here is that our solutions do not attempt to exploit the low cost of transport and efficiency which are so characteristic and appealing in particularly well-designed passive dynamic walkers, such as the Cornell biped [35, 99]. Our purpose here was to demonstrate the theoretical performance possible for passive-based walking with an inherently underactuated ankle. Further study can be done using a cost function which takes into account both the distance traveled and the cost of transport.

5.10 Conclusions

The passive compass gait walker is inherently quite susceptible to noise. Small variations in terrain flatness or small impulsive disturbances (being shoved about lightly) tend to topple it over within a few steps for all but the most trivial (to humans, qualitatively) of noise levels.

However, a controlled compass gait biped can successfully negotiate an impressive range of rugged terrain using a surprisingly simple control strategy, even with ankles which remain completely unactuated (no torque) in rotation at the contact point. We have shown that a simple yet effective low-level control strategy can be obtained by combining a toe-off just before each ground collision with a PD control loop on the desired inter-leg angle. The magnitude of the toe-off can be a constant value for all footsteps, so that the high-level control at each step consists solely of the desired inter-leg angle. The action of the toe-off insures energy is added efficiently at each step, while actively controlling the step width tends to regulate the dynamics around a nominal trajectory in state space.

Chapter 6

Toward Metastable Bounding for LittleDog

This brief chapter presents early results toward modifying the dynamic lunge presented on page 85 of Chapter 3, to produce a more stable, cyclical, bounding gait. We present a general design approach but also note that additional feedback is required to make this gait truly metastable.

One of the goals of this chapter is to investigate application of the concepts introduced in Chapters 4 and 5 on systems of higher dimensionality; specifically, we revisit the underactuated, 18-degree-of-freedom LittleDog robot. The same modeling approximations made throughout the kinodynamic planning presented in Chapter 3 reduce the dynamics of the system to a small number of degrees of freedom, and this in turn simplifies our analysis of the experimental results. Specifically, we again assume a 2D (planar) model. We model the robot as a “brick” of constant inertia, and we assume it contacts the ground at a single point through a massless leg. On the real robot, this modeled ground contact point corresponds to the midpoint between the two support feet on the robot during a given double support phase. The planar model has three degrees of freedom: x , z , and the pitch angle (α).

Results presented in this chapter are preliminary and are primarily included to emphasize that the methodology used in the preceding two chapters is also relevant to higher degree-of-freedom robots, as well. Perfecting the stabilization of this gait is outside the

scope of this thesis.

6.1 Bounding gait development

This chapter focuses on the toy example of developing a reliable bounding gait for LittleDog which can negotiate relative benign (flat) terrain. However, the same general ideas presented could be extended to develop reciprocating, sequential double-support motions on rougher terrain. A basic approach to do so is to sequentially decide (1) which part(s) of the motion should be adjusted, (2) what signals (state information over time) to use to define a policy function, and (3) how to design a family of possible control actions. The remaining task is simply to tune algorithmic parameters appropriately so that robot observes the results of adjusting the policy function and updates the policy to improve reliability.




Figure 6-1: LittleDog video frames during rocking

The top row show sequential frames in a successful rocking trial, while the bottom 4 frames end in a failed “headstand” pose.

6.2 Experimental results

This section outlines our initial results in designing a continuous bounding gait for LittleDog.

6.2.1 Rimless wheel motion trials

We begin with a standard dynamic lunging motion, of the type initially described in Section 3.7.5¹. Next, we need to achieve a rocking motion to support the robot with only the front legs and to bring the rear legs forward for the next lunge. One way to accomplish this nose-down rocking is to position the front legs just a few centimeters forward of the center of body during touch-down, so that the angular momentum of the robot about these legs is not entirely lost when the front feet land. Our goal is to have the robot rock onto the front legs in a motion similar to a rimless wheel which rocks up only partway – without flipping too far head-over-heels.

Figure 6-2 shows data for multiple trials of rimless-wheel style rocking. Only about 50% of the trials are “successful”, such that the robot returns to zero pitch with all four feet on the ground. The other half of the time, the robot rotates too far and lands “on its head” at a final pitch angle of about 70° .

In about half of the trials, there is a kink which occurs during the initial pitching motion. The joint trajectories for the rear leg come near a singularity in the direction of loading in the knee, so that in some trials, the loading on the knee reverses sign. The behavior here is nonlinear because of the backlash in the gears. The data in Figure 6-2 are plotted either in red or blue, to classify whether there is a kink caused by this knee-loading singularity. Segregating out only the kink-free trajectories from all 108 trials, we see that the repeatability is quite good: the mean value of the pitch is -26.6° and the standard deviation is 0.84° . Although the trajectories which include a kink cluster about a different peak value, they are also fairly repeatable: the mean value of the pitch for the kinked set is -24.6° , with a standard deviation of 0.95° .

The trial data in Figure 6-2 are represented as a transition map in Figure 6-3. One point is plotted for each trial, with the initial peak nose-up (negative) pitch angle plotted on the x-axis and the secondary peak nose-down (positive) pitch angle plotted on y. There are three distinct regions in the data. For each classification of trial (red or blue), there is a “stochastic transition map” analogous to the return map for the rimless wheel, described in Section 4-4 (on page 109). A peak angle in pitch correlates to potential energy in the system, when

¹Page 85.




Figure 6-2: Data for rimless-wheel rocking trials

all velocities are zero. Stochasticity in the system results in a “noisy” relationship between the successive peak pitch angles (nose-up and nose-down), but they clearly demonstrate a direct correlation.

Trajectories ending at the top of Figure 6-2 correspond to trials where the robot tipped too far and landed on its head, as depicted in the final (lower) frame of Fig. 6-1. From Fig. 6-2, we can see that this happens any time the secondary pitch angle exceeds about 40° .

6.2.2 Triple-rocking trials

In the previous section, all of the joints of the robot were held in place during the secondary rocking motion onto the front legs. This resulted in a noisy correlation between nose-up and subsequent nose-down peak pitch angles and a correspondingly high failure rate. There was simply no particular position to aim the feet to achieve a non-trivial average nose-down pitch peak without also observing a relatively high rate of failure (ending on the head).

To improve repeatability, we wish to have the stochastic transition map data in Figure 6-3 converge to the same average nose-down angle for all nose-up motions. That is, the




Figure 6-3: Stochastic transition map for rimless-wheel rocking trials

dynamics of the controlled system should “forget” the initial (noisy) nose-up pitch apex as rapidly as possible. Pictorially, this would result in a flat, horizontal line on the transition map: all initial conditions would go to the same end distribution.

We find that the peak pitch angle of the nose-down rocking motion can be made more repeatable through the addition of energy at front leg touch-down. We accomplish this through a planned kick-off of the rear legs combined with rotation of the front shoulder. Figure 6-4 shows data recording several “triple-rocking” trials using this approach. In each trial, the robot performs a dynamic lunge, kicks-off at landing to produce a more repeatable forward rocking motion, and then holds all joints in place as the rear legs land again. There is enough angular momentum when the rear legs land that a third rocking motion occurs. Thus, we call this sequence a “triple-rocking” trial. Between sequential trials, there is a brief pause to allow the entire system to come to rest with a known initial condition for the next trial.

Figure 6-5 shows a close-up of the pitch angle for one particular triple-rocking trial, and Figure 6-6 shows the stochastic transition map from nose-up to nose-down apex angle. Note that the second pitch angle (y-axis) is not nearly as correlated to the initial pitch angle




Figure 6-4: Data for metastable bounding, with stabilizing pauses

(x-axis) as it was for the stochastic return map for the stiff-legged rimless wheel landing trial, shown in Figure 6-3. We believe the apex of second pitch angle can be made much more repeatable with additional feedback, but that task is outside the scope of this thesis.

Finally, we note that, as in the case for the rimless-wheel style rocking in Figure 6-2, we again observe near-singularities in the loading direction at the knee which cause kinks in both joint angle and body pitch trajectories in some trials. Improving motion plans to eliminate this unwanted noise source is another suggested task for future work.

6.2.3 Open-loop continuous bounding

Finally, we can replace the third, stiff-legged rocking motion in the triple-rocking motion with another dynamic lunge and create a continuous rocking gait. Here, we simply splice one nose-up/nose-down sequence after another, with careful timing. Figure 6-7 shows data for several seconds of open-loop, continuous bounding. This motion is sensitive to the timing of the push-offs. There is a brief pause before each nose-up dynamic lunge, to ensure the initial pitch of the robot is close to zero. The nose-down motion (supported by the front legs) is simply timed for a nominal impact time, determined experimentally.




Figure 6-5: Zoom view of pitch

Data above show a close-up view of one particular bounding motion. The pitch trajectory shows three distinct apex points: the robot rocks successively in double support onto rear, then front, and then rear feet again. Negative pitch corresponds to a “nose-up” orientation.




Figure 6-6: Apex-to-apex data for initial rocking motion

Because of the stochasticity inherent in the motions, this timing will be significantly off every few steps, and the resulting nose-down pitch will be noticeably too large or too small as a result.




Figure 6-7: Data during continuous bounding

Through careful planning, an open-loop bounding motion can produce a few successful, continuous steps but, it is sensitive to disturbances. The average forward velocity during this motion is about 9.5 cm/sec.

We can typically achieve runs of several bounding steps (5 to 10) using this strategy. We believe the most critical step in stabilizing this motion is the timing of the push-offs for both of the rocking motions (rather than the adjustment of the actual trajectories of the motions). This gait currently has a speed of about 9.5 cm/sec. Adding feedback to improve timing should allow us to remove the pauses before each nose-up dynamic lunge, which would theoretically increase the speed to about 13 cm/sec.

Additional improvements would include adjusting step length and height to avoid hazardous upcoming terrain features, such as gaps and rocks.

6.3 Discussion

This chapter presents initial results toward the development of a continuous bounding gait for LittleDog. Two key concepts in developing such motions are (1) appropriate dimensionality reduction of the high-order dynamic system and (2) modification of the stochastic transition map, to reduce variability in the end state of our dynamic motions. We significantly reduce the dimensionality of the 18-DOF robot in two ways. First, we use a simple planar model for the dynamic lunge (described in Section 3.7.5). Second, the magnitude and timing of the initial pitch apex capture important information about both the energy in the system and the expected time of touch-down by the front feet. We believe predicting the landing time of each rocking motion will prove to be the most important factor in closing the loop to make continuous bounding more robust, and that the time and magnitude of peak pitch can be used to make this prediction accurately.

