PROPOSAL TO NATIONAL SCIENCE FOUNDATION
National Robotics Initiative (NRI "small")
Submitted Dec 11, 2012

Reflex approximation of optimal control for an energy-efficient bipedal walking platform

Andy Ruina, Cornell University

The Ranger robot is 2D, has 3 primary motors and carries no payload. Here we propose to make a robot that is 3D with 10 primary motors and ability to carry a payload. While the methodology of Ranger's control was systematic, it was not formalized. Here we plan to formalize the approach, to extend it, and to make it work with this more complex machine. We think it is a stretch, but within reach. We also propose to fill in some basic theory concerning the essential tradeoffs between energy use, robustness, and versatility.

The submitted proposal summary, and full proposal, are on the following pages.

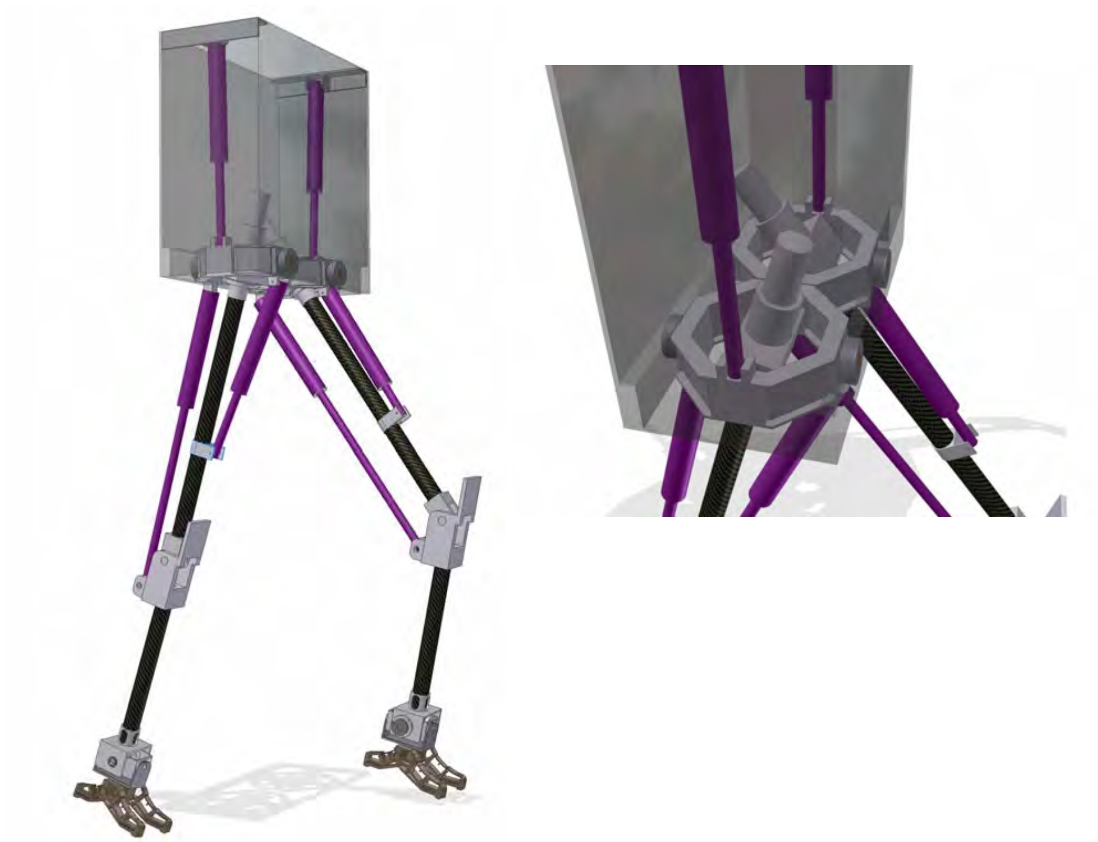

Figure 1: **Preliminary CAD design** Machine with body stub (on left). Details of the hips (on right).

PROJECT SUMMARY. NRI-Small: Reflex approximation of optimal control for an energy-efficient bipedal walking platform — Andy Ruina, Cornell University

The goal of this proposal is to design, build and test a two-legged walking platform and its controller. This will be the first autonomous bipedal robot platform that is plausibly usable for assistive robotics. If successful, this platform could replace wheel and tread based systems used and proposed for service robotics. The problem is to make a machine that is both efficient and robust.

Presently, the most energy effective walking robot is Ranger from the PI's lab at Cornell. Scaled for weight and speed, Ranger matches a human's energy use. But all it does is walk straight on flat ground. The most robust and versatile bipedal walking robots are Boston Dynamics' hydraulic Petman and Atlas and Honda's electric Asimo. However, scaled for weight and speed, these robots use 10 or more times the energy that a human uses. A practical robot needs to match, or improve on, the efficiency of Ranger and on the robustness of Petman and Asimo. It needs to function reliably without massive batteries, hydraulic hoses, or gas engines.

There are two big obstacles to simultaneously achieving robustness and efficiency in a bipedal robot: First, there is no proven approach (either in practice or theory) to controlling such non-linear systems. Walking requires coordinating many moving parts, with more ways to move than there are motors. This motion also involves intermittent contact and complex governing equations that change during the motion. In technical language, a walking robot has many degrees of freedom, is under-actuated, is non-linear and has intermittent hybrid dynamics. Secondly, the basic nature of walking, how humans and other animals walk, itself is not well understood.

The solution we propose is to use a biologically inspired, although not strictly bio-mimetic, reflex-based control. Such a control can be simple because of features of legged linkages. For instance, the motions of one part may have a small influence on many of the other parts. Also, crisp instantaneous multidimensional control is not needed for robustness because the essential instability, balancing the center of mass, is low dimensional. The resulting small number of parameters in the reflex approach makes the simultaneous optimization, of both energy use and stability, practical.

This research aims to bridge the gap between the opposite extremes of no control (passive-dynamics) and high-bandwidth high-gain control. Passive-dynamic machines that are uncontrolled have had no versatility and little stability. But high bandwidth high-gain control is energy hungry. Between these extremes, we propose to use an optimized reflex-based control architecture in which events (collisions or thresholds) trigger smooth control actions. Such smooth implementation of discrete control allows optimization of both energy use and robustness while minimizing need for massive on-line computation, information flow, and memory. The outcome of the project will be a tested functional hardware design and a rational proven-in-practice control concept.

Intellectual merit. The contributions fall into two categories, both of which provide scientific *foundations* from which engineering or therapeutic *methods and tools* should eventually follow. *1) Advances in robotics,* namely the development of a low-bandwidth, low-computation approach to smooth low-energy robust control of legged locomotion; and *2) Understanding of walking.* Walking is a basic human activity which deserves scientific understanding. Basic investigations into issues of robot design and control have led the PI in the past, and should lead again during this research, to basic insights into the nature of legged locomotion.

Broader impact. The basic scientific insights into walking should eventually aid the diagnosis and correction of human walking disabilities, as well as the future design of robots that walk and move in other complex ways. Robot presentations at schools are inspiring to students. Cornell undergraduate students (about 10-25 each year), including women and minorities, are drawn to the subject. They have and will continue to participate in the research in the PI's lab. Some go on to careers related to robotics, prosthetics and orthotics. The broader public is also educated and inspired by media exposure and public demonstrations.

Keywords: bipedal walking robot; reflex-based control; low-energy; robust-control.

NRI-Small: Reflex approximation of optimal control for an energy-efficient bipedal walking platform

Andy Ruina, Cornell University

PROJECT DESCRIPTION

1 Track record

The PI (Ruina) and his lab have experience with the design and testing of autonomous electrical robots [25, 21, 19, 66, 65]. Most notable recently has been the development of the Ranger robot which uses far less power, scaled by weight and speed, than any other autonomous robot. Ranger walked 65 km in an indoor gym without recharge or human touch [11, 12, 9, 10, 8]. Ranger and its control were developed as a step towards the present proposed research. The PI has also developed biomechanics and locomotion theory related to the balance and energetics of locomotion[26, 25, 29, 58, 20, 1, 64, 63, 5, 15, 46, 47, 40, 6, 28]. The external funding for this locomotion and robotics research has been three 3-year NSF grants and their extensions:

1. *Passive, Nonlinear-Dynamic Study of Walking: Simulation, Analysis, and Experiment*, 7/1/98–7/31/03, for a total of $283,515 (total includes an REU supplement and a $10K research supplement), BES 9806612.
2. *An energy-efficient bipedal walking robot*, 8/1/2004–7/31/2008, $364,000, IIS 0413139.
3. *RI: Robust implementation of foot placement for balance of 3D bipedal walking*, 8/15/07–7/31/12, for a total of $853,075 (total includes REU supplements in the amount of $6,000 and $24,750), IIS 0705390.

2 Introduction: the needs for a platform, for understanding and for a control architecture

In the same way that humanoids developed for human interaction aim to evoke the geometry, look and motions of human arms, hands and faces, such service robots should, perhaps, also walk like humans[41]. Also, if service robots are to mimic or extend human abilities they should also have capacities for rough terrain that might, ultimately, best be achieved with legs. But the mechanics and control of walking is not yet well-enough understood to lead to systematic design of effective humanoid walkers, nor for engineering-level understanding of human walking. For example, there are, as yet, no theories that reasonably predict what actuator strengths and what sensing speeds and accuracies are needed by a human, or robot, to achieve a given efficiency or robustness of walking. To address these needs the project has two fronts: 1) development of a walking platform; and 2) development of walking control theory. Each of these depend fully on the other.

