
Physically Variable Compliance in Running

Jonathan W. Hurst and Alfred A. Rizzi

The Robotics Institute, Carnegie Mellon University
5000 Forbes Ave, Pittsburgh PA, 15213

Summary. This paper explores the role and utility of variable compliance in running behaviors. While it is well understood in the literature that leg compliance plays an important role in running locomotion, our conjecture is that mechanically adjustable leg compliance improves the efficiency and robustness of a running system. We claim that variable compliance can serve as an effective means for gait regulation, and can enable energy efficient locomotion over a wide range of terrains and speeds. We draw inspiration from a number of observations in biomechanics, and use qualitative arguments and a simulation of our variable compliance actuator to support the underlying ideas.

1 Introduction

Our long-term goal is to develop mechanisms and controllers that allow legged robots to run in a robust and efficient manner, over terrain that varies in geometry and dynamic properties. In this paper, we use a combination of literature review and qualitative and quantitative arguments to describe important aspects for implementing a successful running robot. We begin with a literature review showing that the Spring Loaded Inverted Pendulum (SLIP) model is a good representation of animal running as well as a useful framework for analysis and control. We then go on to assert, through literature and qualitative argument, that leg stiffness is a useful control parameter. Finally, we use quantitative arguments to show that a SLIP is most accurately and efficiently implemented using a mechanical spring, where the stiffness is tuned to a specific gait and terrain.

Our prototype Actuator with Mechanically Adjustable Series Compliance (AMASC) is an implementation of the ideas presented in this paper[1, 2]. A simulation of the AMASC, used as the knee actuation for a simulated running robot, is able to vary leg stiffness to influence its running gait. In addition, when the mechanical spring is tuned to match the desired leg stiffness, the simulation expends the lowest motor shaft power and is most energetically efficient.




Fig. 1. The Spring Loaded Inverted Pendulum, with rotor inertia

2 Background

The SLIP model accurately describes the center of mass (CoM) motion of an animal in a running gait, regardless of the number of legs, the size of the animal, or the running gait employed[3, 4, 5]. Successful running robots also exhibit SLIP model behavior, such as the Planar Hopper, ARL Monopod II and CMU Bowleg Hopper, and most likely were designed specifically to do so[6, 7, 8].

2.1 Tuned Leg Stiffness

In general, carefully choosing leg stiffness is important, for a variety of reasons. Leg stiffness that is too high causes energetically wasteful ground deformations, high stresses in the body, and high acceleration of sensors or other sensitive components. Leg stiffness that is too low results in a long stance time and larger deflection of the leg, possibly reaching compression limitations of the leg. A long stance time is more expensive energetically, because the animal or robot must hold its weight against gravity for a longer period of time.

Tuned leg stiffness is also important for passive stability properties of legged locomotion. According to mathematical analysis of the SLIP model, for certain angles of attack, the spring-mass system becomes self-stabilized if the leg stiffness is properly adjusted and a minimum running speed is exceeded[9]. Similar results were observed in computer simulation of a running cockroach, showing passive stability in its running gait, and an optimal leg spring stiffness for maximum passive stability[10].

2.2 Using Leg Stiffness for Gait Control

Varying the leg stiffness is one good method for controlling a running gait. Only three terms are required to describe a cyclic running gait based on the SLIP model, and leg stiffness affects one or more of the three terms[11]. Control of a running gait using the three different terms (hopping height, leg stiffness, and leg angle at touchdown) has been demonstrated experimentally[12].

Animals vary leg stiffness to control their running gait and alter their center of mass motion, within the constraints of the SLIP model. Researchers seek to better understand the role of compliance control in animal running through experiments limiting one parameter of the SLIP model. By restricting obvious control methods (such as hip angle), more obscure methods (such as leg stiffness) can be observed. Ferris and Farley restricted humans to hopping in place, removing hip angle as a control variable. When subjects were tested at constant ground stiffness but instructed to alter hopping frequency, subjects adjusted leg spring stiffness[3]. Another experiment forced runners to run at stride frequencies higher and lower than normal for a given running speed, thus artificially restricting hip angle, again observing that runners adjust leg stiffness[13].

