Semester Project at the
Biologically Inspired Robotics Group,
EPF Lausanne
Spring 08

Development and Test of a Model for the Cheetah Robot

Martin Riess

June 6, 2008

Abstract

This project aims at implementing a close-to-reality model of the recently developed Cheetah robot [1] in Webots [2], a commercial mobile robot simulation software developed by Cyberbotics Ltd. Furthermore, a pre-developed CPG [3] will be applied to the robot to determine feasible gaits and their optimal parameters for the Cheetah robot. Special interest lies in the characterization of the usefulness of the passively dynamic springs in the legs.

Contents

1 Introduction

This is the report summarizing the results of a semester project of the spring semester of 2008.

It builds upon a previous project, see [1], on the construction of Cheetah, a quadruped robot with passive dynamics. Having a realistic model of a newly built robot is always useful, as it can be used to characterize its dynamics and suggest interesting experiments for the real robot.

At the beginning of the project, a model of the Cheetah robot was built using the Webots [2] software, which involved the development of a physics plugin to simulate the knee joint mechanism. Most of the characteristic parameters of the robot could be taken from the construction paper while others had to be approximated. After the modeling phase a pre-existing CPG was ported over to the model and customized so that it could produce four different gaits: walk, pace, trot and bound.

Using this CPG, different experiments were done: A first set of experiments with higher than necessary spring strength in the legs and the default friction parameter between the feet and the ground provided satisfying results for some gaits, while others suffered from the low default friction and produced strange results. This motivated a second set of experiments with a more realistic spring strength in the legs and an infinite friction between the feet and the ground, which also provided insights into the robot's innate dynamics.

2 Task description

This report concludes the work of a semester project in the summer 2008 at the Biologically Inspired Robotics Group at the EPFL. The task was to first produce a realistic model of the pre-developed Cheetah robot [1] using Webots [2], a commercial mobile robot simulation software developed by Cyberbotics Ltd. Motivated through a literature review on CPG controllers for exploiting body dynamics of robots with passive elements and using the model, a previously existing CPG [3] was to be implemented. The goal here was to first implement it as an open loop control to get a feel for possible types of locomotion and their parameters, which then later can be finetuned by including feedback loops into the system.

3 About the robot

Cheetah is a lightweight quadruped robot with three-segment legs, see figure 1 for a side view. Each leg's relative dimensions are designed to be close to those of a real cat with the legs having only two degrees of freedom due to a pantograph mechanism used to keep the first and third leg segment parrallel during motion. This is achieved by having a spring parallel to the second leg segment which is contracted by a cable attached to a servo. The spring also introduces a passive element into knee joint, influencing both the contraction of the leg as a counterforce and supporting its expansion. See section 4.1

for further detail on this mechanism and how it is approximated in the model. We thus end up with eight actuated joints, two per three-segment leg. For more information see Simon Rutishauser's report on the construction of Cheetah [1].

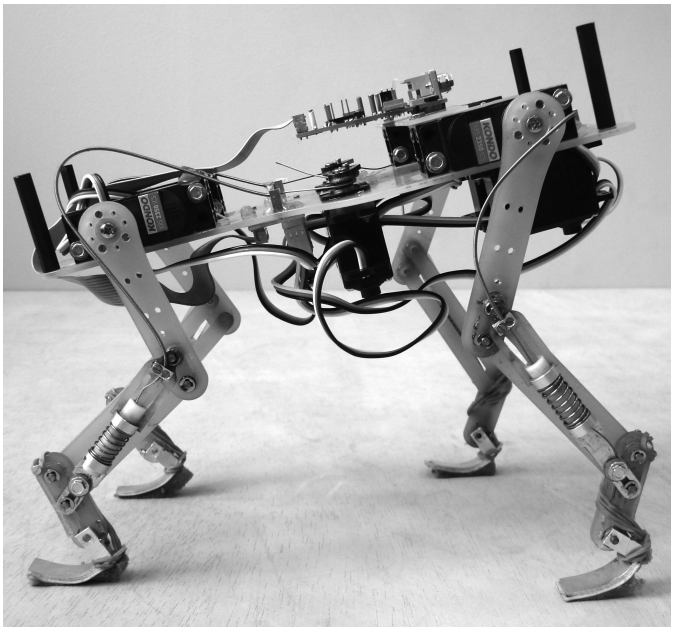

Figure 1: The Cheeath robot - side view

4 Description of the model

4.1 Overview

The model in its steady state can be seen in figure 2. I received the body and one leg pre-modeled from Yvan Bourquin, co-programmer of the Webots software at Cyberbotics [2].

The body's geometry was preserved as much as possible while not introducing too many elements. Physical parameters such as weight, motor torque and spring constants were taken from the construction report of cheetah [1] and held up to date with recent changes to the robot. Hip joints are modeled as normal angular motors. The knee joint mechanism is achieved through a physics plugin that creates a joint to complete the pantographic construction of the leg and controls a slider joint, which replaces the spring in the real robot. The passively dynamic feet are replaced by non-actuated servos with the `springConstant` and `dampingConstant` attribute and limited angles through `minStop` and `maxStop` in Webots.

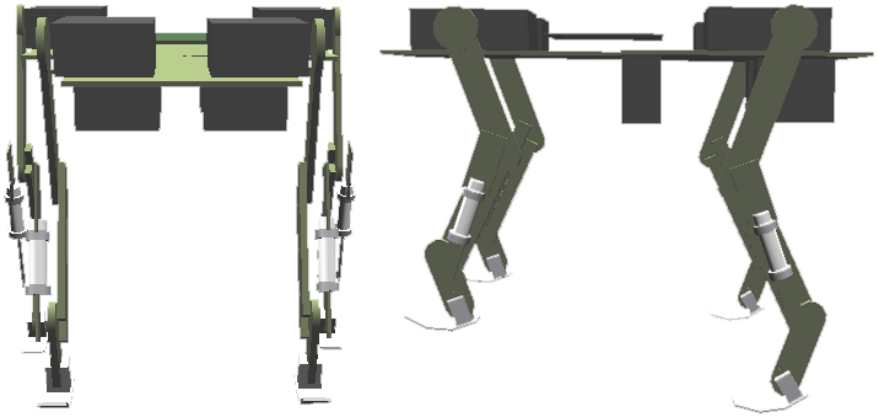

Figure 2: The Cheeath model - front and side view

4.2 Springs and pantograph mechanism - the physics plugin

Webots itself was not able to reproduce the knee mechanism's behavior, so we had to produce an ODE plugin to emulate it. Along with part of the model (see section 4.1), Yvan Bourquin programmed the required ODE functions for one leg to function properly, which I had to expand to all four legs.

The plugin mostly tries to emulate the real robot's legs' behaviour, which is highly nonlinear due to construction. This implicates differences between the model and the real robot, which I had to try to calibrate. This process mostly involved changing the numbers of the spring constants and motor torques to approximate effective values, which resulted in a motor torque of $0.95Nm$ for all joints and spring constants for the fore- and hind limbs of $3000\frac{N}{m}$ and $5000\frac{N}{m}$ or $2100\frac{N}{m}$ or $3500\frac{N}{m}$ respectively. See section 7 for an explanation of which value was used in the different experiments.

The first thing the plugin has to do is create a joint between the second and third leg segments, since the pantographic construction with the two-part second leg segment poses a problem to Webots' modeling feature. A linear slider-joint is used in the model to replace the springs. Then, during simulation, the plugin receives knee-joint data from the controller through an emitter object – see figure 3 – and controls the slider-joint accordingly.

The principle behind its calculations is that expanding movement on the slider can only be the result of the spring force and contraction can be due to the servo pulling the cable attached to the spring or an external force pushing the leg from below, which happens for example during touchdown of the leg after the swing phase. It thus looks at the direction of the movement required each timestep and applies the appropriate force. If a contraction due to external force took place, the leg thus also moves back into a stable position due to the spring force working in the opposite direction.