A significantly more robust continuous bounding motion would likely require more precise feedback timing than permitted by our current control architecture for LittleDog. We are optimistic that such a gait can eventually be designed, however, and may negotiate rough terrain with a significant speed advantage over a crawling walk gait.

Chapter 7

Conclusions and Future Work

Below, we summarize the contributions of this thesis in developing and evaluating reliable control solutions for legged locomotion. We also outline possible future directions for related work and conclude by discussing some particular implications of our results on the future development of highly-dynamic robotics.

7.1 Contributions

An overriding goal throughout this thesis has been to develop principles by which legged (and other highly dynamic) robotics can be advanced, scientifically. As we create robots which interact with and perform in our (human) environments and which mimic our dynamic capabilities, it is natural to request of them the same stability guarantees we ask of ourselves: exceptional performance most of the time, with only occasional failures (falling). We believe the underlying perspectives in this thesis are fundamental and hope other researchers will adopt (and adapt) the stochastic methods presented for evaluation and optimization of robot locomotion.

Below, the major contributions of this thesis are summarized briefly. Results are grouped into three areas: dynamic motion planning for the LittleDog robot, the application of stochastic methods in analyzing walking, and policy optimization for passive-based walking on rough terrain.

7.1.1 Kinodynamic motion planning in an underactuated regime

Our goal in Phase 2 of Learning Locomotion was to explore the dynamic performance capabilities of the LittleDog robot, described in detail in Chapter 3. Dynamic gaits are usually associated with more compliant actuators and dynamics, and one of our main contributions is the demonstration that this high-impedance robot can execute repeatable dynamic motions. Our repeatability is due to careful kinodynamic planning during intentionally underactuated phases of motion.

These underactuated motions fall into two regimes, categorized by the magnitude of pitch (α) which occurs while support is provided by only two legs. In both regimes, we use planar (2D) models of the robot to capture the fundamental dynamics. By appropriately observing physical limits in coefficient of friction, kinematic range, and joint velocities, we can plan motions for the 18-DOF robot by essentially planning in terms of just one or two state variables over time.

Pacing and trot-walking are continuous gaits where $\ddot{\alpha} \approx 0$ and the body does not pitch much. Our trot-walking is generalized to work on rough terrain, providing the staple gait for traversing four of the seven terrain types in Phase 2. Motion planning in this regime is a direct extension of a Preview Control planning method for controlling ZMP, which was presented in [86]. This method reduces our motion planning to a solution for a single degree of freedom over time: the motion of the center of body in x .

In the second regime, $\ddot{\alpha} \neq 0$, and we now plan for a non-zero moment about the support legs over time. By studying the dynamics and geometry of this class of motion, we have identified that angular accelerations are coupled much more strongly to the acceleration in x than to acceleration in z . As a result, motion planning in this regime focuses on planning \ddot{x} over time to achieve a particular apex in pitch angle, α ; we only need to incorporate the effects of \ddot{z} secondarily, to refine our prediction of the entire motion trajectory. Our dynamic lunge is particularly effective in crossing a gap or Jersey barrier. On these two terrains, our speeds are well above any of the other five teams and are over twice the required speed metric set by DARPA for Phase 2.

Finally, initial experiments in developing a continuous bounding gait, which were pre-




Figure 7-1: Dynamic motions for LittleDog.

The planar (2D) model at left captures the dynamics accurately enough to plan desired joint trajectories for the full, 18-DOF robot to execute a “dynamic lunge”, shown at right, with high repeatability. The lefthand image also appears in Figure 3-12 on page 78.

sented in Chapter 6, provide evidence that our motion planning ideas can be developed to negotiate a greater range of terrain; particular ideas for future work are discussed in Section 7.2.1.

7.1.2 Stochastic methods to quantify walking stability

As motivated in Chapter 2, we believe our metrics for walking systems in realworld (noisy) environments should attempt to quantify their reliability in the most direct way possible. We propose the mean first-passage time (MFPT) as a measure to do this and believe its introduction as a metric for walking is a significant contribution in itself.

In Chapter 4, we used two passive walking models to demonstrate tools for the estimation of the MFPT: the rimless wheel and the compass gait biped. Each system can be represented as a Markov process by discretizing the dynamic system in state space, using the post-collision states to define a natural discretization in “time”. Appropriate manipulation of the transition matrix, \mathbf{T} , for the (stochastic) dynamics provides one-shot estimates of MFPTs for all states in the system. We note briefly that this calculation requires taking the inverse of the submatrix of \mathbf{T} which excludes the rows and columns corresponding to any self-absorbing “fallen-down” states, which establishes an inherent numerical limit to the size of the discrete mesh which can be used in modeling. This make these tools most

directly applicable if low-dimensional models can represent the fundamental dynamics of a legged robot; e.g., if the model is already low-dimensional, like the compass gait robot, or if a low-dimensional model captures our dynamics, as was the case for our dynamic motions for LittleDog.

We note that an eigenanalysis of the transition matrix can reveal important *system-wide* dynamics of the system. For realistic walking systems, we find that the first eigenvalue is unity, and its corresponding eigenvector has all weight in the fallen state. This occurs when there is some eventual route (no matter how improbable) from every state to failure (fallen): if it can happen, it eventually will, with probability one. All other eigenvalues have magnitude less than 1, and each has an associated time constant¹, given as $\tau_n = \frac{-1}{\log(\lambda_n)}$.

Metastable limit cycles can exist, however, even if eventual failure is guaranteed. Such metastable walkers can be expected to exhibit extremely long periods of continuous walking before they fall. In such systems, we observe that the magnitude of the second eigenvalue is *very nearly* equal to 1, meaning the associated time constant will be very, very long. The weights in the non-fallen states of the corresponding eigenvector can be rescaled (to sum to unity) to give a PDF of the “metastable neighborhood” in which the dynamics will remain after initial conditions are forgotten.

Finally, we note that the MFPT for a metastable system remains relatively flat across a particular region in state space. This occurs because τ_2 is much larger than τ_3 (or any of the other, smaller time constants). Correspondingly, initial conditions will be forgotten on a time scale which is much smaller than the MFPT. As a result, metastable walking systems have a *system-wide* MFPT².

7.1.3 Policy optimization for CG walker on rough terrain

In Chapter 5, we applied the tools for *quantifying* stochastic stability toward *optimization* of the MFPT of a system. We modeled a controlled version of a compass gait walker on rough terrain. By using a hierarchical control strategy, we could solve for optimal control

¹Note, for systems with cyclical dynamics, one finds some eigenvalues are complex, as mentioned in Section 4.4.3.

²This system-wide MFPT is τ_2 , and for metastable systems, $\tau_2 \approx \frac{1}{1-\lambda_2}$, as noted in Section 4.2.2.

policies using value iteration. Two significant results emerged. First, we demonstrated the capabilities of an underactuated³ compass gait on rough terrain. These results show better performance than has been documented for passive-based walkers, which are generally notoriously sensitive to initial conditions and perturbations.

A second finding was that near-optimal policies could be found using only a single step lookahead on the upcoming terrain. Our results on stochastic terrain for a one-step lookahead far out-performed an optimal policy which knows only the “statistics” of terrain ahead (without knowing the next, particular terrain step height). We also compared a one-step policy to an infinite-lookahead policy on rough terrain which wraps periodically, and we found the one-step policy performed nearly as well as the one which knew all terrain ahead perfectly. We believe this result is coupled to the fact that initial conditions are forgotten so quickly in walking systems, due to the energy dissipation as each new step is taken. Our results are highly suggestive that near-optimal planning for legged robots may often be possible using only a short lookahead of one (or a few) steps on upcoming terrain.

7.2 Suggested directions for future work

This research can be furthered in several directions. Three particular areas are outlined below: further development of dynamic motions for LittleDog, implementation of a realworld version of our actuated compass gait model on rough terrain, and a theoretical comparison of the efficiency of walking versus “rolling” on rough terrain.

7.2.1 LittleDog: Ongoing work

Phase 3 of the DARPA Learning Locomotion program begins as I conclude this thesis. Work during Phase 2 generated fast, dynamic motions which differentiated MIT’s performance from other participating teams and provided benchmark results on specific terrain obstacle types. Significant challenges which remain now include adapting our kinodynamic planning for highly dynamic motions to traverse more generalized terrain and to operate with greater autonomy.

³The toe pivot remains passive in the model.



Figure 7-2: LittleDog lunging onto rough terrain.

To visualize these two challenges, consider Figure 7-2, which shows LittleDog as it executes a dynamic lunge and then carefully places the front feet into pre-planned footholds. Similar shots of LittleDog were also shown in Figure 3-23 on page 95. First, note the dynamic lunge is initiated here from *flat* terrain here. Traversal on the rock-like terrain itself is currently done using less-dynamic motions - trot-walking at best and perhaps only moving at a true crawl. We want to extend our dynamic motions to work on such terrain. Second, reliably climbing *onto* the rock (after the last frame of Figure 7-2) currently depends on having good, solid footholds planned for the front feet, on identifying what style of motion to employ in the first place, and on recovery and/or feedback tactics to avoid toppling once feet slip. All of this depends on a good, high-level planning architecture, to produce greater autonomy.

In addressing the first issue (of extending our repertoire of dynamic motions) there are two obvious directions of research to pursue. First is the development of a continuous, metastable bounding motion, of the type discussed throughout Chapter 6. The second suggested advancement is the generalization of our dynamic lunge to begin and to end with arbitrary foot placements; i.e., to lunge on rough terrain. Initial work toward this second goal was presented in Section 3.8.2⁴. Ideally, these two goals will then (eventually) be combined, to create a continuous bounding gait *on rough terrain*.