Most biomechanists claim that animals prefer to change speed by changing stride length (perhaps by controlling hip angle at touchdown), while leg stiffness remains constant[3, 14, 15]. Others claim that step length changes little with speed for normal running on a hard surface, implying stiffness control of the leg[16]. The discrepancy may arise from testing at different running speeds; most experiments involve running at low to moderate speeds, where runners seem to prefer a stride length change over a leg stiffness change. Only in experiments with high-speed running do subjects maintain a relatively consistent stride length over a range of speeds. Common wisdom in the athletic running community is that faster running is achieved by increasing stride length, increasing stride frequency (likely achieved by increasing leg stiffness), or both[17].

2.3 Using Leg Stiffness for Disturbance Rejection

Discussion has thus far focused on control of a running gait on flat, hard floors, but leg stiffness control is also useful when running over terrain with varying stiffness. Without leg stiffness control, the center of mass displacement and ground contact time will change, affecting the gait significantly. When hopping in place, humans adjust leg stiffness to maintain constant global stiffness in response to changes in surface stiffness at a given hopping frequency[3, 18]. During unrestricted running over varying terrain, human runners compensate for ground surface changes by varying leg stiffness, and maintain a relatively consistent center of mass motion[19, 20, 21].

Maintaining a constant center of mass motion is beneficial, despite the fact that changing leg stiffness costs energy. Hopping higher than necessary requires more energy storage in the leg, leading to higher spring restitution losses. Hopping high also increases flight time, allowing trajectory errors to build. Alternatively, hopping too low increases the risk of hitting obstacles and requires faster leg swing recoveries, which may require high power at the hip. Although these reasons are largely speculative, animals do control leg stiffness to maintain a consistent center of mass motion, indicating some benefits. Most current running robots have manually tuned leg stiffness, and are not capable of varying leg stiffness on the fly. However, these robots are not subject to disturbances in ground stiffness; therefore, the reduced efficiency caused by center of mass motion changes is not important. Many values of leg stiffness will result in successful (though not optimal) running, so a robot with no leg stiffness adjustment can still control its forward speed, hopping height, or stride length[6].

3 Implementation on a Robot

In this section, we discuss implementation strategies for robots. Discussion of the merits of specific robot behaviors, such as spring-like stiffness, does not explicitly address the implementation of the behavior. An independent argument must be made for design choices which can affect the performance, efficiency, and ultimately the ability of a machine to exhibit the desired behavior. The implementation of variable stiffness in animals is still debated; co-contraction of antagonistic muscles results in a stiffness change, but neural control of muscles can also result in behavior similar to a spring stiffness.

While the analysis in this paper is generalized and applies to a range of actuation technologies, it presumes the use of electric motors and springs. We also assume that energy efficiency is an important goal.

3.1 Electric Gearmotors and Inelastic Collisions

A simple design for a legged robot would utilize an electric gearmotor at each joint. Several groups have built bipedal robots using this design, and some intended to make the robots run as well as walk. The problem with this approach to running is that most of the kinetic and potential gait energy is lost with each hop to an inelastic collision with the ground.

A gearmotor-actuated running robot is represented in Figure 2(a) as a spring-free, vertically-hopping representation of the SLIP model. The absence of a spring is an important distinction from a spring of infinite stiffness, because it represents the difference between an elastic collision and an inelastic collision. For the model to portray realistic SLIP model behavior, the foot should stick to the ground on impact without chattering, implying an inelastic collision. Spring behavior should be implemented behind the foot, either by the gearmotor or by a physical spring, so there is some stance time during which the robot can apply control forces. Another important distinction is that the reflected rotor inertia, represented by $\frac{I}{r^2}$, does not add to the overall mass of the system. The entire mass is represented by M , while $\frac{I}{r^2}$ may be any arbitrary quantity, perhaps larger than M .