1) Platform. The engineering development aspect of this research is to design, build, test and document a 12 kg, 1m waist-height, energetically-autonomous electrical bipedal walking robot, up to the waist. A support platform sits atop these legs. That platform is to carry an upper body (arms, head, and trunk) payload of up to 28 kg. The upper body is not part of this proposal. The functional goal is to match the robustness and efficiency of a Segway [48, 59] but with the ability to largely match the functionality of Petman [23], Atlas or Asimo. The key novelty is to achieve this with one twenty-fifth the energy use, with no power cords, and with no possibility of dangerous mule-kicks. The target is a fall frequency of effectively zero for ground that matches typical sidewalks in smoothness, a maximum speed of at least 2 m/s, the ability to stand, sit and climb (known) stairs, and a total energy use of 100 watts. The platform plus humanoid upper body payload should be able to walk up to 2 hours and 12 km for each kg of batteries, yielding a total specific energetic cost of transport of about 0.15. This work will be coordinated by the PI with a professional engineer, and carried out with large efforts from undergraduate and Master of Engineering students.

2) Theory. The primary scientific advancements will come from graduate students working with the PI. The primary intellectual approach is to advance the understanding of walking by trying to re-invent it. Although we marvel at nature's wondrous solutions, understanding these solutions is impossible without understanding nature's problems. But this marveling is premature as there are essentially no quantitative models for the difficulties of legged locomotion. We plan to advance this understanding by quantifying the effects that various given design strategies, sensing capabilities, actuator limits, reflex-loop-time delays and computational limitations have on the ability to achieve reliable and efficient walking.

The main general questions are 1) What are the limits on walking performance (efficiency, robustness, versatility) for given mechanical and electrical hardware? and 2) What is a controller design which approaches these limits? With such models one has metrics for performance of a given design, animal or machine. These metrics can be tools for the design of control. On the practical side, this basic science should aid both the design of robots and eventually be useful for those attempting to understand human disabilities and their remediations. Although we have in hand a reflex based approach that we think, based on its performance with Ranger, can lead good performance by these metrics, it has not been formalized or generalized for a complex machine. The theoretical models to be developed are analogous to the simple energetics and stability results from our past research [26, 6, 1, 64]. These will aid the formalization of the reflex approach and provide development of metrics by which to evaluate it. The ultimate test of the theories will be their ability to aid the design of predictably robust hardware.

The physical robot platform development provides a motivation and focus for the theory, a debugging environment, and an eventual means for demonstrating the control-design approach. It provides an open-source platform for use by others with mission-centric upper body designs, and a demonstration for public outreach. The theory development depends on the specific hardware implementation, and the hardware development needs the new control to function.

The grand hypothesis in this proposal is this: *It is possible, using available hardware, to make a walking robot that is as robust as the best present-day robots, but uses one tenth the power.* We will either prove the hypothesis by example, or disprove it by means of developed theory. In the latter case, the best possible hardware will be developed anyway.

3 Approaches to the control of walking

Our plans are informed by the promises and limitations of existent ideas. We describe some of these now, in the context of these questions:

1. How can one control legged locomotion, including bipedal walking, so as to be simultaneously versatile, robust, safe and energy-stingy? and
2. What hardware requirements (e.g. sensors, motors, stiffness, ...) are necessary to do this?

ZMP. Walking is, in one view, standing in motion. While standing, undesired body tip is primarily corrected with ankle torques (*e.g.,* [50]). These ankle torques also control the center of pressure (or the Zero Moment Point) acting on the bottom of the feet. In the ZMP approach to walking, the center of mass is held in place over the middle of the feet by the center of pressure chasing it in. The ankle torques move the ZMP. Like a sheep dog, the ZMP can corral the center of mass towards the center of the foot (or feet) polygon (*e.g.,* [72]).

The ZMP approach is limited in that it can't address walking on stilts or on curved rails (used on some prosthetic shoes [2]). People can do both of these well, presumably with a control approach similar to what they use for normal walking. In practice, ZMP robots, like Honda's Asimo series [36], have characteristic attributes: they walk with bent legs (this allows the controllers to have authority over all of the upper body degrees of freedom at all times); they have large flat-bottomed feet; they walk smoothly (with the center of mass moving on a nearly straight line); all joint angles are carefully controlled at all times; they use a lot of power; and they often work. The best ZMP robots seem, based on publicly released information, to have much of the functionality we seek, but with 10-20 times too much energy use. There is no evidence that the needed reduction in energy use is possible with the ZMP approach.

Powered passive dynamics. McGeer noted that the Wright brothers powered flight followed easily just after their mastery of gliding [45]. They just added an engine. Following this paradigm we added power the simplest way we could [66] to a successful downhill purely passive-dynamic machine [35]. The resulting robot walked with little energy use but was fragile. The flight analogy is more subtle. Although the Wright gliders didn't have motors, they were controlled by pilots. And that control worked with the addition of a motor. A bit of (ironically) bad luck made it possible for us to build uncontrolled passive robots that were slightly stable: adding a motor to such a marginally stable passive robot yields a powered robot with the same marginal stability. That passive-dynamic robots walking looks like humans walking seems to imply that the two are related. But it seems to be a mistake to think that one is just a small perturbation of the other. Certainly that is not the evolutionary history; advanced nerve systems far preceded the upright human bipedal posture. Rather, it seems, human walking is highly controlled in a manner that uses small actuation forces for much of the walking cycle, resulting in a gait visually similar to the passive gaits, but with fundamentally different stability.

Intrinsic plant stability, while inspiring, now does *not* seem critical for walking robots. The machine does not need to be stable. Absent controls, a robot only needs to fall slowly enough for the control system to catch. Mastering passive dynamics helps understanding of the mechanics of walking, but gives little guidance for how to use power to increase robustness. While we have advocated passive-dynamics as a basis for control in the past, our present approach is almost the opposite. It is to use massive control and then as best possible, to titrate it out.

***Ad hoc* control.** The hopping and running robots from Raibert's Leg Lab at MIT were controlled using insights about one-legged hopping (*e.g.,* [56]), not an abstract control theory. Raibert noted that there were three main things to control (hop height, pitch and speed) and that these need only be controlled step to step. Further, he conceived simple actions that affected each of these things almost independently, in a sense diagonalizing the control problem: hop height was regulated by push-off in the second half of the bounce phase, hip orientation was regulated by torques on the leg during stance, forward speed was regulated by the leg position at landing. The Optimally Scaled Hip-Force (OSHF) approach of [71], used successfully for a tethered cheetah robot, seems to be similar to the Raibert approach, but with the ground reaction shared over 4 legs. These running controls are examples of the reflex control proposed here. Unfortunately Raibert's running control scheme is *ad hoc* and not directly generalizable to other problems like walking. Similarly, also in the MIT Leg Lab, Jerry Pratt honestly called his control system 'intuitive control' [53, 55]. Pratt's intuitive control worked well for control of a 2D robot on a boom and, in simulation, for a 3D robot [54]. In a sense, our control design is a partial automation of the Raibert/Pratt intuitive process. The resulting control architecture is similar.

Fast and powerful actuation. Raibert's robots at MIT and again at Boston Dynamics (BigDog, Petman) have been designed to function, not to save energy. But does robust robotic locomotion really need to use much more power than the animals they mimic? Does the stability of Petman or BigDog depend on the large control authority of high-pressure hydraulics? We contend, using animals as inspiration, that robust control of a horse-sized robot need not use 10-20 horsepower.

Optimal trajectory control. Given a model of a robot one can seek controls that minimize a given cost, for example energy use per distance travelled. A calculated optimal trajectory gives functions of time (a trajectory and time history of controls) as output. If performed numerically, these functions are discretized using one or another scheme [7, 14]. Thus the output is a set of numbers: *e.g.,* coefficients or weights of a Fourier series, a polynomial, or of a spline representation. These coefficients represent one or more functions of one independent variable (*e.g.,* time). Optimal trajectory control is not concerned with stability. And the stability of a given trajectory is different for different representations of a given optimal solution. For example, a given trajectory will have different stability depending on whether it uses time or, instead, one of the dynamical state variables as the independent variable in the function representation. Westervelt, Grizzle and Chevallereau

[73, 74, 33, 17] have implemented optimal trajectory control on 2D and 3D robot simulations, parameterizing the control in terms of the ankle angle(s) (rather than time), and optimized for an idealized energy cost (see 'zero dynamics' below). Stability can then be obtained by with an additional feedback layer [34, 31].

Our approach here explicitly uses optimal trajectory control, but with reflex actions which are a non-explicit paramaterization. In our implementation the optimal trajectory is not generated, but only shows when the control is implemented on a model of the plant. As in[34, 31], stabilization is an added layer. In our reflex based formalism the stabilization uses the same encoding as the trajectory optimization, so the trajectory and it's controller, although designed separately, are not visibly separate in the final implementation. This way we have no need for detailed representation of the trajectory, no need for fast control, and a unified description of the trajectory and its stabilizer.