Achieving all these goals may depend on a ground-up reorganization of the control structure for LittleDog, which is a part of the second challenge: obtaining greater autonomy. In particular, achieving a reliable bounding cycle will depend on better *timing* of the leg motions than we can currently achieve. At present, short time-length joint trajectories

⁴In particular, see Figure 3-24 on page 96 and Figure 3-25 on page 97.

are planned in MATLAB based on a “snapshot” of the current state of the robot. Then, a high-level control loop sends commands to the PD controller on the robot itself. The feedback loop of motion planning, joint-trajectory regulation and sensor feedback needs to be more tightly entwined to connect the nose-up and nose-down rocking motions described in Chapter 6.

Better autonomous locomotion also requires an improved high-level planner, as previously mentioned. Better quantification of the low-level motion performance (reliability and speed, given particular desired initial and final states) should facilitate better autonomy, making the problems of high-level planning more tractable. Additionally, the high-level planner should identify when *recovery motions* are required, to respond to unexpected terrain collisions or foot slippage.

At present, Alec Shkolnik and Michael Levashov are continuing work at MIT on controlling LittleDog during Phase 3, under the guidance of Russ Tedrake. This short section on future work is of course incompletely, and their achievements will undoubtedly follow directions and achieve goals outside the scope of what has been mentioned above.

7.2.2 Verification of compass gait results on a real robot

Given the success of our simulations, we plan to implement a similar control strategy on a real compass-gait robot. This robot is shown in Figure 7-3 and was designed and built by Stephen Proulx and Mario Bollini at the Robot Locomotion Group. The design is intentionally similar to the idealized model studied in Chapter 5. Specifically, it has a direct-drive motor at the hip and a constant-magnitude (post-collision) toe-off at each sensed ground collision. It is mounted on a boom, providing lateral stability but also introducing some additional, unmodeled dynamics. The same motion capture environment and known-profile terrain boards used for LittleDog will also be used for this work, allowing us to know both the state of the robot and the upcoming, rough terrain profile with good accuracy.

Promising initial work has been done by Fumiya Iida in designing open-loop, limit-cycle walking on this robot, and Ian Manchester performed some initial characterization of the robot dynamics (while visiting in the summer of 2008). We are hopeful that the inherent



Figure 7-3: Compass gait robot posed on rough terrain.

stability demonstrated by the simple low-level control strategy described in this thesis can also be demonstrated on the real-world system. We expect that a controller based on the same practical principles will provide a good initial control policy, which can then be tuned online using model-free, gradient-based optimization methods.

7.2.3 Efficiency of passive walking on rough terrain

One particularly promising characteristic of passive-dynamic based robots is their high efficiency on flat terrain [35], which occurs because some energy is carried over from one step to the next. On flat terrain, however, wheels clearly win over legs in a head-to-head battle in efficiency of locomotion.

We would like to examine the theoretical efficiency for both dynamic legged robot gaits and wheeled transportation on rough terrain. Legged robots can select desired contact points with the ground, while wheeled vehicles lack this flexibility. As terrain becomes more extreme, there surely exists a point at which using legs provides not only greater capability but also greater efficiency in locomotion as well. We have not seen such an analysis in the literature to date and hope to address this question in the near future.

7.3 Implications for development of highly dynamic robots

We conclude by highlighting the implications our results have toward the development dynamic, autonomous robots in the coming years. First, we reiterate that global stability will not typically exist for our walking machines, and that our goal should be to optimize *stochastic stability*.

The success of short-sighted strategies, discussed in Section 5.8, has important implications in kinodynamic planning for legged robots. It means that near-optimal strategies may simply require good low-level control, and that selecting a “greedy” short-term action may often be a good policy for long-term stability. For a robot such as LittleDog, this hints at the potential for compact high-level planning strategies. We can already demonstrate high reliability of short-duration, low-level dynamic motions for LittleDog, and we suspect the proper control framework may allow us to combine such motions while using only a relatively short lookahead.

Finally, we note that most of the work in this thesis applies more generally to a broader class of highly dynamic robots (flying, swimming, etc.) in realworld environments, and that we have presented powerful tools which can be adapted quite naturally for machine learning.

Appendix A

Pendulum Foot Model

This appendix derives the force balance and equations for the instantaneous location of the center of pressure (COP), equivalently called the zero-moment point (ZMP).




Figure A-1: Inverted pendulum (2D) balancing on a weighted “foot” platform
This image is reproduced from Figure 2-1 on page 40.

The position, velocity and acceleration of the mass m in the z direction are given by:

$$z = a \cos \theta$$

$$\dot{z} = -a \sin \theta \dot{\theta}$$

$$\ddot{z} = -a \sin \theta \ddot{\theta} - a \cos \theta \dot{\theta}^2$$

When both F_1 and F_2 are positive in value, so that the foot remains flat on the ground, there is a single equation of motion to describe the pendulum:

$$J\ddot{\theta} = \tau + mg \sin \theta$$

where $J = ma^2$. Solving for angular acceleration,

$$\ddot{\theta} = \frac{\tau}{ma^2} + \frac{g \sin \theta}{L}$$

We now express the simple force balance, $F_z = m(g + \ddot{z})$, in terms of state variables and inputs by substituting for $\ddot{\theta}$:

$$F_z = mg - \frac{\tau}{a} \sin \theta - mg \sin^2 \theta - ma\dot{\theta}^2 \cos \theta$$

$$F_z = mg \cos^2 \theta - \frac{\tau}{a} \sin \theta - ma\dot{\theta}^2 \cos \theta$$

In the zero-torque case ($\tau = 0$) this force, F_z , and the additional force from the static mass, m_f , will each be distributed onto the point contacts as:

$$F_{1,\tau=0} = \frac{b_2 F_z + (b_2 - a_f) m_f g}{b_1 + b_2}$$

$$F_{2,\tau=0} = \frac{b_1 F_z + (b_1 + a_f) m_f g}{b_1 + b_2}$$

To balance the torque source, we must add some force, $k_2 = k$, to F_2 and add an equal but opposite force, $k_1 = -k$, to F_1 to generate a net *torque* without generating a net *force in z*. To balance the additional torque source, we require:

$$\tau = k_2 b_2 - k_1 b_1$$

$$\tau = k b_2 + k b_1$$

$$k = \frac{\tau}{b_1 + b_2}$$

and so:

$$F_1 = \frac{-\tau + b_2 F_z + (b_2 - a_f) m_f g}{b_1 + b_2} \quad (\text{A.1})$$

$$F_2 = \frac{\tau + b_1 F_z + (b_1 + a_f) m_f g}{b_1 + b_2} \quad (\text{A.2})$$

as given in Equations 2.2 and 2.3 on page 40.

Appendix B

Compass Gait Implementation Details

This appendix provides details on the compass gait model simulations on both stochastic and (known) wrapping terrain. Included are equations of motion and details on meshing and on implementation of interpolation for post-collision states in state space.




Figure B-1: Four states of the CG walker during continuous dynamics
The biped model is depicted here immediately after an impact. A pre-collision impulse occurs at each step, and only the torque source at the hip actuates the walker until the next step occurs.

If one considers only the continuous dynamics *between* two steps, the dynamics of the compass gait walker are *identical* to those of the underactuated acrobot. Figure B-1




Figure B-2: Compass gait meshing parameters for stochastic terrain




Figure B-3: Compass gait meshing parameters for wrapping terrain

depicts the state variables and model parameters which define the continuous phase of the dynamics. With the stance leg of the compass gait modeled as a passive pivot joint, the entire state of the biped is defined by the two angles and two angular velocities of the legs. We use the absolute angle of the stance leg, θ_1 , and the relative angle of the swing leg (with respect to the stance leg), θ_2 , along with their time derivatives, $\theta_3 = \dot{\theta}_1$, $\theta_4 = \dot{\theta}_2$. The model includes a point mass m_h at the “hip” and a point mass of magnitude m along each of the (otherwise massless) legs, all at the geometric locations shown.

Meshing details for stochastic terrain modeling

Figure B-2 illustrates the states used in meshing for optimal control on *stochastic* terrain. Only *post-collision* states are used in meshing, meaning both legs are in contact with the ground. For our stochastic terrain model, each upcoming step height is drawn from a Gaussian distribution. The figure illustrates the somewhat artificial nature of the terrain. The entire post-collision state on stochastic terrain is defined by the four *meshing* parameters shown. As compared with the state definitions in Figure B-1, the first two states are now the horizontal and vertical distances between the two feet (rather than the two angles of the legs). The meshing states for stochastic terrain are also defined in the equations below:

$$\begin{aligned}
 X_{m1} &= \Delta z = z_{st} - z_{sw} \\
 X_{m2} &= -\Delta x = x_{sw} - x_{st} \\
 X_{m3} &= \dot{\theta}_1 \\
 X_{m4} &= \dot{\theta}_2
 \end{aligned}$$

Table B.1 gives the meshing discretization used for the stochastic terrain case.

Table B.1: Meshing for compass gait model on stochastic terrain

Parameter	num. of elements	min	max	[units]
X_{m1}	19	-1	.1	[m]
X_{m2}	10	-7	-.16	[m]
X_{m3}	10	-2.1	-1.1	[rad/s]
X_{m4}	10	-1	1.5	[rad/s]
α_{des}	11	15	40	[deg]

The spacing in the mesh for Table B.1 is *not* linear. The exact values are listed below:

$$X_{m1} = [-.1, -.075, -.04, -.03 : .005 : .03, .04, .06, .1]$$

$$X_{m2} = [-.7 : .06 : -.16]$$

$$X_{m3} = [-2.1 : .1 : -1.4, -1.25, -1.1]$$

$$X_{m4} = [-1, -.7, -.5 : .25 : .75, 1.1, 1.5]$$

Meshing details for wrapping terrain modeling

Figure B-3 illustrates the states used in meshing for optimal control on *wrapping* terrain. On wrapping terrain, the terrain profile over a repeating, 7-meter long stretch of terrain is known exactly and is included in the simulation of step-to-step dynamics. Therefore, the entire post-collision state of the walker can be defined using the absolute x position of the (upcoming) stance foot on terrain, the relative x distance to the swing foot (which also touches the terrain instantaneously at collision), and the two angular velocities (given in both Figures B-1 and B-2). Meshing states for wrapping terrain are given below:

$$X_{m1} = x_{st}$$

$$X_{m2} = -\Delta x = x_{sw} - x_{st}$$

$$X_{m3} = \dot{\theta}_1$$

$$X_{m4} = \dot{\theta}_2$$

Table B.2: Meshing for compass gait model on wrapping terrain

Parameter	num. of elements	min	max	[units]
X_{m1}	140	0	7	[m]
X_{m2}	15	-.85	-.15	[m]
X_{m3}	14	-3.0	-.4	[rad/s]
X_{m4}	14	-.1	5.1	[rad/s]
α_{des}	13	10	40	[deg]

Table B.2 gives the meshing discretization used for the stochastic terrain case.