Because the rotor inertia and the robot mass are unrelated, the pin joint of the reflected rotor inertia may be moved to the ground without affecting the model, as shown in Figure 2(b). The rotor begins at rest, and after collision, has some speed that matches that of the mass. If the kinetic energy just prior to impact is represented as T_0 , the rotor inertia is represented as I , the effective gear ratio is r , and M is the robot's total mass (including motor mass), then the energy lost to an impact is:

$$T_{loss} = \frac{I}{Mr^2 + I}T_0$$

and the remaining energy, stored in the downward motion of the robot and the rotation of the motor, is:

$$T_{final} = \frac{Mr^2}{Mr^2 + I}T_0.$$

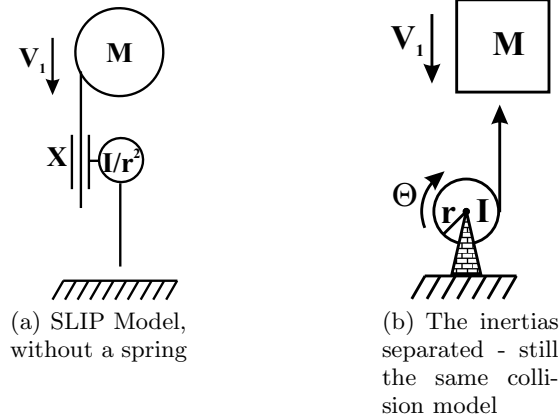


Fig. 2. Figures representing the Spring Loaded Inverted Pendulum model, with the physical spring removed. The inertia of the motor is represented by I , the mass of the robot by M , the gear ratio on the motor by r , the velocity of the robot just before collision by V_1 , and the position of the motor by X , or Θr .

If the effective inertia of the motor rotor ($\frac{I}{r^2}$) is the same as the robot's mass, then half the kinetic energy from flight will be lost to an inelastic collision with the ground. Even after the collision, any remaining energy must be converted through the motor and transmission inefficiencies, which are compounded when energy must pass into the system and then out. In effect, very little energy can be recovered. Most will be lost instantly during collision, much of the rest to motor inefficiencies.

3.2 Improving on the Basic Gearmotor

Minimizing rotor inertia or adding a spring in series between the motor and the leg will reduce the energy lost to inelastic collisions. Both methods are used in force-control applications[22, 23, 24, 25], which are similar in many ways to implementation of a spring rate. A series spring system such as that depicted in Figure 3 has the added benefit of disconnecting the rotor inertia from the leg inertia, such that energy is lost to inelastic collision with only the mass of the leg. Our analysis shows that for a series spring system being used to implement a SLIP, performance and efficiency are maximized when the physical spring stiffness matches the desired spring stiffness.

For a series-elastic actuator with a physical spring rate that does not match the desired total spring rate, the motor must apply torques so the overall system exhibits the desired total spring rate. Given a zero-inertia rotor (the best-case scenario), a proportional controller will behave like a spring, creating two springs in series - a software spring and a physical spring, as shown in Figure 4(a). This is relatively simple to analyze, and provides a conservative estimate of energy use and power output due to the assumption of no inertia. Therefore, further analysis will assume perfect force control of a massless rotor, providing an ideal software spring in series with a physical spring.



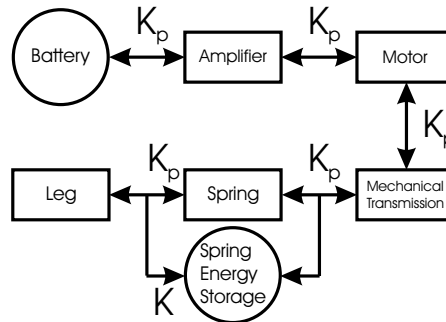


Fig. 3. A SLIP, identical to that shown in Figure 1, but hopping vertically

Power Density of a Series Spring System



(a) SLIP model, with software spring depicted and labelled P, hardware spring labelled K



(b) Energy flow diagram; K_p represents the energy path of stiffness behavior implemented through software control, K represents the mechanical spring stiffness.

Fig. 4.