Optimal feedback control encompasses optimal trajectory control. In optimal feedback control the optimal trajectories are found from all possible starting points. At each starting point there is a set of control values (*e.g.,* torques and currents). Thus if there are n controls and the system dynamical state is m-dimensional the result of the optimization calculation is n functions each of m variables. Thus, in action, the optimal feedback controller measures the dynamical state, calculates the controls, and feeds them back. Optimal feedback control fully encompasses stability as the system is always guided to the goal. Further, the rate of convergence to the goal, and resistance to disturbances, can be included as part of the cost function in the optimization. Although extremely promising as a general descriptor of the motor control in animals [70] this approach has not yet been used successfully on a walking robot. And it might be more than is needed. For the most part, a walking robot must deal with nearly periodic trajectories. Thus the full optimal feedback control might be adequately represented by an optimal trajectory with local rules near that trajectory. This is done somewhat formally in the LQR trees approach [69] which we do not pursue because of its reliance on a possibly high-gain continuous feedback loop, because it has not been tested on a legged robot, and because it seems too complex for the immediate needs. Our approach here, as in Ranger, uses reflex loops as a simple approximate discretization of an optimal feedback policy. As judged from behavior, this approximate optimal control may have features of optimal control.

Representation with neural nets and learning via *e.g.,* evolutionary algorithms. An optimal feedback controller maps dynamical states to control values. In practice, such maps (functions) need a finite-dimensional representation. One such representation uses weights on a neural net. Then, rather than using formal optimization, one or another learning or heuristic optimization methods can be used to tune the weights and topology to better achieve some optimization criterion (*e.g.,* [76]). One problem with physical implementation is that learning (training) generally involves falling and most bipedal robots are not built to withstand hundreds or thousands of falls.

The small RunBot robots used neural nets and learning (*e.g.,* [44]) to make robust control of two-dimensional walking. Perhaps, however, success was partially through hand tuning by competent researchers. There is no doubt that live learning will eventually be a key part of locomotion control.

Our contention is that neural nets are too general purpose to be most information-efficient in representing motor control. Rather, we believe, that our minimalistic representation of coordination may provide a better substrate for learning and evolutionary algorithms.

CPGs. Many animals have neural circuits that, without muscles, generate patterns similar to the gaits of those animals, and there is some evidence of such in humans as well [22]. These are called central pattern generators (CPGs). One view of locomotion that is philosophically close to passive dynamics, is to view gait as emergent behavior from the mechanics and neural circuits, a kind of coupled oscillation. CPG-based robots have worked in 2D walking simulations [67, 68] and in 3D [57]. But we do not anticipate that such a robot will generally have much more stability than a passive-dynamic robot (unless the CPG is supplemented with substantial feedback).

There are, however, more subtle ways of thinking about CPGs. For example, a finite state machine that ramps up sensitivity to a state change sensor, and then eventually transitions with or without sensor input, would behave as a CPG, running through states corresponding to the phases of gait. Still another overlapping interpretation of a CPG is as an internal model for, in controls language, model-based state estimation. In the presence of noisy sensors, system state is better determined with the help of a model, and what is observed as a CPG may be, in one way of thinking, just such a model.

In that are controller uses a finite state machine, and that within each state a feedforward control is used, our control is only distinguishable from a CPG in that we have not yet implemented state transitions based on passage of time after an expected sensor input is late.

'Robust control' has been adopted by control theorists to mean control which has been optimized for performance given uncertainty in modeling parameters, sensor properties or environmental forcing. Some attempt has been made to apply standard robustness measures of stability to walking (the 'gait-sensitivity norm' [38]). Certainly any functioning walking robot controller has to plan for uncertainty whether using known formal 'robust control' methods or not. In a domain as poorly understood as walking is now, robustness has to include robustness not just to unknown parameters, but to unknown model features (e.g., joint play, motor friction) and unknown degrees of freedom (e.g. excited vibrations) and thus must adhere to heuristics like KISS (keep it simple) perhaps more than to classical linear robust-control methods. Nonetheless, our methodology includes philosophical alignment with "robust control" in that control parameters will be optimized for insensitivity to modeling errors, sensor noise and physical disturbances.

Dimension reduction via zero dynamics (HZD). One way to reduce the degrees of freedom is to make all joint angles track pre-determined functions of the ankle angle (in 2D) [32]. Successful tracking thus reduces the robot to a 1 DOF mechanism This approach is somewhat opposite of the biological idea of the 'unconstrained manifold' [42] wherein inessential degrees of freedom are barely controlled (rather than strictly controlled). When applied to 3-dimensional walking [16] there are 2 uncontrolled degrees of freedom (assuming scrub torques prevent rotation about a vertical axis) and when combined with series-elastic actuators there can be more. However, even without the proofs of stability afforded by the 1D reductions originally proposed, additional discrete control can be added and be highly functional [34, 34, 62, 49].

In this sense, the architecture we propose here overlaps with the recent HZD approaches. The key differences are that 1) we use indirect description of trajectories, rather than direct tracking, and 2) we implicitly ignore as many degrees of freedom as possible, rather than of constraining them away. The modern HZD approach, including a higher-level discrete control, overlaps with our approach and is a viable competitor. At present no HZD robot has had close to the low energy use of Ranger. However, so far no reflex-based robot has had close to the robustness of the HZD robot MABEL. We have chosen our reflex-based approach over HZD because of its intrinsic compliance, its lack of need for tight feedback looks, and its low-information implementation.

Dimension reduction via SLIP. Running animals have center of mass motion that is close to that of a point-mass on a spring, like a pogo-stick: a so-called Spring-Loaded Inverted Pendulum (*e.g.,* [13]). This *description* might be accurate because animals have bodies bigger than their legs, and they can push down with their feet harder than they can kick to the side. Thus the motion of most animals and of most biomimetic robots cannot be far from that of a point mass with an axial actuator. So it has to be a reasonable description for most motions. However, the point-mass model has also been proposed as a robotic *prescription* [24]. In this SLIP control the internal degrees of freedom are controlled so that the center of mass would follow the trajectory predicted by the spring-mass model. This approach is different in detail from the zero-dynamics control approach above. But both methods do dimension reduction by quickly controlling internal degrees of freedom to accurately track a simple goal. As discussed further below, we subscribe to the descriptive aspects of SLIP, but not the prescriptive aspects. As a prescription SLIP misses opportunities for saving energy while spending control effort (energy and computation) to control degrees of freedom that have little effect on balance.

On the other hand, point mass models have been, and will be fruitful for us for testing and developing basic models of control. That is, we expect that the various basic trade-offs between power and stability will best be illuminated by use of such simple models that we plan to develop.

Dimension reduction through decoupling. The various parts of a complex robot are not necessarily all deeply mechanically coupled to each other. For example, the motion of the foot on a swing leg has little effect on the motions of any other body parts. And side-to-side motion control might be decoupled from fore-aft control [3]. So the control of the foot, when off the ground, can safely be decoupled from that of the rest of the robot. To greater or lesser extent, this idea can reduce the, say, 10 DOF problem to a collection of problems each of which has, say, no more than 2-3 DOFs. This is somewhat opposite of, for example, the formal approach to dimension reduction of Gregg and Spong [30] which effectively add virtual constraints to simplify the control. Our approach is to explicitly exploit this decoupled feature of legged locomotion. That is, the sparsity of the mechanical coupling will be used to generate sparsity in the control.

Capture points and capture regions. While one foot is on the ground a person is an inverted pendulum of sorts. Kajita's 'linear inverted pendulum' [39] model for walking has linear governing equations in 2D, and no collisions, so is particularly tractable. Pratt [52] exploits this model to consider capture points and capture regions. Where, he asks, should a robot place its foot so as to end up standing upright and still at the next step? Capture points are said to be part of the control in Petman and Atlas. The recent success of these robots suggests that this is a functional concept. Various ideas related to capture points are unavoidable the same way SLIP is unavoidable. Whether one pays attention to them or not, these point-mass ideas must be reasonably accurately descriptive, because they describe the actual (for legged people and bipedal robots) near-decoupling of the center of mass motion and foot placement from other degrees of freedom.

We believe that our proposed approach, to optimize discrete feedback, will encompass the capture point idea in that any reasonable sensitivity analysis of a complex robot will have to discover the key role of foot placement. For example, the control design of Ranger had no explicit use of capture points for foot placement. Nonetheless, we found in control development, that issues related to foot placement were necessarily paramount. Furthermore, as stated, point mass models will always be used for test of concept and generation of theory.

Reflex control. Tuned reflexes constitute a form of feed-forward predictive control. We believe that this is an effective approach to motor control, especially in the context of delays and essential nonlinearities such as play and friction. In 2D it has been shown that a bio-mimetic simulation can walk stably using only reflex loops that are known to exist in people [27]. Similar ideas have been used in 2D and 3D by Hartman and others [61, 43] in the context of evolutionary robotics where there is only one reflex trigger per step. Also related is the computer graphics approach in SIMBICON [75] where the motion is generated by just a few triggered switches between simple compliant controllers.

As opposed to [27] our reflexes will be 3D, not 2D, and, as in [9] they will be designed for our particular machine rather than directly mimicking known human reflexes. Finally, our reflexes will be designed (optimized), as described later, for both minimization of energy and maximization of reliability (robustness/stability).