Unlike the stochastic case, the spacing in the mesh for the wrapping case (given in Table B.2) is linear. These simulations were done *before* the simulations on stochastic terrain, and the stochastic mesh was refined based on the results on wrapping terrain (to refine more-frequently visited regions in state space).

The rest of the modeling details given in this appendix apply to *both* the stochastic and wrapping terrain cases.

Swing phase dynamics

During the swing phase of the leg, the dynamics of the system are the so-called “acrobot” dynamics; this is a classic underactuated model which has been studied in depth [159, 14, 18, 161]. We intentionally choose to present the two equations of motion below using a familiar definition of the parameters and states:

$$d_{11}\ddot{\theta}_1 + d_{21}\ddot{\theta}_2 + h_1 + p_1 = 0 \quad (\text{B.1})$$

$$d_{12}\ddot{\theta}_1 + d_{22}\ddot{\theta}_2 + h_2 + p_2 = \tau \quad (\text{B.2})$$

where

$$\begin{aligned}
d_{11} &= m_1 l_{c1}^2 + m_2 (l_1^2 + l_{c2}^2 + 2l_1 l_{c2} \cos \theta_2) + I_1 + I_2 \\
d_{12} &= d_{21} = m_2 (l_{c2}^2 + l_1 l_{c2} \cos \theta_2) + I_2 \\
d_{22} &= m_2 l_{c2}^2 + I_2 \\
h_1 &= -m_2 l_1 l_{c2} \sin \theta_2 \dot{\theta}_2^2 - 2m_2 l_1 l_{c2} \sin \theta_2 \dot{\theta}_2 \dot{\theta}_1 \\
h_2 &= m_2 l_1 l_{c2} \sin \theta_2 \dot{\theta}_1^2 \\
p_1 &= (m_1 l_{c1} + m_2 l_1) g \cos \theta_1 + m_2 l_{c2} g \cos (\theta_1 + \theta_2) \\
p_2 &= m_2 l_{c2} g \cos (\theta_1 + \theta_2)
\end{aligned}$$

and

$$\begin{aligned}
m_1 &= m_2 = m + \frac{1}{2} m_h \\
l_1 &= l_2 = a + b \\
l_{c1} &= L - \frac{bm}{m_1} \\
l_{c2} &= L - l_{c1} \\
I_1 &= I_2 = m (b - l_{c2})^2 + \frac{1}{2} m_h l_{c2}^2
\end{aligned}$$

Solving explicitly for the time derivatives of the state variables, we obtain:

$$\ddot{\theta}_1 = \frac{d_{22}\eta_1 - d_{12}\eta_2}{d_{11} * d_{22} - d_{12}^2} \tag{B.3}$$

$$\ddot{\theta}_2 = \frac{-d_{12}\eta_1 + d_{11}\eta_2}{d_{11} * d_{22} - d_{12}^2} \tag{B.4}$$

where

$$\begin{aligned}
\eta_1 &= m_2 l_1 l_{c2} \sin \theta_2 \dot{\theta}_2^2 + 2m_2 l_1 l_{c2} \sin \theta_2 \dot{\theta}_2 \dot{\theta}_1 + \dots \\
&\quad m_2 l_{c2} g \cos (\theta_1 + \theta_2) + (m_1 l_{c1} + m_2 l_1) g \cos \theta_1 \\
\eta_2 &= -m_2 l_1 l_{c2} \sin \theta_2 \dot{\theta}_1^2 - m_2 l_{c2} g \cos (\theta_1 + \theta_2) + \tau
\end{aligned}$$

Collision dynamics

We assume collisions are instantaneous and inelastic. For convenience and for the reader to compare with typical notation, e.g. [58], we will define absolute angular velocities of the legs:

$$\dot{\theta}_{ns} = \dot{\theta}_3 + \dot{\theta}_4 \quad (\text{B.5})$$

$$\dot{\theta}_s = \dot{\theta}_3 \quad (\text{B.6})$$

Geometry and conservation of angular momentum yield the following relationships between the pre-collision and post-collision states, which are given the superscripts $-$ and $+$, respectively.

$$\theta_1^+ = \theta_1^- + \theta_2^- - \pi \quad (\text{B.7})$$

$$\theta_2^+ = 2\pi - \theta_2^- \quad (\text{B.8})$$

$$Q_{11}^- \dot{\theta}_{ns}^- + Q_{12}^- \dot{\theta}_s^- = Q_{11}^+ \dot{\theta}_{ns}^- + Q_{12}^+ \dot{\theta}_s^- \quad (\text{B.9})$$

$$Q_{21}^- \dot{\theta}_{ns}^- + Q_{22}^- \dot{\theta}_s^- = Q_{21}^+ \dot{\theta}_{ns}^- + Q_{22}^+ \dot{\theta}_s^- \quad (\text{B.10})$$

where the last two equations can be expanded immediately, as given below. We also present an expanded form below, which has been useful for (vectorized) MATLAB implementation; many collision relationships can be calculated at once by using expanded forms of the matrix inverse and other algebraic expressions. Eqn. B.9, which represents the angular momentum balance about the swing toe immediately before collision, will be modified later to account for pre-collision impulse later in the appendix.

$$Q_{11}^- (\dot{\theta}_1^- + \dot{\theta}_2^-) + Q_{12}^- \dot{\theta}_1^- = Q_{11}^+ (\dot{\theta}_1^+ + \dot{\theta}_2^+) + Q_{12}^+ \dot{\theta}_1^+ \quad (\text{B.11})$$

$$Q_{21}^- (\dot{\theta}_1^- + \dot{\theta}_2^-) + Q_{22}^- \dot{\theta}_1^- = Q_{21}^+ (\dot{\theta}_1^+ + \dot{\theta}_2^+) + Q_{22}^+ \dot{\theta}_1^+ \quad (\text{B.12})$$

where

$$\begin{aligned}
Q_{11}^- &= -mab \\
Q_{12}^- &= -mab + (m_h L^2 + 2maL) \cos(2\alpha) \\
Q_{21}^- &= 0 \\
Q_{22}^- &= -mab \\
Q_{11}^+ &= mb(b - L \cos(2\alpha)) \\
Q_{12}^+ &= mL(L - b \cos(2\alpha)) + ma^2 + m_h L^2 \\
Q_{21}^+ &= mb^2 \\
Q_{22}^+ &= -mbL \cos(2\alpha)
\end{aligned}$$

where $2\alpha = \pi - \theta_2^-$ is the interleg angle. We can solve for the two new angular velocities by using the matrix inverse of \mathbf{Q}^+ , $\mathbf{R}^+ = (\mathbf{Q}^+)^{-1}$, which is simply:

$$\begin{aligned}
R_{11}^+ &= \frac{Q_{22}^+}{Q_{11}^+ Q_{22}^+ - Q_{12}^+ Q_{21}^+} \\
R_{12}^+ &= \frac{-Q_{12}^+}{Q_{11}^+ Q_{22}^+ - Q_{12}^+ Q_{21}^+} \\
R_{21}^+ &= \frac{-Q_{21}^+}{Q_{11}^+ Q_{22}^+ - Q_{12}^+ Q_{21}^+} \\
R_{22}^+ &= \frac{Q_{11}^+}{Q_{11}^+ Q_{22}^+ - Q_{12}^+ Q_{21}^+}
\end{aligned}$$

Now let $\mathbf{A} = \mathbf{R}^+ \mathbf{Q}^-$:

$$\begin{aligned}
A_{11} &= R_{11}^+ Q_{11}^- + R_{12}^+ Q_{21}^- \\
A_{12} &= R_{11}^+ Q_{12}^- + R_{12}^+ Q_{22}^- \\
A_{21} &= R_{21}^+ Q_{11}^- + R_{22}^+ Q_{21}^- \\
A_{22} &= R_{21}^+ Q_{12}^- + R_{22}^+ Q_{22}^-
\end{aligned}$$

The matrix \mathbf{A} now relates the *absolute* angular velocities of the two legs. That is, the matrix

A defines the following two relationships:

$$\begin{aligned} A_{11} \begin{pmatrix} \dot{\theta}_s^- \\ \dot{\theta}_{ns}^- \end{pmatrix} + A_{12} \begin{pmatrix} \dot{\theta}_s^- \end{pmatrix} &= \dot{\theta}_s^+ + \dot{\theta}_{ns}^+ \\ A_{21} \begin{pmatrix} \dot{\theta}_s^- \\ \dot{\theta}_{ns}^- \end{pmatrix} + A_{22} \begin{pmatrix} \dot{\theta}_s^- \end{pmatrix} &= \dot{\theta}_s^+ \end{aligned}$$

We can simply rearrange these relationship to find a direct relationship given $\dot{\theta}_2 = \dot{\theta}_{ns} - \dot{\theta}_s$ is actually a *relative* velocity. Note that these relationship have also accounted for the fact that the two legs “swap roles” during impact (with stance becoming swing leg, and vice versa). We will call this final transformation matrix **B**:

$$\begin{aligned} B_{11} &= A_{21} + A_{22} \\ B_{12} &= A_{21} \\ B_{21} &= A_{11} + A_{12} - A_{21} - A_{22} \\ B_{22} &= A_{11} - A_{21} \end{aligned}$$

This final relationships relating pre- and post-collision velocities are then:

$$\dot{\theta}_1^+ = B_{11}\dot{\theta}_1^- + B_{12}\dot{\theta}_2^- \quad (\text{B.13})$$

$$\dot{\theta}_2^+ = B_{21}\dot{\theta}_1^- + B_{22}\dot{\theta}_2^- \quad (\text{B.14})$$

PD controller

During a step, the interleg angle is regulated with respect to a desired set-point, which is selected as a high-level control action, once per step.

$$\alpha = \theta_{sw} - \theta_{st} \quad (\text{B.15})$$

$$\tau = K_p(\alpha_{des} - \alpha) + K_d(0 - \dot{\alpha}), \quad \theta_{st} < 0 \quad (\text{B.16})$$

where $K_p = 100$ and $K_d = 10$. These two equations appear in Chapter 5 as Equations 5.5 and 5.6.