In a cyclical system such as a hopping SLIP, energy is transferred from external sources (kinetic energy of motion, potential energy of height) to internal sources (physical spring energy or chemical battery energy) and vice-versa, repeatedly. This transfer of energy is represented in Figure 4(b), where the energy may go into and out of the physical spring as an energy storage element (compression and extension), or through the physical spring as merely a power-transmission element (the spring translating, with no deflection). The power output will be divided between the software spring and the physical spring, depending on their stiffnesses.

If the series spring system is deflecting at some rate, the power output attributed to the software spring, P_{K_p} , is

$$P_{K_p} = \frac{K}{K_p + K} P(t),$$

where $P(t)$ is the total power output of both springs in series, K_p is the software proportional gain, and K is the physical spring constant. If the physical spring is perfectly tuned to match the desired stiffness, P becomes infinite, and it can be seen from the equation that the motor (exhibiting the software spring) exerts zero shaft power.

Because springs have higher power density than electric motors, it makes sense to design a system such that the physical spring transfers as high a proportion of the power as possible. A physical spring can have nearly infinite power density, depending on its stiffness; therefore, a comparison between the power density of a spring and that of a motor must be made in the context of an application. Choosing reasonable values for a hopping robot of leg stiffness $K = 5000 \text{ N/m}$, hopping height of $h = 0.25 \text{ m}$, and robot mass of $m = 30 \text{ kg}$, the highest power output during stance is approximately 1 kW (RMS power is 680 W) and the maximum work stored is about 75 J. With an efficient fiberglass spring, such as those used on archery bows which have energy capacity around 1000 J/kg, a 75 g spring can store the required energy and output the desired power. In contrast, a brushless, frameless motor that can output 600 W of continuous power (found through internet search) weighs approximately 2.2 kg, almost 30 times the mass of the spring. Adding the necessary electronics, batteries, and motor housing would add to the mass considerably. It may be that current electric motor technology cannot supply sufficient power density for a running robot. One study on the gearmotor-driven Honda humanoid robot found that the motors would have to be 27-54 times as powerful, without increasing weight, to make the robot run[26].

Energy Efficiency of a Series Spring System

Although the power requirements are the most compelling reasons to use physical springs, energetic efficiency is also improved through the use of tuned physical springs. Given previous assumptions of a perfect software spring and an inertia-free rotor, spring efficiency e_k and motor efficiency e_m , the equation for energy returned is

$$E_{ret} = \frac{P}{2(K+P)} K x^2 e_k + \frac{K}{2(K+P)} P x^2 e_m.$$

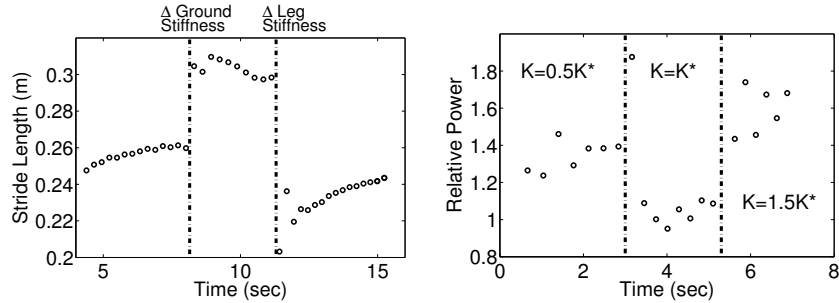
If spring efficiency is higher than motor efficiency, it makes sense for the physical spring stiffness to be as close to the desired spring stiffness as possible. If our assumption of zero rotor inertia is false (and for any real system, it is), then the motor must transmit power to change the momentum of the rotor, and it will expend (and lose) more energy than in this idealized example.

This effect of energy efficiency on a mis-tuned system has been demonstrated on humans in ‘‘Groucho Running.’’ Human subjects were instructed to run with knees bent more than normal, thus consciously decreasing their leg stiffness. Oxygen consumption increased by as much as 50%, indicating a significant increase in energy use[27].