Better actuators and biomimetic actuators. To date, people and animals walk and run much better than robots while generally using less energy. Animals have muscles and robots have electric motors or pressure-driven pistons. So, one might reason, robotics could be much better if engineers had actuators with the 'amazing properties' of muscles. Thus there has been a search for artificial muscles or for actuators that mimic properties of muscles (*e.g.,* [4, 18, 51] and recent non-NSF RFPs). We question this reasoning. It is true that Raibert has done well with powerful actuators. But the theory of locomotion control is still too primitive to attribute his successes to his actuators. Nevertheless, the motors and transmissions we propose to use far surpass the peak muscle power of 200 Watts/kg and peak muscle tension of about 20 N/cm^2. We agree with [60] that electric motors are adequate for the task, although we don't feel a need for custom motors.

4 Proposed biped platform

Above we have placed our proposed work in context. From here on we describe details of our approach. First we describe aspects of the proposed work related to this NSF NRI call to *establishing a common hardware platform* by *creating a testbed* that has *functional capabilities.* This engineering aspect of the proposed work not just for the functional goal of a platform for use by others, but also for the development, testing and demonstration of the more theoretical control-design concepts.

Target capabilities. We plan to develop behaviors including: walking at speeds up to 2 m/s with total specific cost of transport (TCOT) = 0.2 at 1.5 m/s; standing in one place; climbing stairs and curbs (ie, via single step control to pre-destined foot-placements); robust recovery from side forces via high-speed foot placement at any time during most tasks; sitting and getting up.

Hardware considerations

Reliable hardware depends on attention to detail in design and construction detail. Ultimate success depends on improvement based on testing and simulation and testing again. And again. When Ranger first walked it had a passive-dynamic looking gait and a TCOT of 1.6. Through incremental design improvements, control optimization, and electronics debugging facilitated by testing, simulations, and comparisons between the two, Ranger was made to walk more robustly and with a TCOT of 0.19, a factor of 8+ reduction in energy use. Ranger's development took 5 years even though version 2 was a functioning walking robot after only five months. Similarly here we propose to have a working machine in 2 years, with 3 years of refinement.

Note that the majority of Ranger's improvement was not due to the addition of specialized energy-efficiency mechanisms like clutches, variable-ratio transmissions, or adjustable-rate springs. Rather, simulation and testing, led to control and basic design improvements that were greater than one could even hope to come from special transmissions, etc. In contrast, the Collins biped did use of specialized energy-saving mechanisms, but it was far from robust. Ranger, by contrast, eventually achieved a somewhat better TCOT, while having the robustness to walk 186,000 steps without a fall. It's all in the control, enabled by robust and consistent, rather than ingenious, hardware. Another way to say this is that legged robots are far from the theoretical minimum of energy use, which is zero for level ground legged locomotion [29]. So robot designers are far from needing to make use of the factor of two that humans get from elastic energy recovery in running. The added complexity of control design, and the loss of controllability, makes it so inappropriately placed springs can *increase* energetic cost [61, 43].

Joint play, joint friction and overall compliance in the leg are to be minimized (failure to attend to such may be one reason for the lack of success of all three copies of the Delft Tulip Robot [37] — private communication w/ Tomas de Boer, Oct 2011). We will not use tight feedback loops which could cause chatter due to excitation of unmodelled degrees of freedom.

Mechanical design

The basic bipedal platform is with two legs and a robotic trunk. Each leg will have 6 degrees of freedom with extended angular range to allow sitting, stair climbing and large-disturbance rejection. In total there will be 5 controlled degrees of freedom for each leg, 10 in total. The upper body is a trunk stub for carrying a payload (head, arms, batteries, high-level sensing). The target weight is 12 kg with a 28 kg payload. The leg length is 80 cm.

Motors. Electric motors were chosen over hydraulics or pneumatics because they are far more efficient, easier to manage and we have in-house know-how. Each powered joint will use a Maxon 4-pole 200W brushless motor. However, based on data sheet values, adding water jackets to the motors and a miniature pump and radiator to the system could permit sustained motor power outputs of over 1000 watts, enabling stair climbing and possibly running. Motor characterization through bench testing is key for use in modeling and optimization. We will use custom GaN FET-based motor drivers allowing higher efficiency, at both low and high output powers, than easily available silicon-based MOSFET motor drivers.

Transmissions will use ball screws, planetary gears, harmonic drives, chains, and stainless steel flexible cable, depending on the joint, to deliver potentially large motor powers to the joints. Driving the knees and ankles through cables and linkages will facilitate fast leg swing.

Hip joint. This is the most difficult aspect of the mechanical design. We plan for a pair of gimbals, one for each leg, driven by a cable differential. Cable differentials for the hip gimbals will use Sava Cable 7x49 nylon-coated cable; each cable has 343 stainless steel wires, giving them the flexibility needed to work reliably around compact pulleys (such cables worked for $\approx 5 \times 10^5$ cycles on Ranger). Detailed design, testing and development is part of the proposed work.

Knee. A motor-driven ball screw running down the thigh will apply force to a carbon-fiber linkage at the knee. The actuation lever arm at the knee is at an angle such that the effective gear reduction ratio is low for when the knee is nearly straight, and higher when the knee is highly flexed (the opposite of the kinematic locking advocated by some). This decreases the energy needs for swing during normal walking, but allows more torque for squatting, sitting, etc.

Ankle. Ankle torques for powering pushoff and for fore-aft balance when standing will be provided by a motor-driven ball screw running down the thigh actuates a carbon-fiber linkage running from the knee to the ankle, arranged such that the pivot point of the linkage element running down to the heel is in line with the knee axis, and thus is minimally coupled to the knee flex. The desired degree of coupling between ankle pushoff force and knee flex will be refined by optimization in modeling. Lateral ankle flex will be passive.

Foot. Whether to use a curved-rail foot design as in Petman, or a design that has more contact for control of steering at one foot, will be settled by optimization.

Electronics design Rangers electronics were designed for porting to this more complex robot. In brief, we will use a network of ARM7 microcontrollers, one for each controlled degree of freedom (and some for major sensors) communicating over four separate CAN buses with a high-level processor used for sensor fusion and dynamical state estimation.

Sensing Each degree of freedom will have 7 sensors (motor angle, joint angle, joint rate, joint torque and motor voltage, current and temperature). These are in addition to at least 6 inertial measurement (IMU) channels, foot contact sensors, battery voltage and any navigation aids (e.g., GPS). Force/torque sensing at all actuators will use a combination of off-the-shelf miniature load cells and displacement sensors on integrated elastic elements (on Ranger a $1 optical displacement sensor measured foot force reliably). Angle sensors at all joints, will be Netzer Precision high-resolution capacitive absolute angle sensors, giving both angle and angular velocity. Whether or not a more expensive fiber-optic laser IMU is needed will be studied in the modeling phase.

4.1 Hardware overview

We do not believe that legged robotics is primarily now held back by the lack of performance, or inappropriate properties, of available engineered structures, springs, transmissions and motors. Thus we have no radical mechanical nor electrical plans. Nonetheless, legged robotics *has* often been held back by the lack of appropriateness of the hardware design for the task. With so many people trying to make walking robots it is disappointing that none are adequate for our research, or for the needed platform. But all known platforms are too energy hungry (e.g. hydraulic or using highly-reduced electric motors) or not able to be modeled accurately (e.g., sloppy joints) or not sufficiently versatile (no ability to steer).

Control architecture

The overall architectural plan here was developed on Ranger for the purpose of extension to the proposed robot. The basic scientific hypothesis to be tested with this proposal is this: *the reflex-based control described here can, with the control-development procedure outlined later, lead to unprecedented legged locomotion performance.*

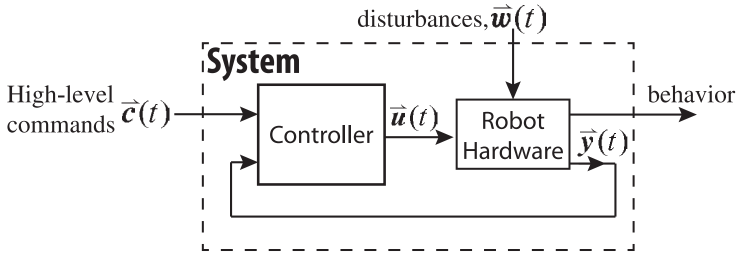

Figure 1: **The legged platform System.** High level commands $\vec{c}(t)$, e.g. to walk at a specified speed, are processed by the software and hardware which, with disturbances $\vec{w}(t)$, leads to (hopefully stable) walking behavior. The inner blocks are expanded in later figures.