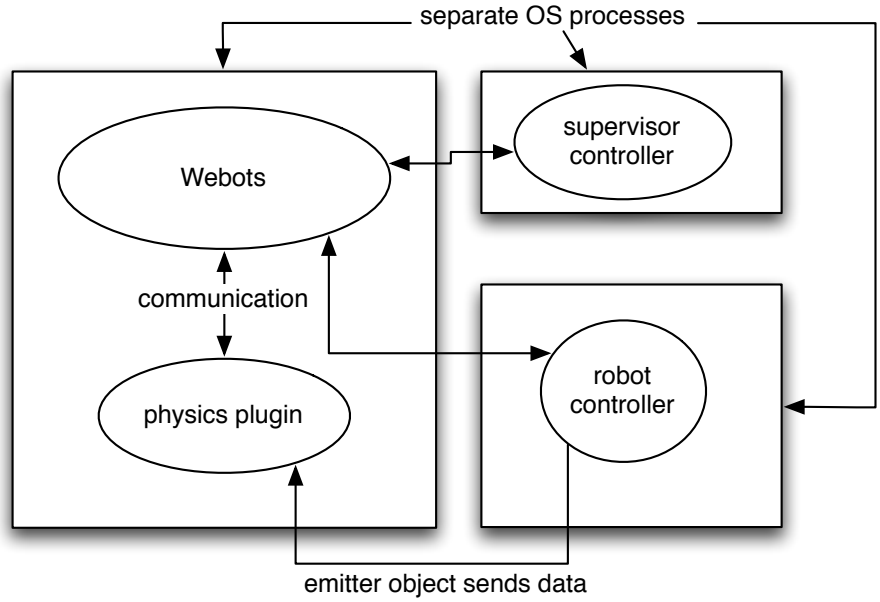

Figure 3: Webots - controller - plugin interaction. While the controllers run in separate OS processes and communicate with the Webots application through pipes, the plugin itself runs in the same process and memory space as the main application. The robot controller and the plugin are able to communicate through a Webots emitter object and the `dWebotsSend()` and `dWebotsReceive()` utility functions provided by the physics library. For more information, see the Webots user guide and reference manual, to be found at the Cyberbotics website [2].

5 Literature review: Robots with passive body dynamics and methods of controlling them

Controlling robots with passive body dynamics has been a field of study for some time, be it with one 3-segment leg [4], quadruped robots [5] or hexapods [6]. Inspired by research done in biology, it was established that when observing the behavior of a system driven by motors one has to take into account both the physical and the neural dynamics applied [7]. As Full points out in [8], the coupling of these dynamics seem to only become apparant in appropriately fast locomoting species, as the neural control dominates over physical dynamics in slower locomotors like crabs. It is this coupling of control and passive physical dynamics that inspired various control methods to be developed that try to cope with different problems one encounters when working with compliant robots and make use of this compliance.

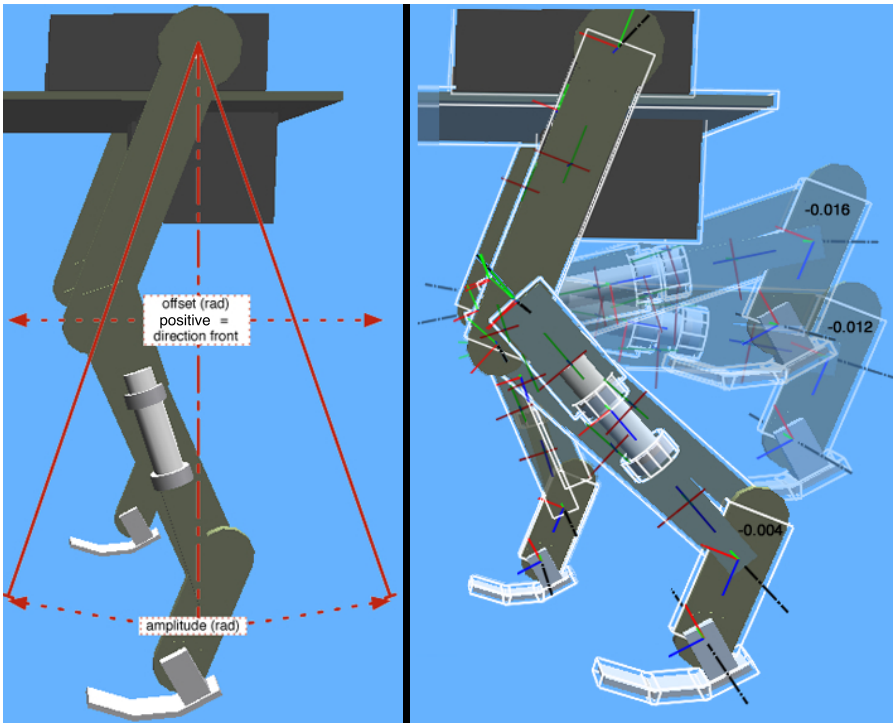

Figure 4: Leg actuation clarification. Both sides of the figure show the left hind leg: on the left at position 0 for both leg compression and offset; on the right example values for compression with 0 offset, with the physics plugin taking values in meters. The simulations in 7 and 8 were made with a minimum compression of -0.004m and a maximum compression of -0.012m.

5.1 Robot control on difficult terrain

Over time, many people built different robots with passive body dynamics and applied different ways of controlling these. A deeply studied robot with a CPG- and reflex control tailored specifically to its body dynamics is the Tekken quadruped robot (and later revisions of it), whose dynamics are studied in detail by Kimura [5] and Fukuoka [9]. Focus of these studies was stability of the robot on irregular terrain through introduction of different reflexes into the system. The model used for the CPG worked on the neuron level, with each oscillator consisting of one flexor and one extensor neuron, and was first introduced by Matsuoka [10]. Through various sensors on the robot, a high adaptability to the environment was achieved, so that the robot was able to walk in differently sloped terrains or go down steps. However, its passive dynamics are also mostly used to improve the generation of reflexes that facilitate walking on irregular terrain, no studies were done trying to optimize performance or energy consumption on flat terrain using passive body dynamics.

Similar, terrain-based research was done by Cham with the hexapodal robot Sprawlita [6]. Sprawlita is a very small hexapod with one degree of freedom per leg. In the cited

paper, Cham studied the robot's performance on terrain with different slopes. The simple open-loop controls used are much less sofisticated than those used for the Tekken robot, but also allow adaptation of the stride period by sensory feedback. A simple stride period minimizing adaptation was used and realized by measuring foot contact time during one movement cycle, however the results showed that this adaptation only leads to optimal speed on flat terrain and fails to do so on a sloped one.

5.2 Biped robots with passive dynamics and their controls

The usage of very simple controls for exploiting passive body dynamics was also shown by Collins et al. with biped robots inspired by passive-dynamic walkers [11]. The group constructed three robots that mimicked the passive walker's ability to walk down a slight slope only by using gravitational energy and their passive body dynamics, but introduced slight actuation to enable them to walk on flat ground. The robots learned the required control signals by reinforcement learning and were able to walk on flat ground without putting significantly more energy into the system through actuation than using gravitation through a slope would.

The same problematic was also discussed by Schiehlen in [12] from a more theoretical point of view. In this paper, the author approaches the problem of having biped passive-dynamic walkers without any active actuation on the one hand and actuated biped robots with an electric power source, but most of the time not more than half an hour of battery life, on the other hand. A model biped walking machine with 7 body parts and 9 degrees of freedom was proposed and theoretically studied through building an inverse kinematics controller for it that tries to exploit the robot's passive dynamics as much as possible. As a comparison measure to other walking machines, the specific resistance was chosen, the ratio between the robot's power consumption and the product of its weight and speed. As a result, the actuated model biped walking machine with its controller presented in the paper theoretically has a specific resistance comparable to the ones of passive-dynamic walkers. However, these were all theoretical results, a real-world building and testing of the paper's results would have been interesting.

Another approach on biped dynamic walking was taken by Geng et al. with RunBot in [13]. The robot is a lightweight, 23 cm high biped walker with 2-segment legs that is attached to a boom of one meter length and is thus constrained to moving in a circle. Only each of the leg's hip- and knee joints are actuated, its ankles are curved - which proved to be useful in passive dynamic walkers - and not actuated. RunBot's controller is receptor-driven, with stretch- and contact sensors used to signal join limits or gait phase transitions. Receptor- and motor neurons are organized in a network whose parameters are optimized through reinforcement learning. After a relatively short learning period it was possible to reach a speed which - relative to the size - is comparable with that of a human.