4 Simulation

Based on these ideas of maintaining a tuned physical spring, we designed and built an Actuator with Mechanically Adjustable Series Compliance (AMASC)[1, 2]. We also

created a dynamic simulation of the AMASC, and verified the simulation through comparisons with benchtop tests of the prototype. We then used the simulated AMASC to actuate a simulated bipedal running robot, using Raibert-style controllers. Through testing of the simulation, we show that variable compliance can be used as a method of gait control, and also that a mis-tuned spring requires more shaft power than a properly tuned physical spring.



(a) Stride length vs. time, with changing commanded stiffness K^* (b) relative power vs. time, with changing physical stiffness K . Commanded stiffness, K^* , remained the same.

Fig. 5.

The first test involves a variable ground stiffness, and a controllable leg stiffness. After the robot has reached a steady-state running gait, the ground stiffness is changed. In Figure 5(a), it can be seen that the stride length is affected. After several hops with the altered gait, the leg stiffness is changed, bringing the robot closer to its original stride length. This test is not optimized, and the leg stiffness adjustment is only approximate; but it does show that in this simulated system, which is based on the SLIP model and simulates a physical prototype, the leg stiffness can be used as a method of gait control.

The second test of the simulation involves standard running, with a constant leg stiffness K^* , but a changing mechanical spring stiffness K . While running with a mis-matched mechanical spring, the motor must exert additional power to create the desired knee behavior. When the mechanical spring matches the desired stiffness, power output by the motor is minimized. Assuming some motor inefficiencies, a graph of energy expenditure would follow the graph of the power output.

5 Conclusion

Observations of animal behavior provide insights to possible control for robotic legged locomotion. One observation is that animals vary leg stiffness to control their running gait within the constraints of the SLIP model. Based on arguments put forth in this paper, we believe that robots would also benefit from this strategy. Control of leg stiffness can be used to vary the gait and influence the forward speed, hopping height, or stride length; it can also be used to maintain a consistent CoM motion in

the presence of outside disturbances, such as a ground stiffness change. A properly tuned leg stiffness produces passive stability properties, and likely minimizes energy use. Most importantly, the leg stiffness behavior should be implemented primarily through mechanical springs, rather than through software control of the motors. It is far more energetically efficient, requires significantly lower power from the motors, and will result in lower shock loads to the entire robot.

The final goal is to be able to design and build robots capable of efficient and disturbance-tolerant running on variable terrain, mimicking much of the efficiency, agility, and robustness of animals. To achieve these goals, leg actuators of running robots should have mechanically adjustable series compliance, and the software control algorithms should be designed to appropriately tune the leg compliance in the full variety of running situations.

Acknowledgements

This work is supported in part by an NSF fellowship held by the first author and by The Robotics Institute of Carnegie Mellon University. Many thanks to Joel Chestnutt, for ongoing discussions and for his contribution to the software simulation tool.

References

1. Jonathan W Hurst, Joel E Chestnutt, and Alfred A Rizzi. An actuator with physically adjustable compliance for highly dynamic legged locomotion. In *IEEE International Conference on Robotics and Automation*, 2004. (to appear).
2. Jonathan W Hurst, Joel E Chestnutt, and Alfred A Rizzi. An actuator with mechanically adjustable series compliance. Technical report, Carnegie Mellon Robotics Institute, 2004.
3. Robert J. Full and Claire T. Farley. Musculoskeletal dynamics in rhythmic systems - a comparative approach to legged locomotion. In J. M. Winters and P. E. Crago, editors, *Biomechanics and Neural Control of Posture and Movement*. Springer-Verlag, New York, 2000.
4. R. Blickhan and R. J. Full. Similarity in multilegged locomotion: Bouncing like a monopode. *Journal of Comparative Physiology*, pages 509–517, 1993.
5. Michael H. Dickinson, Claire T. Farley, Robert J. Full, M. A. R. Koehl, Rodger Kram, and Steven Lehman. How animals move: An integrative view. *Science*, 288:100–106, April 2000.
6. Marc Raibert. *Legged Robots That Balance*. MIT Press, Cambridge, Mass., 1986.
7. M. Ahmadi and M. Buehler. The ARL monopod II running robot: Control and energetics. In *IEEE International Conference on Robotics and Automation*, pages 1689–1694, May 1999.
8. Garth Zeglin and H. Benjamin Brown. Control of a bow leg hopping robot. In *Proceedings of the IEEE International Conference on Robotics and Automation*, Leuven, Belgium, May 1998.
9. Andrew Seyfarth, Harmut Geyer, Michael Gunther, and Reinhard Blickhan. A movement criteria for running. *Journal of Biomechanics*, 35:649–655, November 2001.