Pre-collision Impulse

Our control was much more successful if we used a pre-collision impulse in addition to a PD-controlled torque at the hip, as described in Section 5.6.2 (p. 144). We assume the impulse occurs instantaneous and happens exactly before the step-to-step collision takes place.

If the imparted impulse is very large, the entire compass gait walker will become airborne. Calculation of new state in this “ballistic” case simply involves solving two equations for conservation of angular momentum and another two to conserve linear momentum. For the magnitudes of impulse we use in simulation, there is typically not enough momentum added for the upcoming stance leg to lose contact with the ground. In this “non-ballistic” case, we simply assume that momentum directed toward the ground is perfectly absorbed in an instantaneous collision.

For the case where the upcoming stance foot does not become ballistic, we simply add the effect of the impulse as an additional term in the equations of conservation of angular momentum for the pre-to-post collision dynamics. Specifically, we rewrite Equation B.11 to include a term which scales the impulse to give its contribution to the angular momentum about the pre-collision swing toe. Defining our impulse as having magnitude p , directed axially from the toe to the hip of the pre-collision stance leg, Equation B.11 then becomes:

$$Q_{11}^- \dot{\theta}_{ns}^- + Q_{12}^- \dot{\theta}_s^- - 2p (\sin \alpha) (\cos \alpha) = Q_{11}^+ \dot{\theta}_{ns}^- + Q_{12}^+ \dot{\theta}_s^- \quad (\text{B.17})$$

Propagating these terms and solving algebraically, Equations B.13 and B.14 are now replaced by the following:

$$\dot{\theta}_1^+ = B_{11} \dot{\theta}_1^- + B_{12} \dot{\theta}_2^- - R_{21} (2 \cos \alpha \sin \alpha) p \quad (\text{B.18})$$

$$\dot{\theta}_2^+ = B_{21} \dot{\theta}_1^- + B_{22} \dot{\theta}_2^- - (R_{11} - R_{21}) (2 \cos \alpha \sin \alpha) p \quad (\text{B.19})$$

Bibliography

- [1] Yeuhi Abe and Jovan Popovic. Interactive animation of dynamic manipulation. Technical report, MIT CSAIL, feb 2006.
- [2] Abhishek Agrawal and Sunil K Agrawal. An approach to identify joint motions for dynamically stable walking. *Journal of Mechanical Design*, 128(3):649–653, May 2006.
- [3] R McN Alexander. Bipedal animals, and their differences from humans. *J. Anat.*, 204:321–330, 2004.
- [4] R. Altendorfer, N. Moore, H. Komsuoglu, M. Buehler, H.B. Brown Jr., D. McMordie, U. Saranli, R. Full, and D.E. Koditschek. RHex: A biologically inspired hexapod runner. *Autonomous Robots*, 11(3):207–213, Nov 2001.
- [5] R. Altendorfer, U. Saranli, H Komsoglu, D. Koditschek, H. B. Brown, M. Buehler, N. Moore, D. McMordie, and R. Full. Evidence for spring loaded inverted pendulum running in a hexapod robot. In *Proceedings of the 7th International Symposium on Experimental Robotics (ISER)*, 2000.
- [6] Frank C. Anderson and Marcus G. Pandy. Dynamic optimization of human walking. *Journal of Biomechanical Engineering (AMSE)*, 123(5):381–390, October 2001.
- [7] S. O. Anderson, M. Wisse, C. G. Atkeson, J. K. Hodgins, G. J. Zeglin, and B. Moyer. Powered bipeds based on passive dynamic principles. *Proc. IEEE Humanoids*, pages 1–7, 2005.

- [8] Shinya Aoi and Kazuo Tsuchiya. Self-stability of a simple walking model driven by a rhythmic signal. *Nonlinear Dynamics*, 48(1–2):1–16, Apr 2007.
- [9] F. Asano, M. Yamakita, N. Kamamichi, and Luo Zhi-Wei. A novel gait generation for biped walking robots based on mechanical energy constraint. *IEEE Transactions of Robotics and Automation*, 20(3):565–573, June 2004.
- [10] Alan Asbeck, Sangbae Kim, M. R. Cutkosky, W. R. Provancher, and M. Lanzetta. Scaling hard vertical surfaces with compliant microspine arrays. *Int. Journal of Robotics Research*, 25(12):1165–1179, 2006.
- [11] S K Au. Probabilistic failure analysis by importance sampling markov chain simulation. *Journal of Engineering Mechanics*, pages 303–311, Mar 2004.
- [12] S K Au. Reliability-based design sensitivity by efficient simulation. *Computers and Structures*, 83(14):1048–1061, 2005.
- [13] Joaquin Blaya and Hugh Herr. Adaptive control of a variable-impedance ankle-foot orthosis to assist drop-foot gait. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 12(1):24–31, Mar 2004.
- [14] Gary Boone. Minimum-time control of the acrobot. volume 4, pages 3281–3287. IEEE International Conference on Robotics and Automation (ICRA), 1997.
- [15] Anton Bovier. Metastability and ageing in stochastic dynamics. In A. Maas, S. Martinez, and J. San Martin, editors, *Dynamics and Randomness II*, pages 17–81. Kluwer, 2004.
- [16] Anton Bovier, Michael Eckhoff, Veronique Gayraud, and Markus Klein. Metastability and small eigenvalues in markov chains. *J. Phys. A: Math. Gen.*, 33(46):L447–L451, Nov 2000.
- [17] S. Boyd, P. Diaconis, and L. Xiao. Fastest mixing Markov chain on a graph. *SIAM Review*, 46(4):667–689, Dec 2004.

- [18] Scott C. Brown and Kevin M. Passino. Intelligent control for an acrobat. *Journal of Intelligent and Robotic Systems*, 18:209–248, 1997.
- [19] Brett Browning, Jeremy Searock, Paul E. Rybski, and Manuela Veloso. Turning segways into soccer robots. *Industrial Robot*, 32(2):149–156, 2005.
- [20] M. Buehler and D.E. Koditschek. From stable to chaotic juggling: Theory, simulation, and experiments. In *Proceedings of the IEEE International Conference on Robotics and Automation*, volume 3, pages 1976–1981. IEEE, May 1990.
- [21] M. Buehler, R. Playter, and M. Raibert. Robots step outside. *Int. Symp. Adaptive Motion of Animals and Machines (AMAM)*, Ilmenau, Germany, pages 1–4, Sep 2005.
- [22] Stephen P. Buerger and Neville Hogan. Complementary stability and loop shaping for improved human-robot interaction. *IEEE Transactions on Robotics*, 23(2):232–244, Apr 2007.
- [23] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek. Sequential composition of dynamically dexterous robot behaviors. *International Journal of Robotics Research*, 18(6):534–555, June 1999.
- [24] Katie Byl, Alec Shkolnik, Sam Prentice, Nick Roy, and Russ Tedrake. Reliable dynamic motions for a stiff quadruped. In *Proceedings of the 11th International Symposium on Experimental Robotics (ISER)*, 2008.
- [25] Katie Byl and Russ Tedrake. Stability of passive dynamic walking on uneven terrain. In Art Kuo, editor, *Proceedings of Dynamic Walking 2006*, May 2006.
- [26] Katie Byl and Russ Tedrake. Approximate optimal control of the compass gait on rough terrain. In *Proc. IEEE International Conference on Robotics and Automation (ICRA)*, 2008.
- [27] Katie Byl and Russ Tedrake. Metastable walking on stochastically rough terrain. In *Robotics: Science and Systems*, 2008.

- [28] John Canny, Ashutosh Rege, and John Reif. An exact algorithm for kinodynamic planning in the plane. In *SCG '90: Proceedings of the sixth annual symposium on Computational geometry*, pages 271–280. ACM, 1990.
- [29] John F. Canny, Bruce Randall Donald, John H. Reif, and Patrick G. Xavier. On the complexity of kinodynamic planning. In *29th Annual Symposium on Foundations of Computer Science*, pages 306–316. IEEE, 1988.
- [30] Y Choi, B-J You, and S-R Oh. On the stability of indirect ZMP controller for biped robot systems. *Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems.*, pages 1966–1971, Oct 2004.
- [31] Michael J. Coleman. *A Stability-Study of a Three-Dimensional Passive-Dynamic Model of Human Gait*. PhD thesis, Cornell University, 1998.
- [32] Michael J. Coleman and Andy Ruina. Motions of a rimless spoked wheel: a simple 2d systems with impacts. August 2002.
- [33] Steven H. Collins, Andy Ruina, Russ Tedrake, and Martijn Wisse. Efficient bipedal robots based on passive-dynamic walkers. *Science*, 307:1082–1085, February 18 2005.
- [34] Steven H. Collins, Martijn Wisse, and Andy Ruina. A three-dimensional passive-dynamic walking robot with two legs and knees. *International Journal of Robotics Research*, 20(7):607–615, July 2001.
- [35] Steven Hartley Collins and Andy Ruina. A bipedal walking robot with efficient and human-like gait. *Proc. IEEE International Conference on Robotics and Automation, Barcelona, Spain*, Apr 2005.
- [36] A F Cordero, H J F M Koopman, and F C T van der Helm. Mechanical model of the recovery from stumbling. *Biol. Cybern.*, 91:212–220, 2004.
- [37] Harry Dankowicz and Petri T. Piiroinen. Exploiting discontinuities for stabilization of recurrent motions. *Dynamical Systems*, 17(4):317–342, 2002.