We start with a high-level description. The controller determines the high-level commands and sensor data to the actuator outputs given , as shown in Fig. 1. The control 'architecture' is the algorithmic form of this map. In a later section we will how we determine particular values for free parameters within this form. For a linear (continuous, autonomous, observable, controllable, noise and disturbance free) linear system we could describe a reasonable control architecture as

$$\underbrace{\vec{u} = K(\vec{x}^{ref} - \vec{x});}_{\text{control}} \qquad \text{and} \qquad \underbrace{\dot{\vec{x}} = A\vec{x} + B\vec{u} + L(\vec{y} - C\vec{x})}_{\text{estimation}} \tag{1}$$

where $\vec{x}$ is the estimate of the present system dynamical state (sometimes called $\hat{x}$), $\vec{y}$ is the sensor values, matrices A, B and C represent a model of the system and its sensors and L is the estimator gain. K is the gain matrix which we must design. Generally in linear control it is assumed that the best model estimation (A, B, C and L) is independent of the best choice of gains K. To keep this somewhat general, $\vec{x}$ could include estimates about the environment (for example, the value of a slowly varying disturbance force). Leaving aside model estimation, we would then say the linear control 'architecture' is $\vec{u} = K\vec{x}$. The implementation of the control would be the determination of the entries of the K matrix.

Assuming separability of estimation and control, a general non-linear control architecture is

$$\underbrace{\vec{u} = \vec{f}(\vec{x}, \vec{x}^{ref})}_{\text{control}} \qquad \text{and} \qquad \underbrace{\dot{\vec{x}} = \vec{g}(\vec{x}, \vec{y}, \vec{u})}_{\text{estimation}}. \tag{2}$$

For our hybrid system we also need to account for jumps in $\vec{g}$ (that is, either $\vec{g}$ has delta functions in it or the flow of the estimator has to be described as having separate continuous and discrete phases).

A more general form that captures both control and estimation without separation is

$$\vec{u}(t) = \vec{\mathcal{F}}\left(\vec{y}(t'), -\infty < t' \leq t\right) \tag{3}$$

where $\mathcal{F}$ represents a causal functional of the sensor data history $\vec{y}(t)$. The implementation of $\vec{\mathcal{F}}$ need not contain an explicit dynamical-state estimate $\vec{x}$. Again, however, even if one could find, say, various series approximations to such a general functional, the utility of any given form depends on it's efficiency, as defined above. We add this form because it more easily encompasses approaches that don't so carefully separate the control from the estimation. That is, ultimately the controller maps sensor to actuators and it is not clear that the best way to do that is by means of an explicit internal model.

Efficiency of control representation. It is momentarily liberating to realize that the most general thing a controller can do is evaluate a single function: give the controller the present system dynamical state and it evaluates a single function of several variables to calculate the

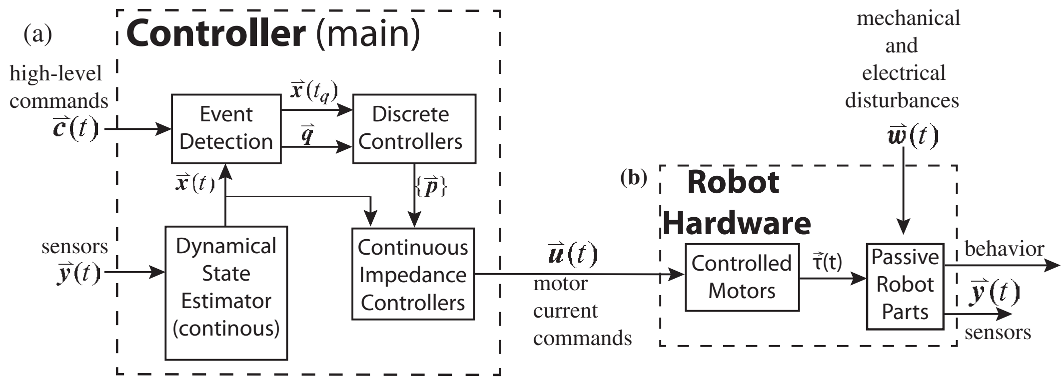

Figure 2: **(a)** The **Controller**, in software, takes high-level commands $\vec{c}(t)$ and sensor data $\vec{y}(t)$ and determines logic-level motor current commands $\vec{u}(t)$. **(b)** The robot hardware takes in motor commands $\vec{u}(t)$ to control the motors which, with the disturbances $\vec{w}(t)$ and the laws of nature act on the robot and sensor properties, thus causing the sensor signals $\vec{y}(t)$. Details of these boxes, including the internal variables, are shown in later figures.

present control values. However, a general form, such as the 'policy' of Eq.2, has little meaning until it is restricted by a particular representation. For example, Taylor Series, Fourier Series and neural networks might all be able to describe any given function with arbitrary accuracy. But what is important about a given representation is its efficiency, how well does it serves the purpose with given implementation cost. In this case, implementation cost is complexity, memory and computation for the design phase and for the real-time calculations.

Reflex-based control

Although we have not attempted mathematical proof, we believe that our form of reflex-based control can describe general non-linear controls of the forms of Eqs. 2 or 3 with arbitrary accuracy. More importantly, the reflex-based form is efficient. That is, with relatively few parameters and relatively little calculation on board, and with relatively few design parameters in the controller design, that reflex-based control can lead to low-energy-use and reliable and versatile walking.

Proposed control architecture

The reflex-based representation we have chosen for eq. 3 is layered from high level commands down through several layers to the motor commands. We have not found a complete and compact form that is nearly as revealing as Eq. 1 is for control. Instead, we describe the layers.

1. **System.** At the highest level this walking robot platform responds to high level commands $\vec{c}(t)$ and physical disturbances $\vec{w}(t)$. These, through the controller software and the robot physics, leads to an observed useful behavior, like walking. See Fig. 1. The high level command is a mixture of discrete and continuous variables, collectively described by

$$\vec{c}(t) = \text{High level command signals.} \tag{4}$$

Ultimately these will come from a supervisor on the upper body (e.g., ROS, Robot Operating System) possibly reflecting tele-operation. $C(t)$ will have discrete variables, selecting behaviors (e.g., standing, walking, sitting, single step mode) and continuous variables (e.g., desired step length, speed, turning radius, next ground point location, etc). For our testing these will either be fixed, or controlled by tele-operation. These high-level commands feed into the overall Software Control which outputs motor commands $\vec{u}$. In our plan,

$$u_i(t) = \text{Command current for motor } i. \tag{5}$$

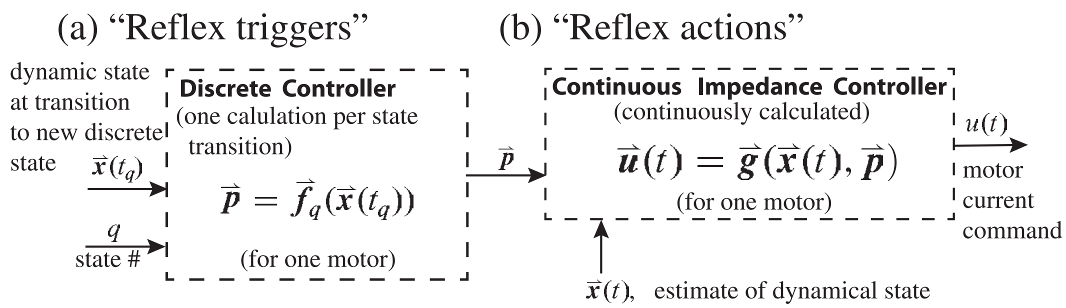

Figure 3: **"Reflex control" for one motor.** Each motor has its own discrete-state machine. **a) "Refex trigger".** At a discrete-state transition the **discrete controller** takes in the discrete-state number q and the value of the dynamical state $\vec{x}(t_q)$ at the time of the discrete-state transition t_q. It calculates a small set of parameters $\vec{p}$ which are held constant for the duration of the discrete state. The collection of all the $\vec{p}$ values from all of the motors is indicated as set $\{\vec{p}\}$ in Fig. 2; **b) "Reflex action".** The now-fixed parameters $\vec{p}$ are used, with the continuous dynamical state estimate $\vec{x}$, to continuously calculate a motor current command $u(t)$.

The physics of the Robot Hardware, including motors, power-electronics, passive mechanical components and sensors, determines the sensor outputs $\vec{y}(t)$, where

$$y_j(t) = \text{raw data from sensor j.} \tag{6}$$

2) The controller. The control software can be divided into 4 parts (see Fig. 2a). The Event Detection and Dynamical State Estimator are run at a high level on a machine-wide process. Each motor i has a separate Discrete Controller and Continuous Impedance Controller whose output is the motor current command $u_i(t)$.

The Dynamical State Estimator will use classical model-based estimation concepts to make a best guess of the present dynamical state, $\vec{x}(t)$. We are on the lookout for information-efficient measurements that can serve as a proxy for dynamical state estimation. For example the most recent step length divided by the most recent step time might be a more robust, or at least simpler to measure, estimate of robot speed than the present speed as estimated by a complex model-based filter. However, the search for such is not a central thrust.

Continuous and discrete-time variables. There is no genuinely continuous-time control; all control commands are determined by computers refreshing output at a rate on the order of 1 khz. For simplicity we use the words 'continuous' or 'dynamical' for things that are used in feedback at a rate far higher than characteristic times of the mechanical robot. For example, anything measured by a control loop running every 2 ms would be regarded as continuous in relation to an action that is smooth over 50 ms or longer. On the other hand, we regard as *discrete* either things which change suddenly or things which are intentionally held constant for many fast control cycles.