5.3 Generic controllers

Other papers describe the development of generic CPG controls for compliant robots. Righetti recently developed and described a CPG that is able to produce both robust locomotion through sensory feedback and different gaits [3]. To do so it uses a network of modified Hopf oscillators. The generality of the approach was shown by implementation on the Aibo, iCub and Ghostdog quadruped robots and its robustness through tests on both flat and uneven terrain. This CPG will also be used in this project to control the Cheetah robot.

A controller designed to specifically calibrate itself to the intrinsic frequency of the robots passive dynamics is described by Buchli in [14]. It uses an adaptive frequency Hopf oscillator to constantly readapt the system using information of sensory feedback. The controller thus tries to adapt the oscillators' intrinsic frequency to that of the robot, converging to parameters that then help developing a stable and efficient controller tailored to it.

6 CPG implementation

6.1 Description of the CPG

The CPG in question is described in detail in Ludovic Righetti's publication [3]. Very briefly, it uses a modified Hopf oscillator that follows these nonlinear differential equations:

$$\dot{x} = \alpha(\mu - r^2)x - \omega y \tag{1}$$

$$\dot{y} = \beta(\mu - r^2)y + \omega x \tag{2}$$

$$\omega = \frac{\omega_{stance}}{e^{-by}+1} + \frac{\omega_{swing}}{e^{by}+1} \tag{3}$$

where $r = \sqrt{x^2+y^2}$, ω is the frequency of oscillations in $rad * s^{-1}$, $\sqrt{\mu}$ is the amplitude of oscillations, ω_{swing} and ω_{stance} are the frequencies of the swing and stance phases respectively. α and β are positive constants that control the speed of convergence to the limit cycle.

6.2 The oscillator-network

The network can be seen in figure 5, it is again taken from [3]. The choice of gait the network produces is made solely through the different couplings of the oscillators, it is thus only required to store a coupling matrix for every gait one wants to produce and change them accordingly.

The equations listed in 6.1 then change to

$$\dot{x_i} = \alpha(\mu - r_i^2)x_i - \omega_i y_i \tag{4}$$

$$\dot{y_i} = \beta(\mu - r_i^2)y_i + \omega_i x_i + \sum k_{ij} y_i \tag{5}$$

with k being the coupling matrix' entry for the oscillator.

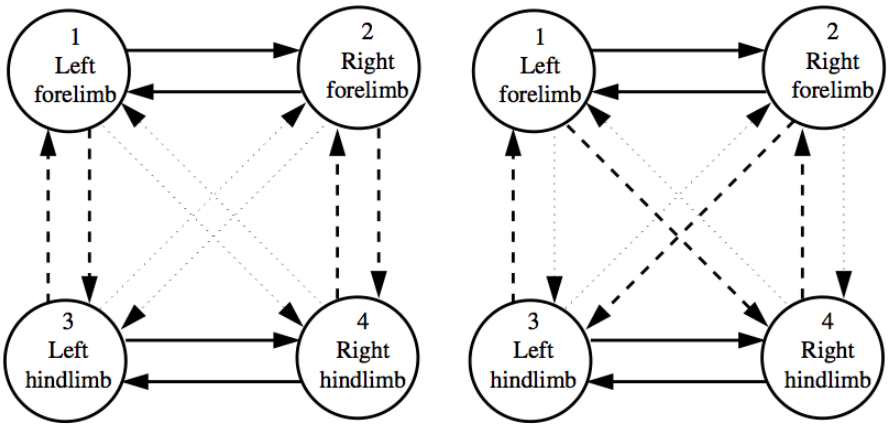

(a) Trot Pace Bound Net. (b) Walk Net.

Figure 5: The oscillator network used. Taken directly from [3].

6.3 Application to the robot

For the CPG output to be used as signals to the controlling servos, adjustments had to be made. This was the case especially for the leg-length controlling servo, this is due to the construction of the mechanism in the physics plugin (see section 4.2). The plugin uses compression values in meters as input, with zero being no spring compression (see also 4). One output of the CPG is thus transformed by the following pseudocode:

For the knee joint (with x oscillating around 0 with amplitude 1):

$$\text{applied value} = \begin{cases} \text{max\_retract} & \text{if } dx > 0 \text{ and } x > 0.8 \quad (1) \\ \text{min\_retract} & \text{if } dx > 0 \quad (2) \\ \text{min\_retract} - 0.1 \cdot (1 - \|x\|) & \text{if } dx < 0 \text{ and } x < -0.9 \quad (3) \\ \text{max\_retract} & \text{else} \quad (4) \end{cases}$$

with appropriate values for max_retract and min_retract. Case (1) represents the end of the stance phase, a backwards moving leg shortly before takeoff, case (2) represents the stretched leg during stance phase with the leg moving backwards and cases (3) and (4) thus represent the stance- to swing phase transition of touch down and the swing phase with a maximally retracted leg respectively.

For the hip joint:

$$\text{applied value} = -x \cdot [1.2 \text{ if forelimb}] \cdot \text{amplitude} + \text{offset}$$

The CPG's output is thus taken and directly applied to the servo, with optional addition of an amplitude increase or an offset. A default amplitude factor of 1.2 is chosen for the forelimbs to compensate for them being shorter than the hind ones.

Typical servo signals are shown in figure 6. The irregular parts around the $y = 0$ line of the hip joint signal is caused by the oscillator coupling.

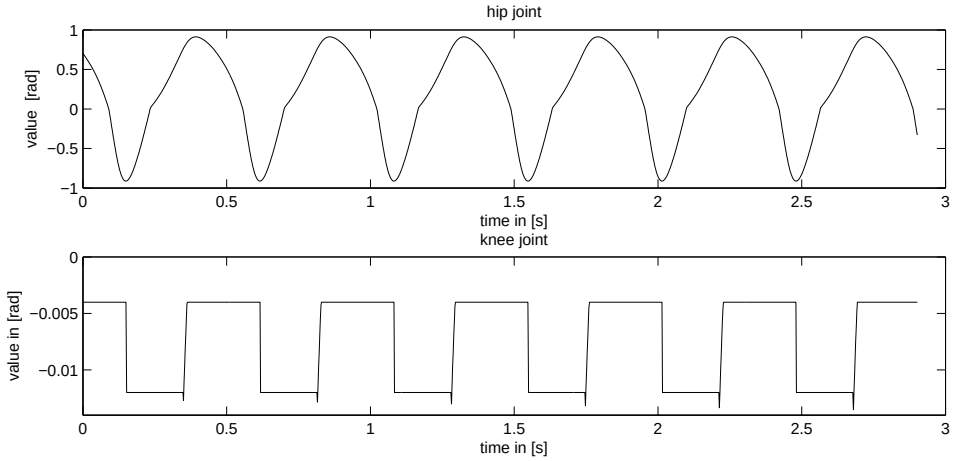

Figure 6: Typical signals obtained from the CPGs oscillation and sent to the servos. The upper picture shows the signal for the hip joint, the lower one that for the knee joint.

7 Naïve experiments in simulation

7.1 Setup

All experiments were made with the robot starting on a plane ground, facing in x-axis direction. Convergence time of the CPG at the beginning of the simulation was excluded from the experiments by not sending signals to the servos for a set time period, i.e. 5 seconds. The robot was then actuated for another time period, after which the distance it covered in x- and z-direction was recorded. Most interesting is the distance in x direction, since we want to achieve a stable gait both in the sense of not falling over and going straight.

A minimum leg compression of -0.004 and a maximum leg compression of -0.012 was chosen. This was done because usage of the full spring length for compression resulted in highly unstable movement and to mimic the real knee servos' behaviour of limiting their amplitude due to the high frequency they have to cope with, which was shown in real experiments. We also used the default values for friction between the feet and the ground, a friction coefficient of 1 for the ODE coulomb friction model. This value, however, turned out to lead to "unnatural" behaviour especially for the walking gait, which motivated a second set of experiments with infinite friction between feet and ground (see section 8). For the spring strengths, we had to differ from the ones used in the real robot because they proved to be too "weak" for the model. This is most likely due to several factors like a higher effective damping in the real robot and friction of the wire attached to the spring inside its cable, which was also mentioned as a problem in [1].