- [38] Gonzalez de santos, Pablo; Garcia, Elena; Estremera, and Joaquin. *Quadrupedal Locomotion: An Introduction to the Control of Four-legged Robots*. Springer-Verlag New York Inc, 7 2006.
- [39] A. P. Dempster. A generalization of bayesian inference. *J. of the Royal Statistical Soc. Series B (Methodology)*, 30(2):205–247, 1968.
- [40] Bruce Donald, Patrick Xavier, John Canny, and John Reif. Kinodynamic motion planning. *J. of the ACM*, 40(5):1048–1066, 1993.
- [41] Bruce R. Donald and Patrick G. Xavier. Provably good approximation algorithms for optimal kinodynamic planning for Cartesian robots and open chain manipulators. In *SCG '90: Proceedings of the sixth annual symposium on Computational geometry*, pages 290–300. ACM, 1990.
- [42] Finale Doshi, Emma Brunskill, Alec Shkolnik, Thomas Kollar, Khash Rohanimanesh, Russ Tedrake, and Nicholas Roy. A supervised learning approach for collision detection in legged locomotion. In *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007.
- [43] J J Eng, D A Winter, and A E Patla. Intralimb dynamics simplify reactive control strategies during locomotion. *Journal of Biomechanics*, 30(6):581–588, Jun 1997.
- [44] Tom Erez and William D. Smart. Bipedal walking on rough terrain using manifold control. *Proc. International Conference on Intelligent Robots and Systems (IROS)*, pages 1539–1544, 2007.
- [45] M. Fujita. How to make an autonomous robot as a partner with humans: design approach versus emergent approach. *Philosophical Transactions of the Royal Society A: Mathematical, physical, and engineering sciences*, 365(1850):21–47, Jan 2007.
- [46] Yasuhiro Fukuoka, Hiroshi Kimura, and Avis H. Cohen. Adaptive dynamic walking of a quadruped robot on irregular terrain based on biological concepts. *International Journal of Robotics Research*, 22(3–4):187–202, March-April 2003.

- [47] R J Full, T Kubow, J Schmitt, P Holmes, and D Koditschek. Quantifying dynamic stability and maneuverability in legged locomotion. *Integ. and Comp. Biol.*, 42:149–157, 2002.
- [48] Elena Garcia and Pablo Gonzalez de Santos. An improved energy stability margin for walking machines subject to dynamic effects. *Robotica*, 23:13–20, 2005.
- [49] Elena Garcia, Joaquin Estremera, and Pablo Gonzalez de Santos. A classification of stability margins for walking robots. In *Proceedings of the International Conference on Climbing and Walking Robots (CLAWAR)*, 2002.
- [50] M Garcia, A Chatterjee, A Ruina, and M Coleman. The simplest walking model: Stability, complexity, and scaling. *Journal of Biomechanical Engineering – Transactions of the ASME*, 120(2):281–288, Apr 1998.
- [51] C.W. Gardiner. *Handbook of Stochastic Methods for Physics, Chemistry and the Natural Sciences*. Springer-Verlag, third edition, 2004.
- [52] Bernard Gaveau and L. S. Schulman. Theory of nonequilibrium first-order phase transitions for stochastic dynamics. *Journal of Mathematical Physics*, 39(3):1517–1533, Mar 1998.
- [53] L. Geppert. QRIO, the robot that could. *IEEE Spectrum*, 41(5):34–37, May 2004.
- [54] Kenneth Y. Goldberg and Matthew T. Mason. Bayesian grasping. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1264–1269. IEEE, 1990.
- [55] A. Goswami. Postural stability of biped robots and the foot rotation indicator (FRI) point. *International Journal of Robotics Research*, 18(6), 1999.
- [56] A. Goswami, B. Espiau, and A. Keramane. Limit cycles and their stability in a passive bipedal gait. pages 246–251. IEEE International Conference on Robotics and Automation (ICRA), 1996.

- [57] Ambarish Goswami and Vinutha Kallem. Rate of change of angular momentum and balance maintenance of biped robots. *Proceedings of the 2004 IEEE International Conference on Robotics and Automation*, 4:3785–3790, Apr 2004.
- [58] Ambarish Goswami, Benoit Thuilot, and Bernard Espiau. Compass-like biped robot part I : Stability and bifurcation of passive gaits. Technical Report RR-2996, INRIA, October 1996.
- [59] Herve Hacot. Analysis and traction control of a rocker-bogie planetary rover. Master’s thesis, MIT, June 1998.
- [60] Peter Hanggi, Peter Talkner, and Michal Borkovec. Reaction-rate theory: fifty years after kramers. *Reviews of Modern Physics*, 62(2):251–342, Apr 1990.
- [61] K Harada, S Kajita, K Kaneko, and H Hirukawa. An analytical method on real-time gait planning for a humanoid robot. *4th IEEE/RAS International Conference on Humanoid Robots*, 2:640–655, Nov 2004.
- [62] Kale Harbick and Gaurav S. Sukhatme. Robustness experiments for a planar hopping control system. *International Conference on Climbing and Walking Robots*, Sep 2002.
- [63] G C Haynes and A A Rizzi. Gaits and gait transitions for legged robots. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pages 1117–1122, May 2006.
- [64] Eishi Hirasaki, Naomichi Ogihara, Yuzuru Hamada, Hiroo Kumakura, and Masato Nakatsukasa. Do highly trained monkeys walk like humans? A kinematic study of bipedal locomotion in bipedally trained japanese macaques. *Journal of Human Evolution*, 46(6):739–750, June 2004.
- [65] Masato Hirose and Kenichi Ogawa. Honda humanoid robots development. *Philosophical Transactions of the Royal Society*, 365(1850):11–19, Jan 2007.
- [66] Shigeo Hirose and Keisuke Kato. Study on quadruped walking robot in tokyo institute of technology. In *Proceedings of the 2000 IEEE International Conference on Robotics and Automation*, pages 414–419. IEEE, April 2000.

- [67] Shigeo Hirose, Hideyuji Tsukagoshi, and Kan Yoneda. Normalized energy stability margin and its contour of walking vehicles on rough terrain. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, volume 1, pages 181–186, 2001.
- [68] I Hiroshi and K Masayoshi. Adaptive gait for a quadruped robot on 3d path planning. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pages 2049–2054, Sep 2003.
- [69] H Hirukawa, S Hattori, K Harada, S Kajita, K Kaneko, F Kanehiro, K Fujiwara, and M Morisawa. A universal stability criterion of the foot contact of legged robots - Adios ZMP. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pages 1976–1983, May 2006.
- [70] Hirohisa Hirukawa. Humanoid robotics projects in Japan. In *Proceedings of the Understanding Humanoid Robots Workshop at IEEE ICRA 2006*, 2006.
- [71] Jessica Hodgins and Marc Raibert. Adjusting step length for rough terrain locomotion. *IEEE Transactions on Robotics and Automation*, 7(3):289–298, June 1991.
- [72] Jessica Hodgins and Marc H. Raibert. Biped gymnastics. pages 5–13. 4th International Symposium on Robotics Research, MIT Press, 1987.
- [73] Neville Hogan. Impedance control: An approach to manipulation. part i - theory. *Journal of Dynamic Systems, Measurement and Control*, 107:1–7, Mar 1985.
- [74] John M. Hollerbach. Dynamic scaling of manipulator trajectories. Technical Report 700, MIT Artificial Intelligence Lab, Jan 1983.
- [75] John M. Hollerbach, Ian W. Hunter, and John Ballantyne. A comparative analysis of actuator technologies for robotics. In O. Khatib, J. J. Craig, and T. Lozano-Peréz, editors, *The robotics review 2*, pages 299–342. MIT Press, 1992.
- [76] I Horenko, E Dittmer, A Fischer, and C Schutte. Automated model reduction for complex systems exhibiting metastability. *Multiscale Model. Simul.*, 5(3):802–827, 2006.

- [77] E T Hsiao and S N Robinovitch. Biomechanical influences on balance recovery by stepping. *Journal of Biomechanics*, 32(10):1099–1106, Oct 1999.
- [78] Vanessa Hsu. Passive dynamic walking with knees: A point-foot model. Master’s thesis, Massachusetts Institute of Technology, February 2007.
- [79] WW. Huisinga, S. Meyn, and C. Schutte. Phase transitions and metastability in Markovian and molecular systems. *Annals of Applied Probability*, 14(1):419–458, 2004.
- [80] Karl Iagnemma and Steven Dubowsky. Traction control of wheeled robotics vehicles in rough terrain with application to planetary rovers. *The Int. Jour. of Robotics Research*, 23(10–11):1029–1040, Oct–Nov 2004.
- [81] F Iida, G.J. Gomez, and R. Pfeifer. Exploiting body dynamics for controlling a running quadruped robot. In *Proceedings of the 12th International Conference on Advanced Robotics (ICAR)*, pages 229–235, July 2005.
- [82] A. J. Ijspeert, A. Crespi, and D. Ryczko and J. M. Cabelguen. From swimming to walking with a salamander robot driven by a spinal cord model. *Science*, 315(5817):1416–1420, Mar 2007.
- [83] Y. Ikemata, A. Sano, and H. Fujimoto. Analysis of stable limit cycle in passive walking. In *SICE Annual Conference in Fukui*, pages 117–122, Aug 2003.
- [84] S Jain and R K Jain. Problems of estimating meantime system failure. *Microelectronics and reliability*, 34(11):1755–1760, Nov 1994.
- [85] Kajita, S. Kanehiro, F. Kaneko, K. Fujiwara, K. Harada, K. Yokoi, K. Hirukawa, and H. Resolved momentum control: humanoid motion planning based on the linear and angular momentum. *Intelligent Robots and Systems (IROS), Proceedings*, 2003.
- [86] S. Kajita, F. Kanehiro, K. Kaneko, K. Fujiware, K. Harada, K. Yokoi, and H. Hirukawa. Biped walking pattern generation by using preview control of zero-moment point. In *ICRA IEEE International Conference on Robotics and Automation*, pages 1620–1626. IEEE, Sep 2003.