10. Robert J. Full, Timothy Kubow, John Schmitt, Philip Holmes, and Daniel Koditschek. Quantifying dynamic stability and maneuverability in legged locomotion. *Integrative and Comparative Biology*, 42:149–157, 2002.
11. William J. Schwind and Daniel E. Koditschek. Characterization of monopod equilibrium gaits. In *Proceedings of the 1997 IEEE International Conference on Robotics and Automation*, pages 1986–1992, Albuquerque, New Mexico, April 1997.
12. Jessica K. Hodgins and Marc H. Raibert. Adjusting step length for rough terrain. In *IEEE Transactions on Robotics and Automation*, volume 7, 1991.
13. Claire T. Farley and Octavio Gonzalez. Leg stiffness and stride frequency in human running. *Journal of Biomechanics*, 29(2):181–186, 1995.
14. Thomas A. McMahon and George C. Cheng. The mechanics of running: How does stiffness couple with speed? *Journal of Biomechanics*, 23:65–78, 1990.
15. Claire T. Farley, James Glasheen, and Thomas A. McMahon. Running springs: Speed and animal size. *Journal of Experimental Biology*, pages 71–86, 1993.
16. T. A. McMahon. The role of compliance in mammalian running gaits. *Journal of Experimental Biology*, 115:263–282, 1985.
17. Florence Griffith Joyner, John Hanc, and Jackie Joyner-Kersey. *Running For Dummies*. For Dummies, 1999.
18. Daniel P. Ferris and Claire T. Farley. Interaction of leg stiffness and surface stiffness during human hopping. *The American Physiological Society*, pages 15–22, 1997.
19. Daniel P. Ferris, Kailine Liang, and Claire T. Farley. Runners adjust leg stiffness for their first step on a new running surface. *Journal of Biomechanics*, pages 787–794, 1999.
20. Claire T. Farley, Han H. P. Houdijk, Ciska Van Strien, and Micky Louie. Mechanism of leg stiffness adjustment for hopping on surfaces of different stiffnesses. *The American Physiological Society*, pages 1044–1055, 1998.
21. Daniel P. Ferris, Mickey Louie, and Claire T. Farley. Running in the real world: adjusting leg stiffness for different surfaces. *Proceedings of the Royal Society London*, 1998.
22. W. T. Townsend and J. K. Salisbury. Mechanical bandwidth as a guideline to high-performance manipulator design. In *IEEE International Conference on Robotics and Automation*, volume 3, pages 1390–1395, 1989.
23. Takeo Kanade and Donald Schmitz. Development of CMU direct-drive arm II. Technical Report CMU-RI-TR-85-05, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, March 1985.
24. J. De Schutter. A study of active compliant motion control methods for rigid manipulators based on a generic control scheme. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 1060–1065, 1987.
25. Gill A. Pratt and Matthew M. Williamson. Series elastic actuators. In *IEEE International Conference on Intelligent Robots and Systems*, volume 1, pages 399–406, 1995.
26. Shuuji Kajita, Takashi Nagasaki, Kazuhito Yokoi, Kenji Kaneko, and Kazuo Taie. Running pattern generation for a humanoid robot. In *Proceedings of the 2002 IEEE International Conference on Robotics and Automation*, pages 2755–2761, may 2002.
27. Thomas A. McMahon, Gordon Valiant, and Edward C. Frederick. Groucho running. *Journal of Applied Physiology*, 62:2326–2337, 1987.