The requirements for the spring strength in the real robot were mainly being able to support the robot's weight at all times while also being maximally usable during movement. This resulted in a spring strength of $900\frac{N}{m}$ in the forelimbs and $2300\frac{N}{m}$ in the

experiment	w1	w2	w3	p1	p2	p3
gait	walk	walk	walk	pace	pace	pace
ω_{swing}	*	8	8	*	10	10
ω_{stance}	*	*	3	*	*	17
amplitude	0.73	*	0.9	0.37	*	3.1
offset forelimbs	0.244	0.244	*	0.37	0.37	*
offset hind limbs	-0.175	-0.175	*	0.12	0.12	*
experiment	t1	t2	t3	b1	b2	b3
gait	trot	trot	trot	bound	bound	bound
ω_{swing}	*	10	10	*	25	25
ω_{stance}	*	*	23	*	*	15
amplitude	0.9	*	1.7	0.5	*	2.1
offset forelimbs	0	0	*	0	0	*
offset hind limbs	0	0	*	0	0	*

Table 1: Experimental setups. Cells marked with a * are being investigated in the experiment. Amplitude and offsets in $[rad]$, ω_{swing} and ω_{stance} in $[\frac{rad}{s}]$. All experiments were done with friction coefficient 1 between feet and ground and a spring strength of 3000 and 5000 $\frac{N}{m}$ for fore- and hind limbs respectively.

hind limbs. The high difference is chosen due to the different length of fore- and hind limbs – the latter being 20% longer – and the uneven mass distribution of the robot.

For the following experiments a very high spring strength of 3000 $\frac{N}{m}$ in the forelimbs and 5000 $\frac{N}{m}$ in the hind limbs was chosen.

Three same experiments were done for the four gaits walk, pace, trot and bound. First, convenient values for the movement amplitude and offsets were chosen and a search for optimal values of ω_{swing} and ω_{stance} was done. The next experiment conserved the optimal value of ω_{swing} and the chosen offsets and searched for an optimal combination of ω_{stance} and amplitude. The last experiment conserved the values for both ωs and amplitude and searched for optimal offsets. Here, only offsets to the front for the forelimbs and offsets to the back for the hind limbs are investigated, since this combination looked the most promising to obtain stable results. A summary of the experimental setups can be found in table 1.

For experiments with a more careful calibration of the spring strength see section 8.

7.2 Walk

The simulation results can be seen in figure 7. Experiment a) shows one optimal peak at low values of ω_{stance} and medium ones of ω_{swing}, with a larger area of acceptable speeds with increasing ω_{swing}. Notable here is that for some parameter combinations the robot moved backwards at about the same speed as the fastest combination for forward movement achieved. Part b) of figure 7 represents the result of the ω_{stance} and amplitude experiment, which express one clear peak at $\omega_{stance} = 3$ and an amplitude

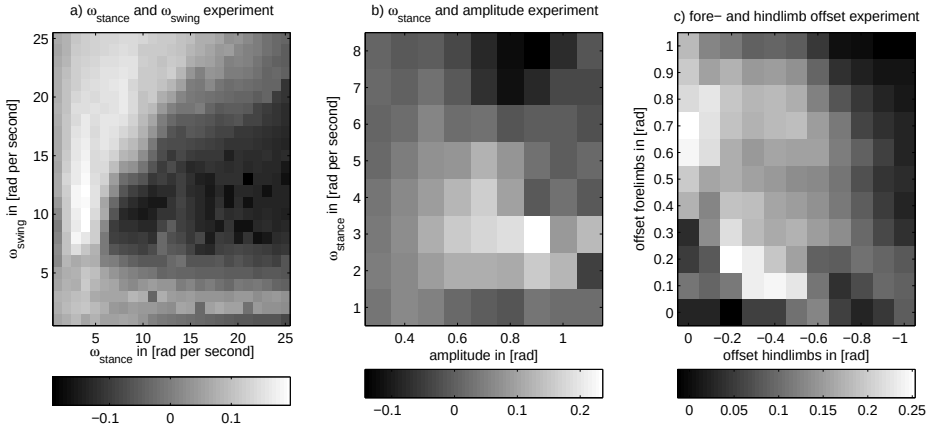

Figure 7: Results of the three naïve experiments for the walk gait. The coloration signifies the speed in x direction in $[\frac{m}{s}]$

of 0.9. The result of the third experiment, however, is not as clear, showing two peaks with comparable speeds, one at an approximately equal, small- offset of fore- and hind limbs to the front and back respectively, and one at a large offset of the forelimbs with no offset of the hind limbs. Further investigation showed, that the peak at $(0.7, 0)$ only yields unstable results, with the robot tilted in one direction and more falling forward than walking. A series of snapshots for both cases can be seen in figure 8. It also has to be noted that both configurations look rather strange and do not resemble what experiments with the real robot and the walking gait looked like. This behavior motivated the experiments in section 8.

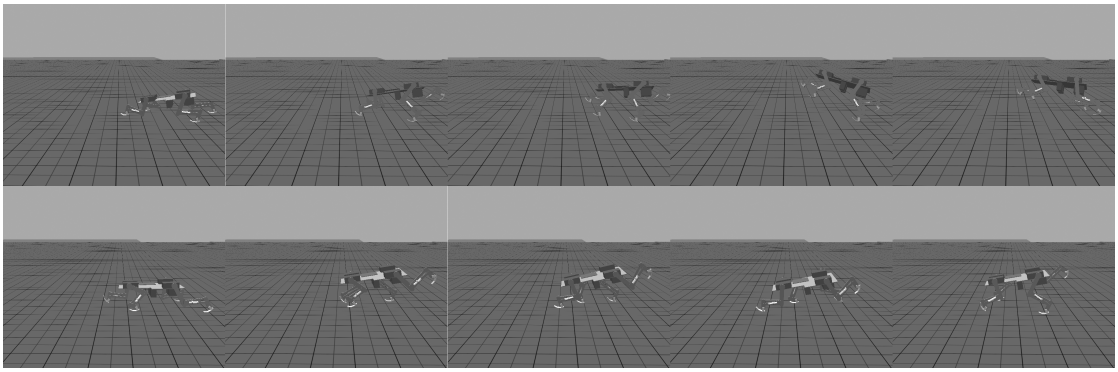

Figure 8: Snapshot series for both the offset optimum at $(0.7, 0)$ (upper row) and the one at $(0.2, -0.2)$ (lower column) for fore- and hind limb offset respectively, 100ms between each shot, robot moving from right to left. Note the robot nearly falling to one side in the upper row, while having relatively stable movement in the lower row.

7.3 Pace

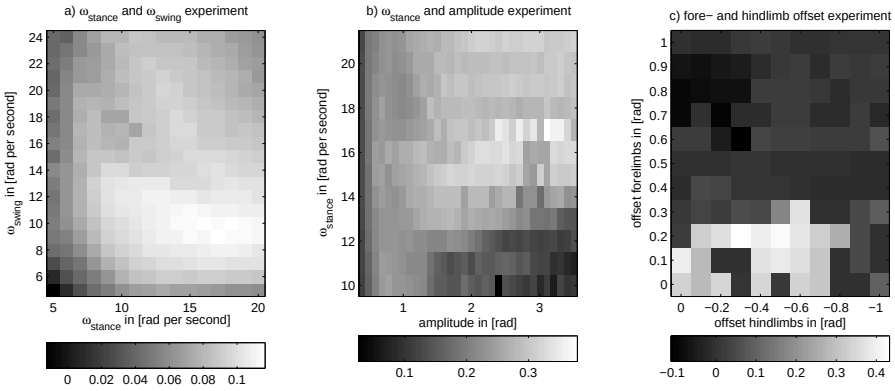

Figure 9: Results of the three naïve experiments for the pace gait. The coloration signifies the speed in x direction in $[\frac{m}{s}]$

As can be seen in part a) of figure 9, the first experiment leads to a large plateau of good values with ω_{swing} ranging between 8 and 12 and ω_{stance} going from about 12 to the end of the investigated range. The most plausible explanation for this is a saturation due to limited servo speeds and torques. Notable here is the very large area of usable values, indicating that pace might be a very stable gait for this robot. The second experiment also shows a large field of usable values for locomotion, but provides one more or less clear peak. The offset experiment shown in part c) of figure 9 however suggests that the pace gait for this robot very susceptible to changes of the limbs's offset: it shows a clear region of acceptable values, but for the most part the robot either does not move forward or falls over.

A snapshot series in figure 10 shows that the obtained optimal gait realistically resembles the pace gait.