- [87] S. Kajita, K. Kaneko, M. Morisawa, S. Nakaoka, and H. Hirukawa. ZMP-based biped running enhanced by toe springs. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3963–3969, Apr 2007.
- [88] S. Kajita, T. Nagasaki, K. Kaneko, and H. Hirukawa. Zmp-based biped running control. *IEEE Robotics and Automation Magazine*, 14(2):63–72, 2007.
- [89] Hiroshi Kaminaga, Taichi Yamamoto, Junya Ono, and Yosihiko Nakamura. Back-drivable miniature hydrostatic transmission for actuation of anthropomorphic robot hands. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, 2007.
- [90] N.G. Van Kampen. *Stochastic Processes in Physics and Chemistry*. Elsevier, third edition, 2007.
- [91] Tohru Katayama, Takahira Ohki, Toshio Inoue, and Tomoyuki Kato. Design of an optimal controller for a discrete-time system subject to previewable demand. *International Journal of Control*, 41(3):677–699, 1985.
- [92] O. Khatib, K. Yokoi, O. Brock, K. Chang, and A. Casal. Robots in human environments: Basic autonomous capabilities. *Int. Journal of Robotics Research*, 18(7):684–696, Jul 1999.
- [93] Oussama Khatib. A unified approach for motion and force control of robot manipulators: The operational space formulation. *IEEE Journal of Robotics and Automation*, 3(1):43–53, February 1987.
- [94] Sangbae Kim, M. Spenko, S. Trujillo, B. Heyneman, V. Mattoli, and M. R. Cutkosky. Whole body adhesion: hierarchical, directional and distributed control of adhesive forces for a climbing robot. *Robotics and Automation, 2007 IEEE International Conference on*, pages 1268–1273, April 2007.
- [95] Seungsu Kim, ChangHwan Kim, and Jong Hyeon Park. Human-like arm motion generation for humanoid robots using motion capture database. In *Proceedings of the*

IEEE International Conference on Robotics and Automation (ICRA), pages 3486–3491, Oct 2006.

- [96] Hiroshi Kimura, Yasuhiro Fukuoka, Yoshiro Hada, and Kunikatsu Takase. Adaptive dynamic walking of a quadruped robot on irregular terrain using a neural system model. pages 88–97. *International Symposium on Robotics Research*, 2001.
- [97] Daniel E. Koditschek and Martin Buehler. Analysis of a simplified hopping robot. *International Journal of Robotics Research*, 10(6):587–605, Dec 1991.
- [98] J. Zico Kolter and Andrew Ng. Learning omnidirectional path following using dimensionality reduction. In *Proceedings of Robotics: Science and Systems 2007 (RSS)*, 2007.
- [99] A. D. Kuo. Choosing your steps carefully. *Robotics and Automation Magazine, IEEE*, 14(2):18–29, 2007.
- [100] A D Kuo, J M Donelan, and A Ruina. Energetic consequences of walking like an inverted pendulum: Step-to-step transitions. *Exerc. Sport Sci. Rev.*, 33(2):88–97, 2005.
- [101] Arthur D. Kuo. Stabilization of lateral motion in passive dynamic walking. *The International journal of robotics research*, 18(9):917–930, Sep 1999.
- [102] Arthur D. Kuo. Energetics of actively powered locomotion using the simplest walking model. *Journal of Biomechanical Engineering*, 124:113–120, 2002.
- [103] Ryo Kurazume, Kan Yoneda, and Shigeo Hirose. Feedforward and feedback dynamic trot gait control for quadruped walking vehicle. *Autonomous Robots*, 12(2):157–172, Mar 2002.
- [104] F Lacquaniti, R Grasso, and M Zago. Motor patterns in walking. *News Physiol Sci*, 14:168–174, Aug 1999.
- [105] H. Larralde and F. Leyvraz. Metastability for markov processes with detailed balance. *Phys. Rev. Lett.*, 94(16), Apr 2005.

- [106] Joseph F. Laszlo, M. Van De Panne, and E. Fiume. Limit cycle control and its application to the animation of balancing and walking. pages 155–162. SIGGRAPH, 1996.
- [107] LaValle, S.M., Kuffner, J.J., and Jr. Randomized kinodynamic planning. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, 1:473–479, 1999.
- [108] S. LaValle and J. Kuffner. Rapidly-exploring random trees: Progress and prospects. In *Proceedings of the Workshop on the Algorithmic Foundations of Robotics*, 2000.
- [109] S. M. LaValle and M. B. Egerstedt. On time: Clocks, chronometers, and open-loop control. In *In Proceedings IEEE Conference Decision and Control, to appear*, 2007.
- [110] Steven M. LaValle, James J. Kuffner, and Jr. Randomized kinodynamic planning. *The International Journal of Robotics Research*, 20(5):378–400, 2001.
- [111] W Lee and M Raibert. Control of hoof rolling in an articulated leg. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pages 1386–1391, Apr 1991.
- [112] J. Liu, Y. Y. Tang, and O. Khatib. Modeling and learning robot manipulation strategies. In *Proceedings of the 5th International Symposium on Experimental Robotics (ISER)*, 1997.
- [113] Ching long Shih. Analysis of the dynamics of a biped robot with seven degrees of freedom. pages 3008–3013. IEEE International Conference on Robotics and Automation (ICRA), 1996.
- [114] Tad McGeer. Passive dynamic walking. *International Journal of Robotics Research*, 9(2):62–82, April 1990.
- [115] Tad McGeer. Passive walking with knees. pages 1640–1645. IEEE International Conference on Robotics and Automation (ICRA), 1990.
- [116] R.B. McGhee and A.A. Frank. On the stability properties of quadruped creeping gaits. *Mathematical Biosciences*, 3:331–351, 1968.

- [117] T.A. McMahon. Mechanics of locomotion. *The International Journal of Robotics Research*, 3(2), Summer 1984.
- [118] Thomas A. McMahon. *Muscles, Reflexes, and Locomotion*. Princeton University Press, May 1984.
- [119] D. A. Messuri and C. A. Klein. Automatic body regulation for maintaining stability of a legged vehicle during rough-terrain locomotion. *IEEE Journal of Robotics and Automation*, RA-1(3):132–141, Sep 1985.
- [120] N. Metropolis. The beginnng of the Monte Carlo method. *Los Alamos Science*, (15):125–130, 1987.
- [121] Nicholas Metropolis and N. Ulam. The Monte Carlo method. In *J. of the Amer. Statistical Assoc.* [120], pages 335–341.
- [122] Shin min Song and Kenneth J. Waldron. *Machines That Walk : The Adaptive Suspension Vehicle*. MIT Press, November 1988.
- [123] K. Mitobe, G. Capi, H. Takayama, M. Yamano, and Y. Nasu. A ZMP analysis of the passive walking machines. In *SICE Annual Conference*, volume 1, pages 112–116. IEEE, Aug 2003.
- [124] Seiichi Miyakoshi and Gordon Cheng. Examining human walking characteristics with a telescopic compass-like biped walker model. In *Proceedings of the IEEE International Conference on Systems, Man and Cybernetics (SMC2004)*, pages 1538–1543, Oct 2004.
- [125] Simon Mochon and Thomas A. McMahon. Ballistic walking. *Journal of Biomechanics*, 13:49–57, 1980.
- [126] Mitsuhashi Morisawa, Shuuji Kajita, Kensuke Harada, Kiyoshi Fujiwara, Fumio Kanehiro, Kenji Kaneko, and Hirohisa Hirukawa. Emergency stop algorithm for walking humanoid robots. In *Proceedings of the 2005 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE/RSJ, 2005.

- [127] Reinhard Muller, Peter Talkner, and Peter Reimann. Rates and mean first passage times. *Physica A*, 247(247):338–356, Jun 1997.
- [128] Remi Munos and Andrew Moore. Barycentric interpolators for continuous space and time reinforcement learning. In M. S. Kearns, S. A. Solla, and D. A. Cohn, editors, *Advances in Neural Information Processing Systems*, volume 11, pages 1024–1030. NIPS, MIT Press, 1998.
- [129] Remi Munos and Andrew Moore. Variable resolution discretization in optimal control. *Machine Learning*, 49(2/3):291–323, November/December 2002.
- [130] S. Nakaoka, A. Nakazawa, K. Yokoi, H. Hirukawa, and K. Ikeuchi. Generating whole body motions for a biped humanoid for captured human dances. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3905–3910, 2003.
- [131] Gabriel M. Nelson. *Learning About Control of Legged Locomotion Using a Hexapod Robot With Compliant Pneumatic Actuators*. PhD thesis, Case Western, May 2002.
- [132] TF Novacheck. The biomechanics of running. *Gait Posture*, 7(1):77–95, Jan 1 1998.
- [133] Colm ODunlaing. Motion planning with inertial constraints. *Algorithmica*, 2(4):431–475, 1987.
- [134] Yu Ogura, Hiroyuki Aikawa, Kazushi Shimomura, Hideki Kondo, Akitoshi Morishima, Hun ok Lim, and Atsuo Takanishi. Development of a humanoid robot WABIAN-2. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 76–81, 2006.
- [135] Hiroki Ohta, Masaki Yamakita, and Katsuhisa Furuta. From passive to active dynamic walking. *International Journal of Robust and Nonlinear Control*, 11(3):287–303, Mar 2001.
- [136] Hun ok Lim and Atsuo Takanishi. Waseda biped humanoid robots realizing human-like motion. pages 525–530. IEEE International Workshop on Advanced Motion Control, 2000.

- [137] Petri T. Piiroinen and Harry J. Dankowicz. Low-cost control of repetitive gait in passive bipedal walkers. *International Journal of Bifurcation and Chaos*, 15(6):1959–1973, 2005.
- [138] Robert Playter and Marc Raibert. Control of a biped somersault in 3d. In *Proceedings of the 1992 IEEE/RSJ International Conference on Intelligent Robots and Systems*, volume 1, pages 582–589. IEEE, July 1992.
- [139] N. S. Pollard, J. K. Hodgins, M. J. Riley, and C. G. Atkeson. Adapting human motion for the control of a humanoid robot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, volume 2, pages 1390–1397, May 2002.
- [140] Jovan Popovic, Steven M. Seitz, and Michael Erdmann. Motion sketching for control of rigid-body simulations. *ACM Transactions on Graphics*, 22(4):1034–1054, 2003.
- [141] Marko Popovic, Ambarish Goswami, and Hugh Herr. Ground reference points in legged locomotion: Definitions, biological trajectories and control implications. *International Journal of Robotics Research*, 24(12):1013–1032, Dec 2005.
- [142] Gill A. Pratt. Low impedance walking robots. *Integ. and Comp. Biol.*, 42:174–181, 2002.
- [143] Gill A. Pratt, Matthew M. Williamson, Peter Dillworth, Jerry Pratt, Karsten Ulland, and Anne Wright. Stiffness isn't everything. In *Proceedings of the 4th International Symposium on Experimental Robotics (ISER)*, 1995.
- [144] Jerry Pratt. *Exploiting Inherent Robustness and Natural Dynamics in the Control of Bipedal Walking Robots*. PhD thesis, Computer Science Department, Massachusetts Institute of Technology, 2000.
- [145] Jerry E. Pratt and Russ Tedrake. Velocity based stability margins for fast bipedal walking. In *Proceedings of the First Ruperto Carola Symposium on Fast Motions in Biomechanics and Robotics: Optimization and Feedback Control*, volume 340, pages 299–324, Sep 2005.