Event detection is based on the sensor data, almost all of which is continuous. That is, we have no plans for explicit mechanical switches. The purpose of the event detector is to take such continuous data and convert it to a switching output. For example, when the sensors on the bottoms of the swing feet exceed a certain threshold for a long enough time the Event Detector will signal a discrete-state transition, namely the start of a double-stance phase. Other commanded discrete-state transitions could be determined by non-sudden 'events'. For example mid-stance seems to be a good time to re-asses the dynamical-state of robot and thus a good time for a discrete-state transition (or reflex trigger). In some sense, the determination of discrete-state transition times t_q and new discrete state q is part of the overall task of state estimation. Broadly speaking, the discrete state is a representation of overall phase in the motion.

3) Reflex control. The controller commands are calculated by "Reflex triggers" and the "Reflex actions". By a mixture of discrete and continuous calculations, the control software, provides

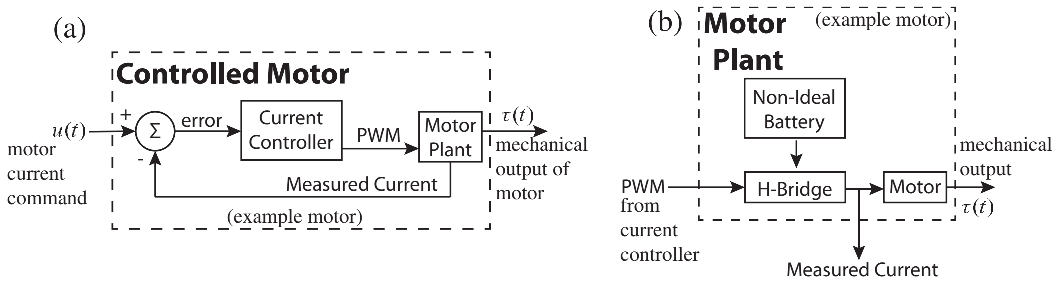

Figure 4: **a) Controlled motor.** This is the lowest level of control. **(a) Current control.** Because the battery voltage is not fixed, in order to have deterministic motor effects a feedback loop is needed to control the motor current, a current sensor is used to regulate a logic-level pulse width modulation (PWM); **(b) The motor "plant"**, the motor hardware without digital control, is s sub-block of the Controlled Motor in (a). It consists of a motor, an H-bridge (power transistors that convert the logic level signal from the computer to a power signal), and a battery whose voltage depletes over time.

a continuous motor current command $\vec{\boldsymbol{u}}(t)$. This command is fed to the robot hardware (Fig. 2b) which consists of power electronics, motors, batteries and passive robot parts (links, hinges, springs, etc). At the mechanical level there are discrete events, namely collisions, but there are no explicitly discrete (that is, impulsive) actuator actions.

Finite state machine. Each motor i is governed by its own finite (meaning discrete) 'state machine'. Technically it is an 'augmented' finite state machine as each fixed state has parameters that can take on various real values. Fig. 3 shows one discrete-state of the state machine for one motor. For a given high-level behavior, there is a block like this for each state q of each motor i.

Discrete controller. A central 'Event Detector' transmits to all motors the occurrence of events. All of the motors' discrete controllers are informed of these discrete state transition events, but an individual motor controller need not make a discrete state transition for every event. That is, a given motor discrete controller in a given discrete state q may or may not be sensitive to a state advance. The independence of the individual motors' finite state machines from each other, that they have independent discrete-state transitions, prevents what might otherwise be a combinatorial explosion of discrete states for the system overall.

For example, when the left heel-strike event is registered the left ankle discrete state q advances from 'prepare for heel-strike' to 'begin stance'. On the other hand, the swing ankle controller will be oblivious to the mid-stance transition and the swing-leg controller will be oblivious to the swing ankle having reached maximum contraction.

At a transition time t_q a motor's discrete controller calculates commands $\vec{\boldsymbol{p}}$ to be fed to the Continuous controller, according to:

$$\vec{\boldsymbol{p}} = \vec{\boldsymbol{f}}_q(\vec{\boldsymbol{x}}(t_q)). \tag{7}$$

That is, for each state q of each motor, it has a different function $\vec{\boldsymbol{f}}_q$ to calculate command outputs to the continuous impedance control. Our hypothesis is that most such control can be accomplished with a sparse (only using one or two of the dynamical states $x_i(t_q)$ at the discrete state transition time, and primarily local variables. For example Eq. 7 should generally simplify to:

$$\vec{\boldsymbol{p}} = \vec{\boldsymbol{p}}_q^{ref} + K_q^{ref} \cdot \left(\vec{\boldsymbol{x}}_q^{ref} - \vec{\boldsymbol{x}}(t_q)\right) \tag{8}$$

where the gain matrix K_q^{ref} should only have a 1-3 non-zero entries.

Continuous impedance controller. Any given motor at any given time is controlled by the output of a function g of the full robot dynamical state $\vec{\boldsymbol{x}}(t)$ and the most recently set values of the discrete commands $\vec{\boldsymbol{p}}$.

$$u(t) = g(\vec{\boldsymbol{x}}(t), \vec{\boldsymbol{p}}) \tag{9}$$

Again, this will generally be a sparse linear function. For example for one joint angle a local impedance (or compliant) control of this form might be adequate:

$$u = g(\underbrace{(\theta, \dot{\theta})}_{\vec{\boldsymbol{x}}(t)}, \underbrace{(P_0, P_1, P_2)}_{\vec{\boldsymbol{p}}}) = P_0 + P_1\theta + P_2\dot{\theta}. \tag{10}$$

The functional form is the same as for a PD trajectory tracker and as for an impedance control with variable stiffness and damping. Here, however, we think of the 'gains' $\vec{\boldsymbol{p}}$ merely as parameters that effect the motion. The functional form is chosen not for trajectory tracking or for impedance imitation, but as a way to have useful control of smooth trajectories with a small number of parameters. Using sparse affine (linear + constant) forms for f and g the overall representation can use very few parameters.

Individual motor control. Each motor gets a command $u(t)$. This logic-level signal needs to have compensation for the non-ideal batteries (voltage decreasing with time). Because at the lowest level we want a machine that is at least able to swing freely, it is nice to have a simple command for that. And, for an ideal motor, torque control (including zero torque, swinging free) is current control. So, even for our real motors we impose an electrical current control at the lowest level (see Fig. 4). One would like to think of this as below the purview of a controller, but because it is implemented in software, it is included as part of the control. Furthermore, unless it is well-managed (that is, made to have a small time constant) it can interfere with the higher-level controls.

Overall, the architecture is that of a hierarchical augmented finite state machine. For any given behavior this is a set of parallel augmented finite state machines, with one finite state machine for each motor. These are synchronized indirectly through sharing of the output of the Dynamical State Estimator and of the Event Detector. The collection of free functions and constants in this architecture are analogous to the elements of a gain matrix in linear control. They are the substrate on which the control design is implemented.

5 Control design

The architecture above evolved in tandem with the evolution of the general control design method. This depends on high-fidelity simulation.

1. With a finest possible grid find the robot trajectory that minimizes energy use for the given constraints.

2. Extract the key features of the trajectory. This was done by hand in the past, we want to develop automatic feature extraction in the project work.

3. Reparameterize the motor current trajectories using the extracted features and the functional form described in the reflex control architecture. The goal is to minimize the number of parameters.

4. Re-optimize using the reflex parameters, check that the energy use is close to the 'true' (that is, fine-grid) optimum.

5. Pick a sparse representation for the reflex (discrete controller) gains. This has been done by hand in the past, and for the new research we want to automate this step.

6. Optimize the gains for one or another measure of stability or robustness.

7. Iterate on all of the above to achieve the optimal tradeoff between a) number of parameters and closeness to the physical limits of machine energy use, and b) the number of gain parameters and the estimated limits on machine-achievable robustness. Again, this was done by hand with Ranger and we want to systematize the generation of such a Pareto front.

8. Test on the robot to check for model fidelity in the context of the control (ability to match gait features, energy use, and response to disturbances).

9. Check for robustness against perturbations not used in development.

6 Summary of Research Plan

Starting in Section 3 on page 2 concepts and details of the research have been described. Here is a summary.