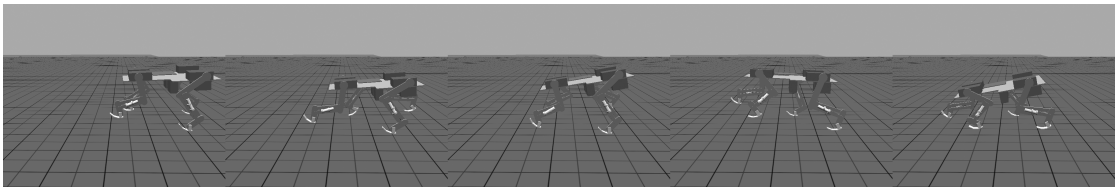

Figure 10: Snapshot series of optimal configuration obtained through the naïve experiments for pace, 100 ms between each shot, robot moving from right to left.

7.4 Trot

The experimental results can be found in figure 11. The first experiment shows that values smaller than 4 or 5 for both ω_{swing} and ω_{stance} are unusable, as the CPG does not

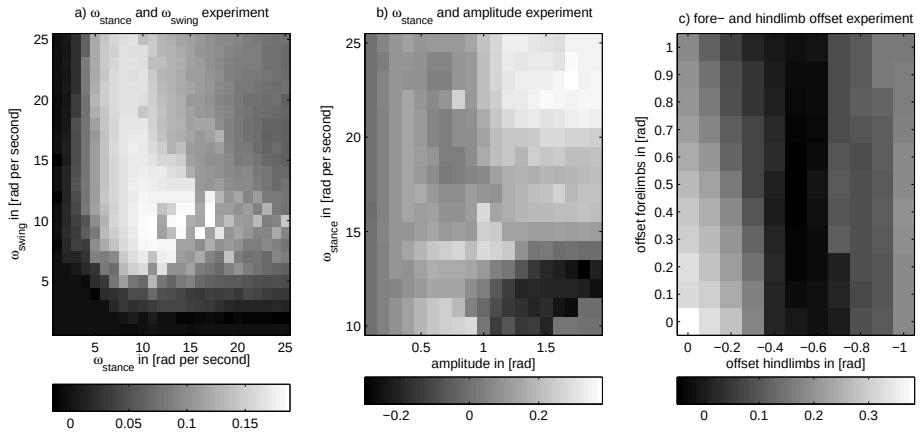

Figure 11: Results of the three naïve experiments for the trot gait. The coloration signifies the speed in x direction in $[\frac{m}{s}]$

converge to limit cycles for these values. When looking at the simulation, the limbs just move to their respective starting values and stay there. The experiment also provides a rather large area in which good results could be obtained, with ω_{swing} ranging from 6 up to 25 and ω_{stance} being between 6 and 12-15, the largest cluster being at $\omega_{swing} \in [7, 15]$ and $\omega_{stance} \in [7, 17]$.

This impression is shifted, however, when we look at the second experiment. The graph exhibits a large plateau at in the high region of both ω_{stance} (tested from 10 to 25) and the amplitude (tested from 0 to 1.9). We can also see a lower peak at values corresponding to those of the first experiment - an amplitude of 0.9 and ω_{stance} being around 15 -, but that seems to be only a local maximum.

The offset experiment shows an optimal value at the (0,0) position, with declining efficiency around it until the robot falls over for too large offsets.

As can be seen in the snapshot series in figure 10, the obtained optimal gait realistically resembles the trot gait.

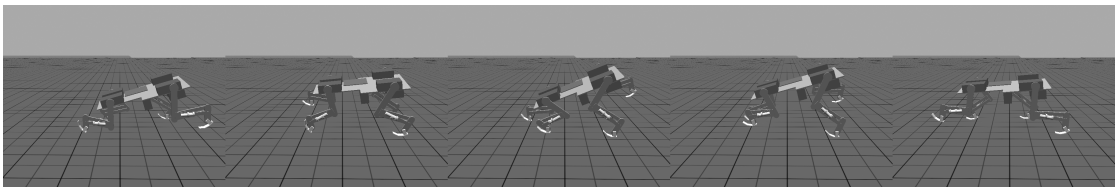

Figure 12: Snapshot series of optimal configuration obtained through the naïve experiments for trot, 100 ms between each shot, robot moving from right to left.

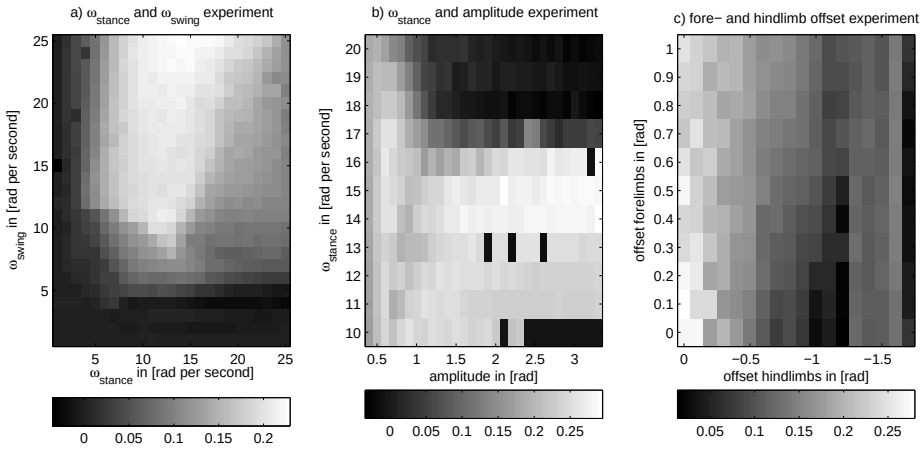

Figure 13: Results of the three naïve experiments for the bound gait. The coloration signifies the speed in x direction in $[\frac{m}{s}]$

7.5 Bound

Results of the simulations are shown in 13. Experiment a) shows similar characteristics as the first experiment for the trot gait, shown in part a) of figure 11. The area of usable values starts when both ω_{swing} and ω_{stance} are above approximately 5, and is expanding as the value of ω_{swing} gets higher. The maximum is actually obtained at the upper edge of the spectrum, with $\omega_{swing} = 25$ and $\omega_{stance} = 14$.

The second experiment confirmes the choice of ω_{stance} being optimal around the value 15 and shows a saturation regarding the amplitude from about 1.4 onwards. This is again most likely due to the limited speed and torque of the servo motors.

The investigation of the offsets provides one optimal region very closely around (0/0) and is more susceptible to changing the offset of the hind limbs as opposed to that of the forelimbs.

A snapshot series in figure 14 shows that the obtained optimal gait realistically resembles the bound that it should.

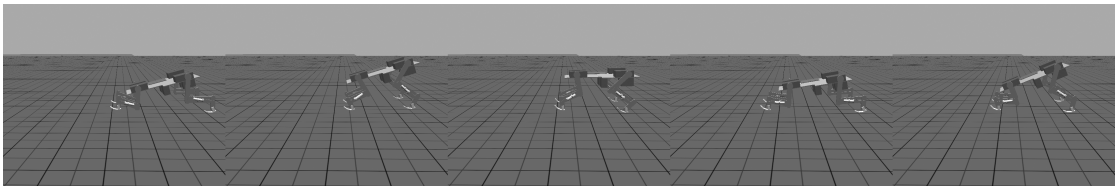

Figure 14: Snapshot series of optimal configuration obtained through the naïve experiments for bound, 100 ms between each shot, robot moving from right to left.

7.6 Discussion

Firstly, the method of parameter search used in these experiments has a few disadvantages: As all experiments build on top of each other, every experiment gives a bias to the next one. There is thus a danger of only finding local maxima instead of globally optimal parameters. This is partly counteracted by investigating ω_{stance} in both experiment 1 and 2 of each gait.

This being said, the results of all gaits show a significant increase in speed from one experiment to the next, with pace as the fastest reaching about $40\frac{cm}{s} \approx 1.7\frac{\text{bodylengths}}{s}$. The direct portability of these results to the real robot is questionable, however, since two major factors were disregarded in these experiments: The spring strength was chosen very high and the standard friction coefficient leads to strange and unrealistic looking behavior with the walking gait, which looks more like the robot sliding over the ground with this setup. This strange behavior of the walking gait motivated a second set of experiments, see section 8. Other gaits don't seem to have these problems and look "normal", though trot suffers from a tilting of the robot to one side, depending on initial conditions.