- [146] Raibert, M. H., Chepponis, M., Brown, and H. B. Running on four legs as though they were one. *IEEE Journal of Robotics and Automation*, 2(2):7082, 1986.
- [147] Marc H. Raibert. *Legged Robots That Balance*. The MIT Press, 1986.
- [148] Arlis Reynolds. Design and control of a clutch for a minimally-actuated biped based on the passive-dynamic simple walker. Mechanical Engineering Undergraduate Thesis, Massachusetts Institute of Technology, June 2006.
- [149] Gideon Sahar and John Hollerbach. Planning a minimum-time trajectories for robot arms. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)* [74], pages 751–758.
- [150] Y. Sakagami, R. Watanabe, C. Aoyama, S. Matsunaga, and N. Higaki and K. Fujimura. The intelligent ASIMO: system overview and integration. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems*, volume 3, pages 2478–2483. IEEE, 2002.
- [151] K. Salisbury. Active stiffness control of a manipulator in Cartesian coordinates. In *Proc. of the 19th IEEE Conference on Decision and Control*, 1980.
- [152] Philippe Sardain and Guy Bessonnet. Forces acting on a biped robot. Center of Pressure – Zero Moment Point. *IEEE Trans. on Systems, Man, and Cybernetics - Part A: Systems and Humans*, 34(5):630–637, Sep 2004.
- [153] Aaron Saunders, D. Goldman, R. Full, and M. Buehler. The RiSE climbing robot: Body and leg design. In *SPIE - Unmanned Systems Technology VIII*, volume 6230, page 623017. International Society for Optical Engineering, 2006.
- [154] A. L. Schwab and M. Wisse. Basin of attraction of the simplest walking model. *Proceedings of the ASME Design Engineering Technical Conference*, 6:531–539, Sep 2001.
- [155] Thomas B. Sheridan. Three models of preview control. *IEEE Transactions on Human Factors in Electronics*, 7(2):91–102, Jun 1966.

- [156] Zvi Shiller and Steven Dubowsky. Global time optimal motions of robotic manipulators in the presence of obstacles. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 370–375, 1988.
- [157] Alexander Shkolnik and Russ Tedrake. Inverse kinematics for a point-foot quadruped robot with dynamic redundancy resolution. In *Proceedings of the 2007 IEEE International Conference on Robotics and Automation*, Apr 2007.
- [158] R. Smith, M. Self, and P. Cheeseman. Estimating uncertain spatial relationships in robotics. In *Autonomous robot vehicles*, pages 167–193. Springer-Verlag New York, Inc., 1990.
- [159] M. W. Spong. Underactuated mechanical systems. In B. Siciliano and K. P. Valavanis, editors, *Control Problems in Robotics and Automation*, Lecture notes in control and information sciences 230. Springer-Verlag, 1997.
- [160] M W Spong and F Bullo. Controlled symmetries and passive walking. *IEEE Transactions on Automatic Control*, 50(7):1025–1030, Jul 2005.
- [161] Mark W. Spong. Swing up control of the acrobot. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2356–2361, 1994.
- [162] Mark W. Spong and Gagandeep Bhatia. Further results on control of the compass gait biped. In *Proceedings of the IEEE International Conference on Intelligent Robots and Systems (IROS)*, pages 1933–1938, 2003.
- [163] Jimmy Li-Shin Su and Jonathan B. Dingwell. Local dynamic stability of passive dynamic walking on an irregular surface. In *ISB XXth Congress - ASB 29th Annual Meeting*, Aug 2005.
- [164] Jimmy Li-Shin Su and Jonathan B. Dingwell. Dynamic stability of passive dynamic walking on an irregular surface. *Trans. of the ASME*, 129(6):802–810, 2007.

- [165] Tomomichi Sugihara, Yoshihiko Nakamura, and Hirochika Inoue. Realtime humanoid motion generation through zmp manipulation based on inverted pendulum control. *IEEE International Conference on Robotics and Automation (ICRA)*, 2002.
- [166] Richard S. Sutton and Andrew G. Barto. *Reinforcement Learning: An Introduction*. MIT Press, 1998.
- [167] A A Taflanidis and J L Beck. Stochastic Subset Optimization for optimal reliability problems. *Probabilistic Engineering Analysis*, 23(2–3):324–338, 2008.
- [168] Kiyotsugu Takaba. A tutorial on preview control systems. *SICE 2003 Annual Conference*, pages 1388–1393, Aug 2003.
- [169] P. Talkner, P. Hangii, E. Freidkin, and D. Trautmann. Discrete dynamics and metastability: Mean first passage times and escape rates. *J. of Stat. Phys.*, 48(1/2):231–254, 1987.
- [170] Russ Tedrake, Teresa Weirui Zhang, Ming-fai Fong, and H. Sebastian Seung. Actuating a simple 3D passive dynamic walker. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, volume 5, pages 4656–4661, New Orleans, LA, April 2004.
- [171] Russell L Tedrake. *Applied Optimal Control for Dynamically Stable Legged Locomotion*. PhD thesis, Massachusetts Institute of Technology, 2004.
- [172] Ivar Thorson, Mikhail Svinin, Shigeyuki Hosoe, Fumihiko Asano, and Kouichi Taji. Design considerations for a variable stiffness actuator in a robot that walks and runs. In *Proceedings of the Robotics and Mechatronics Conference (RoboMec)*, pages 1–4. JSME, May 2007.
- [173] D J Todd. *Walking Machines: An Introduction to Legged Robots*. Anchor Press Ltd., 1985.
- [174] Masayoshi Tomizuka. Optimal continuous finite preview problem. *IEEE Transactions on Automatic Control*, 20(3):362–365, Jun 1975.

- [175] Giovanni Tonietti, Riccardo Schiavi, and Antonio Bicchi. Design and control of a variable stiffness actuator for safe and fast physical human/robot interaction. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 526–531, 2005.
- [176] R. Q. van der Linde. Active leg compliance for passive walking. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 2339–2345, 1998.
- [177] Sethu Vijayakumar, Aaron D’Souza, and Stefan Schaal. Incremental online learning in high dimensions. *Neural Computation*, 17:2602–2634, 2005.
- [178] Miomir Vukobratovic. *Legged Locomotion and Anthropomorphic Mechanisms: A Monograph*. Beograd: Mihailo Pupin Institute, 1975.
- [179] Miomir Vukobratovic and Branislav Borovac. Zero-moment point - Thirty-five years of its life. *International Journal of Humanoid Robotics*, 1(1):157–173, 2004.
- [180] M. R. Walter, R. M. Eustice, and J. J. Leonard. Exactly sparse extended information filters for feature-based slam. *Int. Journal of Robotics Research*, 26(4):335–359, Apr 2007.
- [181] M Weber, S Kube, L Walter, and P Deuffhard. Stable computation of probability densities for metastable dynamical systems. Technical Report ZIB-Report 06-39, Konrad-Zuse-Zentrum für Informationstechnik Berlin, Nov 2006.
- [182] E R Westervelt, J W Grizzle, and D E Koditschek. Hybrid zero dynamics of planar biped walkers. *IEEE TRANSACTIONS ON AUTOMATIC CONTROL*, 48(1):42–56, Jan 2003.
- [183] E. R. Westervelt, B. Morris, and K. D. Farrell. Analysis results and tools for the control of planar bipedal gaits using hybrid zero dynamics. *Autonomous Robots*, 23:131–145, Jul 2007.

- [184] E.R. Westervelt, B. Morris, and K.D. Farrell. Sample-based HZD control for robustness and slope invariance of planar passive bipedal gaits. In *Proceedings of the International Conference on Robotics and Automation (ICRA)*, pages 1–6, September 2006.
- [185] Eric R Westervelt. *Toward a Coherent Framework for the Control of Planar Biped Locomotion*. PhD thesis, University of Michigan, 2003.
- [186] Eric R. Westervelt and Carlos Canudas de Wit. Walking and running biped robots [from the guest editors]. *IEEE Robotics and Automation Magazine*, 14(2):6–7, June 2007.
- [187] Pierre-Brice Wieber. On the stability of walking systems. In *Proceedings of the International Workshop on Humanoid and Human Friendly Robots*, 2002.
- [188] M Wisse, A L Schwab, R Q van der Linde, and F C T van der Helm. How to keep from falling forward: elementary swing leg action for passive dynamic walkers. *IEEE Transactions of Robotics*, 21(3):393–401, Jun 2005.
- [189] Martijn Wisse. *Essentials of dynamic walking; Analysis and design of two-legged robots*. PhD thesis, Technische Universiteit Delft, 2004.
- [190] Martijn Wisse and Jan van Frankenhuyzen. Design and construction of Mike; a 2D autonomous biped based on passive dynamic walking. In *Proceedings of the International Symposium on Adaptive Motion of Animals and Machines*, 2003.
- [191] Fuminori Yamasaki, Ken Endo, Hiroaki Kitano, and Minoru Asada. Acquisition of humanoid walking motion using genetic algorithms: Considering characteristics of servo modules. In *Proceedings of the IEEE International Conference on Robotics and Automation (ICRA)*, pages 3123–3128, 2002.
- [192] K Yoneda, H Iiyama, and S Hirose. Intermittent trot gait of a quadruped walking machine dynamic stability control of an omnidirectional walk. *Proc. of the IEEE Int. Conf. on Robotics and Automation*, pages 3002–3007, Apr 1996.

- [193] M. Zinn, O. Khatib, B. Roth, and J.K. Salisbury. Playing it safe. *IEEE Robotics and Automation Magazine*, 11(2):12–21, June 2004.