1) Detailed hardware design, construction and testing, as informed by simulation. As briefly detailed in Sect. 4 starting on page 7 this does not involve basic scientific development. However, it does depend on scientifically informed engineering. Elements of Ranger's design will be helpful, but more useful is the design approach that we learned while building Ranger. This much more complex robot needs much new thought. Recently developed design features from other robotics and prosthetics labs will be evaluated for applicability. *Claim: careful simulation-inspired design will lead to a robot as robust as the most robust present robots, but with one tenth the energy use.*

2) High-fidelity simulation based on bench testing of robot components. In years of search, we have not found any simulation packages that deal well with our mixture of needs. In particular, fast convergent optimization (below) depends on accurate (or at least consistent) simulation. This is a project in itself done with advise and support from Evan Drumwright (see support letter). *Hypothesis: Effective control design is enabled by a high-fidelity plant model, even if, in the end, the control is to be robust to model perturbations.*

3) Formalizing the control architecture. While the reflex based control architecture has been proven for the 3-motor Ranger, and it is reasonably precisely described in words, it has not yet been reduced to a formal structure with clearly defined parameters. Such is needed for implementation on a more complex machine, for transferring to other researchers, and for easy use by optimization routines. *Hypothesis: two levels of affine (linear + constant) functions f and g can: a) be simply coded into concurrent state machines for use by automatic optimization, and b) provide adequate non-linearity for robust stabilization.*

4) Development of simple models to guide the theory. We do not hope for general theorems, analogous to theorems from linear control theory, concerning the tradeoffs between, for example, needed actuator power and robustness to various disturbances. But we do hope for basic understanding through simple but well-developed examples. In these, the tradeoffs will not be anecdotal, but clearly understood. Point mass models of running and walking have been highly revealing to us concerning legged locomotion energetics and passive stability (e.g., [26, 6, 1, 64]). But we have not yet sufficiently developed such models for understanding the fundamental issues in legged locomotion control. The results from these models will be used to guide the more complex control of the real robot, and will be used to measure performance. These simple models will be the most durable basic science. *Claim: various simple models will indicate the trade-offs, or lack thereof, of power, torque, sensor accuracy, performance (e.g. speed) robustness and versatility.*

5) Implementing and automating the control design procedure described above. Then testing it, in simulation and on the robot, against simple models. This involves selecting optimization techniques. In the past we have selected discrete variables (e.g., number of discrete states in a state machine, or number of non-zero terms in the control gains and compliant controllers) by hand. To automate this selection we will have to move away from the gradient based optimization that has served us so well so far. *Hypothesis: Systematic implementation of numerical optimization will yield unusually-minimal control design that nearly optimizes energy use and is also robust.*

6) Test the robust control philosophy that, by designing-in robustness to some possible forcings and model errors, robustness to other forcings and model errors is automatically covered. In simulations this can be done by testing with disturbances that were withheld in optimization. On the machine, this is tested by the reliability in the field of a controller designed in simulation. *Hypothesis: systematic controller design, in simulation, for "known unknowns" will lead to robust walking in the real hardware world of "unknown unknowns".*

7 Intellectual merit, broader impacts and outreach

Intellectual merit. The intellectual contributions fall in two categories. Both of these provide scientific foundations from which robot engineering and understanding of legged locomotion should be advanced.

Advances in robotics. Development of a low-bandwidth low-computation approach to smooth low-energy robust control of legged locomotion robots.

Understanding of walking. Walking is a basic human activity. It deserves scientific understanding. Detailed investigations into issues of robot and control design will lead to basic insights into legged locomotion more generally, especially the nature of tradeoffs between motor power, sensor and model accuracy, disturbance rejection and versatility.

Broader impacts

Robotics. While the proposed robot platform design does not represent significant science in itself, the platform is a key experimental device for use by ourselves and others in developing and testing controllers. It will also be a 'proof of possibility' showing the frontier in combining efficiency with robustness. Finally, it will serve its place in the lineage leading to practical humanoid robots.

Human health. The basic biomechanics helps with the understanding of human gait. Our greatest hope is that our research will, in the long run, help prosthetics, orthotics and otherwise help the diagnosis and remediation of problems that people have with walking.

General public. Robots are interesting to people of all ages and all backgrounds. Ruina's previous research has been widely featured in the media (New York Times, NPR, Wired, Engadget, Discover, *etc.*). The proposed 3D biped is all the better suited as inspiring scientific research. We also plan public demonstrations. One goal, which we believe is in reach by us, and perhaps by no other robot group in the world, is to publicly complete a full 42.2 km marathon alongside human participants.

Undergraduates. Undergraduates and Masters students have been central to our robotics research, including women and under-represented minorities. Every year we give lab tours to prospective female undergraduates, alumni and freshman advising groups.

Local students. Every year our lab gives lectures to high school physics classes about robots.

Curriculum development and teaching

This grant will support the STEM teaching activities at Cornell through undergraduate involvement in research, and through use of the shared robotics equipment in classes such as Ruina's *Legged Locomotion of Robots and Animals*. Ruina has a strong commitment to undergraduate education. Robots encourage interactive learning and we will leverage this project towards that end.

Public education. Our lab's web sites are a resource for those interested in our robots and papers.

REFERENCES CITED

[1] J.E.A. Bertram A. Ruina and M Srinivasan. A collisional model of the energetic cost of support work qualitatively explains leg sequencing in walking and galloping, pseudo-elastic leg behavior in running and the walk-to-run transition. *Journal of Theoretical Biology*, 237(2):170–192, 2005.

[2] P.G. Adamczyk, S.H. Collins, and A.D. Kuo. The advantages of a rolling foot in human walking. *Journal of Experimental Biology*, 209(20):3953, 2006.

[3] A.D. Ames, R.D. Gregg, and M.W. Spong. A geometric approach to three-dimensional hipped bipedal robotic walking. In *Conference on Decision and Control, New Orleans, LA*, 2007.

[4] Y. Bar-Cohen. *Electroactive polymer (EAP) actuators as artificial muscles: reality, potential, and challenges*. Society of Photo Optical, 2004.

[5] J. E. A. Bertram and A. Ruina. Multiple walking speed-frequency relations are predicted by constrained optimization. *Journal of Theoretical Biology*, 209(4):445–453, 2001.

[6] J. E. A. Bertram, A. Ruina, C. E. Cannon, Y.-H. Chang, and M. Coleman. A point-mass model of gibbon locomotion. *Journal of Experimental Biology*, 202(19):2609–2617, 1999.

[7] JT Betts and I. Kolmanovsky. Practical methods for optimal control using nonlinear programming. *Applied Mechanics Reviews*, 55:B68, 2002.

[8] P. A. Bhounsule. *A controller design framework for bipedal robots*. PhD thesis, Cornell University, 2012.

[9] P. A. Bhounsule, J. Cortell, B. Hendriksen, B.D.G. Karrsen, C. Paul, and A. Ruina. Cornell ranger: Energy-efficient, robust walking via energy-optimization and reflex based control. *International Journal of Robotics Research (in press)*, 2011.

[10] P.A. Bhounsule and A. Ruina. Cornell ranger: Energy-optimal trajectory control. *International Journal of Robotics Research (accepted, in revision)*, 2012.

[11] P.A. Bhounsule and A.L. Ruina. Cornell Ranger: Energy-Optimal Control. In *Proceedings of Dynamic Walking 2009*, 2009.

[12] Bio-Robotics and Locomotion Laboratory. Cornell ranger 2011. World Wide Web electronic publication, 2011.

[13] R. Blickhan. The spring-mass model for running and hopping. *Journal of Biomechanics*, 22(11-12):1217, 1989.

[14] A.E. Bryson and Y.C. Ho. *Applied optimal control*. wiley New York, 1975.

[15] D. Cabrera, A. Ruina, and V. Kleshnev. A simple 1+ dimensional model of rowing mimics observed forces and motions. *Human movement science*, 25(2):192–220, 2006.

[16] C. Chevallereau, JW Grizzle, and C. Shih. Asymptotically stable walking of a five-link underactuated 3d bipedal robot. *IEEE Transactions on Robotics*, 25(1):37–50, 2008.

[17] C. Chevallereau, J.W. Grizzle, and C.L. Shih. Asymptotically stable walking of a five-link underactuated 3-D bipedal robot. *Robotics, IEEE Transactions on*, 25(1):37–50, 2009.

[18] C.P. Chou and B. Hannaford. Measurement and modeling of McKibben pneumatic artificial muscles. *IEEE Transactions on robotics and automation*, 12(1):90–102, 1996.

[19] M. Coleman and A. Ruina. An uncontrolled toy that can walk but cannot stand still. *Physical Review Letters*, 80(16):3658–3661, 1998.

[20] M. J. Coleman, M. Garcia, K. Mombaur, and A. Ruina. Prediction of stable walking for a toy that cannot stand. *Physical Review E*, 6402(2):2901–+, August 2001.

[21] S.H. Collins, M. Wisse, and A. Ruina. A three-dimensional passive-dynamic walking robot with two legs and knees. *The International Journal of Robotics Research*, 20(7):607, 2001.

[22] M.R. Dimitrijevic, Y. Gerasimenko, and M.M. Pinter. Evidence for a spinal central pattern generator in humans. *Annals of the New York Academy of Sciences*, 860(1):360–376, 1998.

[23] Boston Dynamics. Petman - bigdog gets a big brother. World Wide Web electronic publication, 2010.

[24] RJ Full and DE Koditschek. Templates and anchors: neuromechanical hypotheses of legged locomotion on land. *Journal of Experimental Biology*, 202:3325–3332, 1999.

[25] M. Garcia, A. Chatterjee, and A. Ruina. Efficiency, speed, and scaling of two-dimensional passive-dynamic walking. *Dynamics and Stability of Systems*, 15(2):75–99, 2000.

[26] M. Garcia, A. Chatterjee, A. Ruina, and M. Coleman. The simplest walking model: Stability, complexity, and scaling. *ASME Journal of Biomechanical Engineering*, 120(2):281–288, 1998.

[27] H. Geyer and H. Herr. A muscle-reflex model that encodes principles of legged mechanics produces human walking dynamics and muscle activities. *Neural Systems and Rehabilitation Engineering, IEEE Transactions on*, 18(3):263–273, 2010.

[28] M. Gomes and A. Ruina. A five-link 2d brachiating ape model with life-like zero-energy-cost motions. *Journal of Theoretical Biology*, 237(33):265–278, 2005.

[29] M. Gomes and A. Ruina. Walking model with no energy cost. *Physical Review E*, 83(3):032901, 2011.