Another phenomenon observed with gaits at higher oscillation frequencies is shown in figure 15. The test with the trot gait at a frequency of about 2.2Hz suggests that other parameters lead to an unwanted offset of the leg: Even though the signal sent to the servo oscillates around 0, the actual servo movement oscillates around about 0.2. Further tests showed that this is a consequence of high differences between the two ω parameters and high frequency due to high ω values. I did not have enough time to quantify the impact of this phenomenon on the experiments, but its existence should be kept in mind. It can also be seen that especially the hip servo is not at all able to keep up, its amplitude is significantly lower and it has a delay compared to the signal.

8 Experiments with infinite friction

8.1 Setup

This set of experiments consists of one experiment per gait, combining 4 seemingly interesting values for each of the 5 parameters ω_{swing}, ω_{stance}, amplitude, offset of the forelimbs and offset of the hind limbs with each other. The values were chosen from the optimal regions found in the experiments in section 7, but contrary to those experiments every parameter combination is taken into account. For a summary of what values were chosen for the experiments, see table 2.

Furthermore, the spring strength was reduced from 3000 and 5000 $\frac{N}{m}$ to 2100 and 3500 $\frac{N}{m}$ for fore- and hind limbs respectively. This was chosen by looking at runs with several values, looking for one in which the springs were able to support the robot during motion while still making them as soft as possible so that they still influence the motion.

Another change was introducing infinite friction between the feet and the ground. The main goal here is to investigate the walking gait further, as it failed to produce satisfactory realistic results in the set of experiments in section 7. Looking at the optimal walking gait

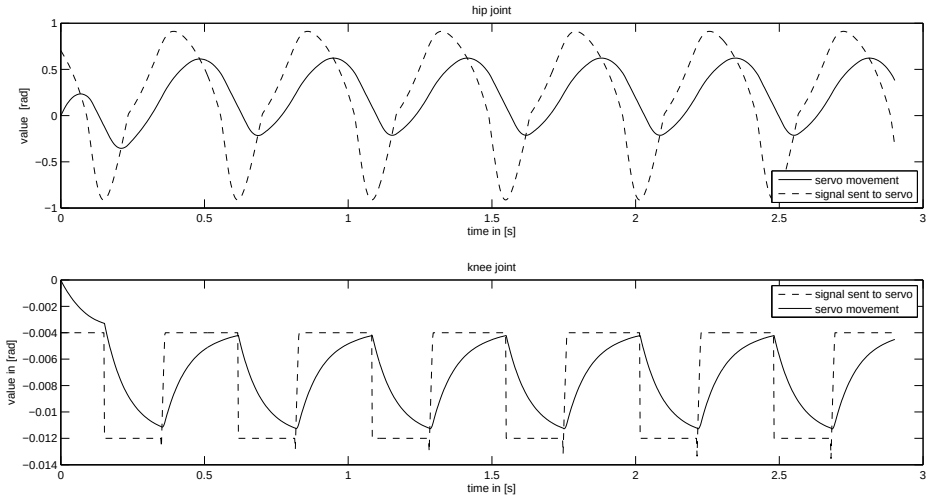

Figure 15: Servo signals (dashed lines) and actual servo movement (normal lines) of the right fore leg during trot gait with parameters: $\omega_{swing} = 10$, $\omega_{stance} = 23$, amplitude = 0.7 and both offsets set to 0. Test done while floating in the air with no gravity.

those experiments found, friction seemed to play a major role in the strange behavior. First tests were made by simply increasing the friction coefficient, which didn't change much of the behavior. This lead to the choice of an infinite friction between the feet and the ground, whose impacts on the different gaits are investigated in the following experiments.

The overall experiments were not changed: the robot starts out facing x direction. 5 seconds pass without signals being sent to the servos so that the CPG can converge properly, then the robot is actuated for 20 seconds. At the end of that time, the position in x and in z direction is recorded, where the speed in x direction is the main quality measure and the deviation in z direction is looked at to find out if the robot is walking

gait	walk 1	walk 2	pace	trot	bound
ω_{swing}	$[6, 9]$	$[3, 5]$ in steps of 0.5	$[9, 12]$	$[6, 9]$	$[22, 25]$
ω_{stance}	$[4, 7]$	$[1, 3]$ in steps of 0.5	$[15, 18]$	$[4, 7]$	$[17, 20]$
amplitude	$[0.7, 1.0]$	$[0, 0.4]$	$[0.7, 1.0]$	$[0.7, 1.0]$	$[0.7, 1.0]$
offset forelimbs	$[-0.2, 0.4]$	$[0, 0.4]$	$[0, 0.4]$	$[-0.2, 0.4]$	$[0, 0.4]$
offset hind limbs	$[0.2, -0.4]$	$[0.2, -0.4]$	$[0.2, -0.4]$	$[0.2, -0.4]$	$[22, 25]$

Table 2: Experimental setups with infinite friction and softer springs. Amplitude and offsets in $[rad]$, ω_{swing} and ω_{stance} in $[\frac{rad}{s}]$. All experiments were done with friction coefficient -1 (meaning infinite friction) between feet and ground and a spring strength of 2100 and 3500 $\frac{N}{m}$ for fore- and hind limbs respectively.

straight. The simulation is then reset and new parameters are set.

8.2 Walk

There were two experiments done for the walking gait: As can be seen in table 2, the first experiment uses an area around the values found to be optimal in section 7 for the parameters, while the second one investigates lower values for the ω parameters and the amplitude.

Experiment one yielded 119 configurations - out of 3520 - that were faster than $0.13\frac{m}{s}$ and 13 configurations faster than $0.15\frac{m}{s}$ with the maximum being at about $0.173\frac{m}{s}$. These configurations are actually diverse, but viewing them shows that they all result in approximately the same behavior, depicted in figure 16. As can be seen in the figure, the optimal behavior the experiment finds for the walking gait resembles more a pace than a walk, which is due to the hind legs touching down very early during the cycle and thus being stopped from fully moving forward. Figure 17 shows the signals sent to the servo and the actual servo movement of the right hind- and foreleg, illustrating the inability of the hind leg of the hind leg to properly move forward.

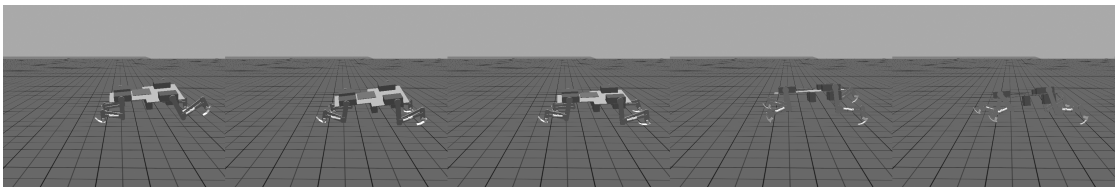

Figure 16: Snapshot series of what the optimal configurations for walk look like, robot moving from right to left.

The second experiment with lower frequency and amplitude shows similar behavior at slower speed: out of 1750 tested configurations, 22 were found to be faster than $0.1\frac{m}{s}$, the maximum being at $0.121\frac{m}{s}$, and all of them have the hind legs bent backwards at least 0.2 radians, so that the same behavior as depicted in 16 is produced. The hind legs only serve to lift the body up a bit while the fore leg on the same side drags the robot forward.

For an analysis of the evolution of the speed during the experiment, see figure 18. It also has to be noted that the higher speeds only occur with both $\omega_{swing} > \omega_{stance}$ with a difference of at least 2 and mostly at higher amplitudes of 0.9 or 1.0. While this makes sense, higher frequency leading to higher speed, the relation between ω_{swing} and ω_{stance} seems to be an important one, as can also be seen from the walk experiment in section 7.