[30] R.D. Gregg and M.W. Spong. Reduction-based control of three-dimensional bipedal walking robots. *The International Journal of Robotics Research*, 29(6):680–721, 2010.

[31] JW Grizzle. Remarks on event-based stabilization of periodic orbits in systems with impulse effects. In *International Symposium on Communications, Control and Signal Processing, Marrakech, Morocco*, 2006.

[32] J.W. Grizzle, G. Abba, and F. Plestan. Asymptotically stable walking for biped robots: Analysis via systems with impulse effects. *IEEE Transactions on Automatic Control*, 46(1):51–64, 2001.

[33] JW Grizzle, C. Chevallereau, and C.L. Shih. HZD-based control of a five-link underactuated 3d bipedal robot. In *Decision and Control, 2008. CDC 2008. 47th IEEE Conference on*, pages 5206–5213. IEEE, 2009.

[34] JW Grizzle and ER Westervelt. Event-based PI control of an underactuated biped walker. In *Decision and Control, 2003. Proceedings. 42nd IEEE Conference on*, volume 3, pages 3091–3096, 2004.

[35] S. H., M. Wisse, and A. Ruina. A three-dimensional passive-dynamic walking robot with two legs and knees. *International Journal of Robotics Research*, 20(7):607–615, 2001.

[36] K. Hirai, M. Hirose, Y. Haikawa, and T. Takenaka. The development of Honda humanoid robot. In *IEEE International Conference on Robotics and Automation*, pages 1321–1326. Institute of Electrical Engineers (IEEE), 1998.

[37] D. Hobbelen, T. de Boer, and M. Wisse. System overview of bipedal robots flame and tulip: Tailor-made for limit cycle walking. In *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems, IROS2008, Nice, France*, pages 2486–2491, 2008.

[38] D.G.E. Hobbelen and M. Wisse. A disturbance rejection measure for limit cycle walkers: The gait sensitivity norm. *IEEE Transactions on robotics*, 23(6):1213–1224, 2007.

[39] S. Kajita, F. Kanehiro, K. Kaneko, K. Yokoi, and H. Hirukawa. The 3D Linear Inverted Pendulum Mode: A simple modeling for a biped walking pattern generation. In *Intelligent Robots and Systems, 2001. Proceedings. 2001 IEEE/RSJ International Conference on*, volume 1, pages 239–246, 2001.

[40] JDG Kooijman, JP Meijaard, J.M. Papadopoulos, A. Ruina, and AL Schwab. A bicycle can be self-stable without gyroscopic or caster effects. *Science*, 332(6027):339, 2011.

[41] AD Kuo. Choosing your steps carefully. *Robotics & Automation Magazine, IEEE*, 14(2):18–29, 2007.

[42] M.L. Latash, J.P. Scholz, and G. Schöner. Motor control strategies revealed in the structure of motor variability. *Exercise and sport sciences reviews*, 30(1):26–31, 2002.

[43] M.A. Locascio, J.H. Solomon, and M.J.Z. Hartmann. Linear reactive control of three-dimensional bipedal walking in the presence of noise and uncertainty. *Adaptive Behavior*, 2012.

[44] P. Manoonpong, F. Pasemann, and H. Roth. Modular Reactive Neurocontrol for Biologically Inspired Walking Machines. *International Journal of Robotics Research*, 26(3):301–331, 2007.

[45] T. McGeer. Passive dynamic walking. *International Journal of Robotics Research*, 9:62–82, 1990.

[46] JP Meijaard, J.M. Papadopoulos, A. Ruina, and AL Schwab. Linearized dynamics equations for the balance and steer of a bicycle: a benchmark and review. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 463(2084):1955, 2007.

[47] JP Meijaard, J.M. Papadopoulos, A. Ruina, and AL Schwab. History of thoughts about bicycle self-stability. *Cornell e-commons*, 2011.

[48] J.B. Morrell and D. Field. Design of a closed loop controller for a two wheeled balancing transporter. In *Intelligent Robots and Systems, 2007. IROS 2007. IEEE/RSJ International Conference on*, pages 4059–4064. IEEE, 2007.

[49] Hae-Won Park, Koushil Sreenath, Alireza Ramezani, and J.W. Grizzle. Switching control design for accommodating large step-down disturbances in bipedal robot walking. In *Int. Conf. on Robotics and Automation (ICRA)*, pages 45–50, May 2012.

[50] R.J. Peterka. Simplifying the complexities of maintaining balance. *IEEE Engineering in Medicine and Biology Magazine*, 22(2):63–68, 2003.

[51] G.A. Pratt, M.M. Williamson, P. Dillworth, J. Pratt, and A. Wright. Stiffness isn't everything. *Lecture Notes in Control and Information Sciences*, pages 253–262, 1996.

[52] J. Pratt, J. Carff, S. Drakunov, and A. Goswami. Capture point: A step toward humanoid push recovery. In *Humanoid Robots, 2006 6th IEEE-RAS International Conference on*, pages 200–207. IEEE, 2007.

[53] J. Pratt and G. Pratt. Intuitive control of a planar bipedal walking robot. *IEEE International Conference on Robotics and Automation*, 3:2014–2021, 1998.

[54] J. Pratt and G. Pratt. Exploiting natural dynamics in the control of a 3d bipedal walking simulation. In *In Proc. of Int. Conf. on Climbing and Walking Robots (CLAWAR99)*, 1999.

[55] Jerry Pratt. *Exploiting Inherent Robustness and Natural Dynamics in the Control of Bipedal Walking Robots*. PhD thesis, MIT, 2000.

[56] M.H. Raibert et al. *Legged robots that balance*. MIT press Cambridge, MA, 1986.

[57] L. Righetti and A.J. Ijspeert. Programmable central pattern generators: an application to biped locomotion control. In *In Proc. of IEEE International Conference on Robotics and Automation, Orlando, Florida, USA*, pages 1585–1590, 2006.

[58] A. Ruina. Nonholonomic stability aspects of piecewise holonomic systems. *Reports on Mathematical Physics*, 42(1–2):91–100, 1998.

[59] Segway. Segway rmp — mobile robot platforms www site, November 2011.

[60] S. Seok, W. Wang, D. Otten, and S. Kim. Actuator design for high force proprioceptive control in fast legged locomotion. In *IROS*. IEEE, October 2012.

[61] J.H. Solomon, M.A. Locascio, and M.J.Z. Hartmann. Linear reactive control for efficient 2d and 3d bipedal walking over rough terrain. *Adaptive Behavior*, 2012.

[62] Koushil Sreenath, Hae-Won Park, and J.W. Grizzle. Design and experimental implementation of a compliant hybrid zero dynamics controller with active force control for running on MABEL. In *Int. Conf. on Robotics and Automation (ICRA)*, pages 51–56, May 2012.

[63] M. Srinivasan and A. Ruina. Idealized walking and running gaits minimize work. *Proceedings of the Royal Society A*, 463(2086):2429, 2007.

[64] Manoj Srinivasan and Andy Ruina. Computer optimization of a minimal biped model discovers walking and running. *Nature Magazine*, 439:72 – 75, Jan 5, 2006.

[65] P. Steinkamp, G. Stiesberg, T. van Oijen, and A. Ruina. A passive-dynamic hopper that seems stable. In *Dynamic Walking, in prep for IJRR*, 2009.

[66] Russ Tedrake Steve Collins, Andy Ruina and Martijn Wisse. Efficient bipedal robots based on passive dynamic walkers. *Science Magazine*, 307:1082–1085, 2005.

[67] G. Taga. A model of the neuro-musculo-skeletal system for human locomotion. *Biological Cybernetics*, 73(2):97–111, 1995.

[68] G. Taga. A model of the neuro-musculo-skeletal system for anticipatory adjustment of human locomotion during obstacle avoidance. *Biological Cybernetics*, 78(1):9–17, 1998.

[69] R. Tedrake. LQR-Trees: Feedback motion planning on sparse randomized trees. In *Proceedings of Robotics Science and Systems (RSS), Zaragoza, Spain*, 2009.

[70] E. Todorov and M.I. Jordan. Optimal feedback control as a theory of motor coordination. *Nature neuroscience*, 5(11):1226–1235, 2002.

[71] A.K. Valenzuela and S. Kim. Optimally scaled hip-force planning: A control approach for quadrupedal running. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 1901–1907. IEEE, 2012.

[72] M. Vukobratovic and B. Borovac. Zero-moment point-thirty five years of its life. *International Journal of Humanoid Robotics*, 1(1):157–173, 2004.

[73] E.R. Westervelt. *Toward a coherent framework for the control of planar biped locomotion.* PhD thesis, University of Michigan, 2003.

[74] E.R. Westervelt, J.W. Grizzle, C. Chevallereau, J.H. Choi, and B. Morris. *Feedback control of dynamic bipedal robot locomotion.* CRC Press, 2007.

[75] K.K. Yin, K. Loken, and M. van de Panne. Simbicon: simple biped locomotion control. In *ACM Transactions on Graphics (TOG)*, volume 26, page 105. ACM, 2007.

[76] V. Zykov, J. Bongard, and H. Lipson. Evolving dynamic gaits on a physical robot. In *Proceedings of Genetic and Evolutionary Computation Conference, Late Breaking Paper, GECCO*, volume 4, 2004.