8.3 Pace

The experiment for the pace gait was done with parameters around the values found to be optimal in section 7. Pace performed the best out of the four gaits with a maximal

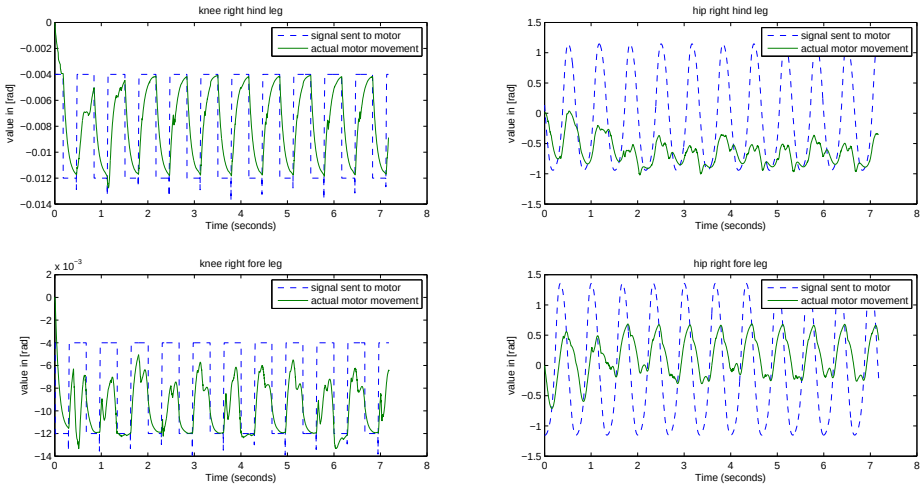

Figure 17: Servo signals (blue, dashed lines) and actual servo movement (green lines) of the right fore leg during the optimal configuration found by experiment one of the walk gait. Forward movement of the hind leg and stretching of the fore leg are significantly hindered by the friction at touchdown, resulting in a more pace-like movement.

speed of $0.286\frac{m}{s}$ and overall 605 out of 2496 configurations running faster than $0.15\frac{m}{s}$. The average speed over all runs is $0.11\frac{m}{s}$.

Figure 19 shows a snapshot series of the optimal parameter set. Something to note here is that counterintuitively, almost all of the fastest configurations have the forelegs bent backwards and the hind legs bent forwards. This seems like it would be a more unstable situation than the other way around, but it produces the best trajectories in this experiment.

8.4 Trot

The trot experiment was run at lower ω values than the ones found to be optimal for this gait in section 7, as can be seen in table 2. From the 2800 runs done for this gait, one optimal configuration could be found that works at around approximately $0.15\frac{m}{s}$ with the next best one at $0.1\frac{m}{s}$; see figure 20 for the full picture.

A snapshot series of the optimal parameter setup can be found in figure 21. It shows that with each cycle the robot is falling onto one foreleg, compressing the spring pretty far. The usage of the springs here, at least to stop the fall, is evident, whether this behavior can be reproduced on the real robot or not remains an open question for now.

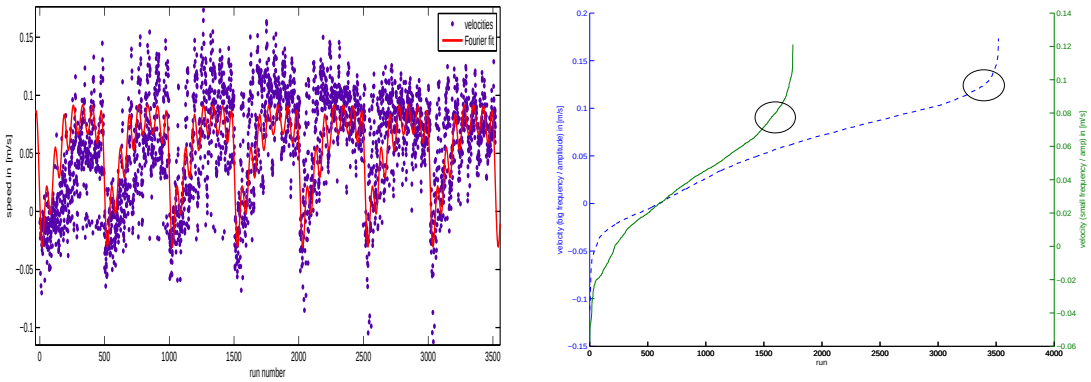

Figure 18: Left: evolution of speed during experiment 1 with the walking gait. The red line is a Fourier fit through the data points. Right: speeds obtained by both walk experiments sorted from slowest to fastest (dashed line being the first experiment). The circles marks the approximate point from which on the "pace-like" gait is done efficiently. When going to the slower configurations, one still encounters the "pace-like" behavior, although not as efficient, i.e. the legs don't have enough ground contact, too much ground contact etc.

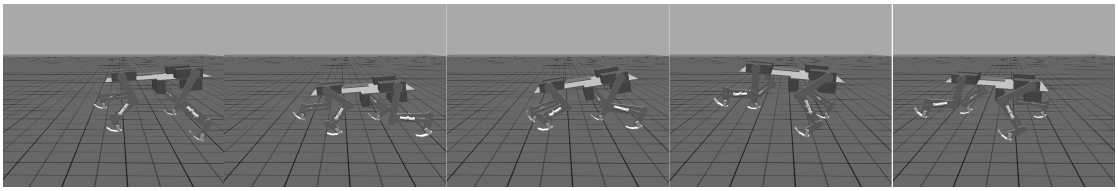

Figure 19: Snapshot series of optimal configuration for pace, 100 ms between each shot, robot moving from right to left.

8.5 Bound

An attempt to do the bound experiment at lower values of ω_{swing} and ω_{stance} did not bring any usable results, thus the values depicted in table 2 were chosen as an area around the values found to be optimal in section 7. 42 out of the 1600 tested configurations were found to be faster than $0.15\frac{m}{s}$ with the maximum at $0.173\frac{m}{s}$.

When looking at the robot's behavior, the bound gait is quite clearly expressed, as can also be seen in the snapshot series in figure 22. With the good parameter values there even is a short flying phase, where no leg touches the ground.

8.6 Discussion

See table 3 for a summary regarding the number of configurations found to be faster than different speeds for each gait. Overall, when sorted from slowest to fastest, the speed evolution of all experiments showed the same characteristics as the ones depicted in plots

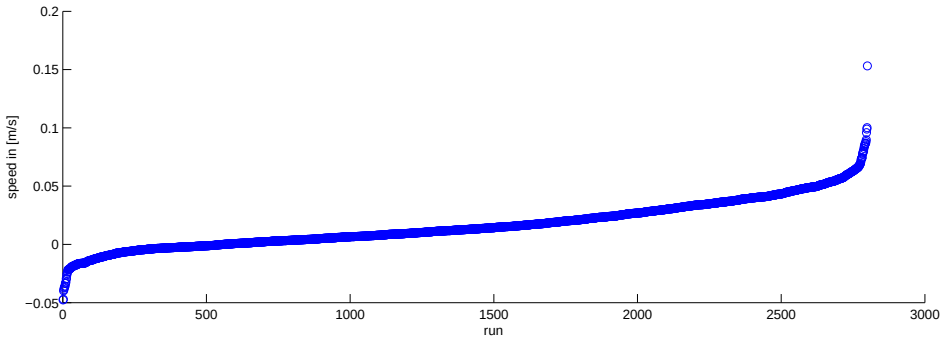

Figure 20: Scatterplot of the velocities obtained with the trot gait, sorted from slowest to fastest.

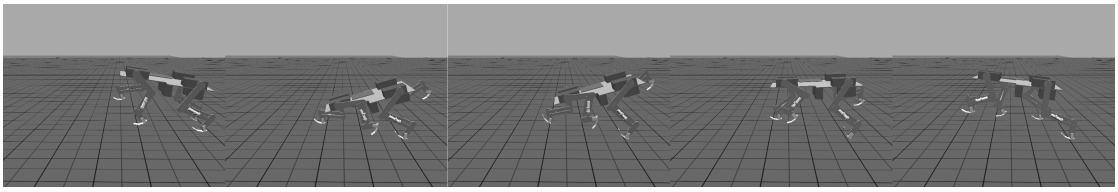

Figure 21: Snapshot series of optimal configuration for trot, 100 ms between each shot, robot moving from right to left.

for some gaits. The table already tells us different things:

Walk performed quite well, with over one third of the configurations being faster than $0.09\frac{m}{s}$. Looking at faster speeds, however, the number of configurations that can keep up quickly becomes less, showing that for faster speeds the gait becomes very dependent on the exact right parameter set. Experiment two showed that walk also works with lower frequency, in contrast to some other gaits. As can be expected, the achieved speed is lower than that of the first experiment, but it still performs quite well.

Regarding trot, the chosen parameter range has been suboptimal, even though it still does fall into an area which yielded good results in section 7. Only 6 out of 2800 tested configurations ($\approx 0.2\%$) achieved a speed of over $0.09\frac{m}{s}$. Unfortunately, I did not have enough time to run the same test for higher frequencies, which should bring better

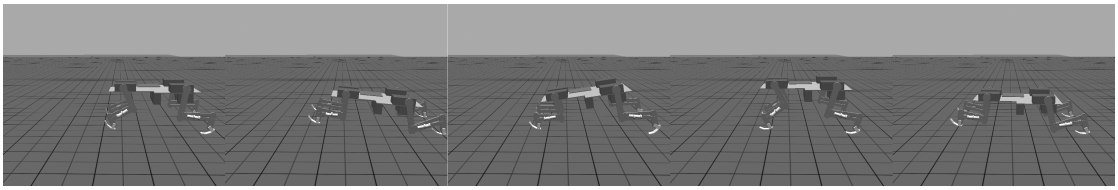

Figure 22: Snapshot series of optimal configuration for bound, 100 ms between each shot, robot moving from right to left.

gait	walk 1	walk 2	pace	trot	bound
number of configurations tested	3520	1750	2496	2800	1600
max speed	$0.173\frac{m}{s}$	$0.121\frac{m}{s}$	$0.286\frac{m}{s}$	$0.15\frac{m}{s}$	$0.173\frac{m}{s}$
configurations $> 0.09\frac{m}{s}$	939	64	1731	4	724
configurations $> 0.12\frac{m}{s}$	185	1	1101	1	264
configurations $> 0.15\frac{m}{s}$	13	0	605	1	43
configurations $> 0.18\frac{m}{s}$	0	0	279	0	0

Table 3: Summary of the results for the infinite friction experiments regarding the number of configurations found to be faster than different speeds.

results. Trot did not seem like a very stable gait overall, though, so I would not expect it to perform very well even with adapted parameters.

Bound on the other hand proves to be a very stable gait, parameter-wise: Almost half of the configurations are faster than $0.09\frac{m}{s}$ and still $43 \approx 3\%$ are faster than $0.15\frac{m}{s}$ (as opposed to $\approx 0.4\%$ with walk).

Another interesting point is the analysis of the robot's behavior for the different gaits: the walking gait does not really produce a walk as an optimal configuration, but something that resembles a pace. This is achieved by exploiting the very high friction, essentially stopping the hind legs from performing their whole planned movement. From this, one could conclude that walk is no very good gait for this robot and that the fast configurations of the gait try to "convert" the walk into a pace to reach higher speed.

This behavior seems reasonable when looking at the performance of the pace gait: Not only did it reach the highest speed, but it is also the most stable gait from a parameter point of view and produces the most realistic looking trajectory, together with bound.

The same can not be said for the trot gait. With each cycle, the robot "falls down" on one of the fore legs with its front body, while the hind part of the body does not move much up and down, which gives the impression of a not very stable gait which also, if it does work that way on the real robot, might not be supportable by the mechanical parts of the real robot.

9 Conclusion

From a modeling point of view the Webots world that was developed during this project provides every possibility needed to simulate the Cheetah robot. With proper calibration, the physics plugin realistically emulates the real robot's leg's behavior, which proved to be the biggest challenge in building the simulation.

As one main goal of the project was to try to characterize the dynamics of the Cheetah robot through parameter searches and trajectory observation of the model, I will concentrate on this point. The experiments showed, that the robot behaved differently with each gait:

For the walking gait it was not possible to produce a trajectory that resembles a walk in either experimental setup. The results of the first experiment suffered from the low

friction and were thus more of a forward sliding motion, while the optimal parameters found in the second experiment resembled more a pace gait, which was a result of the very high friction this time.

The trot gait produced a more convincing trajectory, as the results of both experiments was a trot. On the other hand, it also expressed unrealistic characteristics like excessive up- and down movement of the front body with each cycle. This could be caused by a disadvantageous parameter in the model, however, and should be investigated further.

Opposed to that, the good performance of the pace and bound gaits in both experiments point to another conclusion: Due to its construction, the robot seems to profit from gaits that alternate in putting pressure on each side or on the front and the back. The passive dynamics in the legs would then help in making the transition from one pair of legs to the other, more so than they can for the walk and trot gaits. Furthermore, the model seems to currently support this behavior by having pronounced actuation in the knee joint - which can be seen in the previous experiments -, more so than the real robot, as tests with it showed.

10 Future work

The experiments show that the model used in this project still assumes too much optimality in the physical and mechanical dynamics of the robot. Even though some of the relevant values are taken directly from the construction paper of Cheetah, many are still only approximated. The model would greatly benefit from more extensive calibration experiments, particularly with regard to the knee-joint mechanism, which has lots of unknown parameters that the model somehow needs to combine to a single force. Another parameter discussed here in the report is the friction of the feet on the ground, which proved to have a big impact on the model's behavior. I only presented two extremes, the truth should lie somewhere in between.

I also did not have time to implement the feedback loops into the model. This step would help exploiting the passive dynamics of the robot and should make a significant impact on the speed and stability of locomotion.

11 Acknowledgements

I gratefully thank Yvan Bourquin for providing me with the partially prepared model and all the help with the Webots software, Prof. Auke Jan Ijspeert, Ludovic Righetti and Alexander Spröwitz for offering support, advice and explanation whenever I needed it, Simon Rutishauser for helpful talks about the Cheetah robot and Dr. Alessandro Crespi for technical support with the BIRG computers.

A Appendix

Videos of the experiments will be made available at the project's homepage `http://birg.epfl.ch/page68730.html`.

References

[1] S. Rutishauser, *Cheetah - compliant quadruped robot*, 2008.

[2] Webots, "http://www.cyberbotics.com." Commercial Mobile Robot Simulation Software.

[3] L. Righetti and A. J. Ijspeert, *Pattern generators with sensory feedback for the control of quadruped locomotion*, 2008.

[4] A. Seyfarth, M. Gunther, and R. Blickhan, "Stable operation of an elastic three-segment leg," *Biological Cybernetics*, vol. 84, no. 5, pp. 365–382, 2001.

[5] H. Kimura and Y. Fukuoka, "Biologically inspired adaptive dynamic walking in outdoor enviornment using a self-contained quadruped robot: Tekken2," in *Proceedings IROS 2004*, 2004.

[6] J. Cham, J. Karpick, and M. Cutkosky, "Stride period adaptation of a biomimetic running hexapod," *Int. Journal of Robotics Research*, vol. 23, no. 2, pp. 141–153, 2004.

[7] N. Hatsopoulos, "Coupling the neural and physical dynamics in rhythmic movements," *Neural Comp.*, vol. 8, no. 3, pp. 567–581, 1996.

[8] R. Full and D. Koditscheck, "Templates and anchors: neuromechanical hypotheses of legged locomotion on land.," *Journal of Experimental Biology*, vol. 202, pp. 3325–3332, 1999.

[9] Y. Fukuoka, H. Kimura, and A. Cohen, "Adaptive dynamic walking of a quadruped robot on irregular terrain based on biological concepts," *The International Journal of Robotics Research*, vol. 3–4, pp. 187–202, 2003.

[10] K. Matsuoka, "Mechanisms of frequency and pattern control in the neural rhytm generators," *Biol. Cybern.*, vol. 56, pp. 345–353, 1987.

[11] S. Collins, A. Ruina, R. Tedrake, and M. Wisse, "Efficient bipedal robots based on passive-dynamic walkers," *Science*, vol. 307, no. 5712, pp. 1082–1085, 2005.

[12] W. Schiehlen, "Energy-optimal design of walking machines," *Multibody System Dynamics*, vol. 13, no. 1, pp. 129–141, 2004.

[13] T. Geng, B. Porr, and F. Wörgötter, "Fast biped walking with a sensor-driven neuronal controller and real-time online learning," *The International Journal of Robotics Research (IJRR)*, vol. 25, pp. 243–259, March 2006.

[14] J. Buchli, F. Iida, and A. Ijspeert, "Finding resonance: Adaptive frequency oscillators for dynamic legged locomotion," in *Proceedings of the IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 3903–3909, IEEE, 2006.