ELSEVIER

Neurocomputing 38–40 (2001) 1409–1421

NEUROCOMPUTING

www.elsevier.com/locate/neucom

Control of a robot leg with an adaptive aVLSI CPG chip

M. Anthony Lewis[a,*], Mitra J. Hartmann[b], Ralph Etienne-Cummings[c], Avis H. Cohen[d]

[a]*Iguana Robotics, Inc., P.O. Box 628, Mohomet, IL, 61853, USA*
[b]*Jet Propulsion Laboratory, Caltech MC 303-300 4800, Oak Grove Drive, Pasadena, CA, 91109, USA*
[c]*Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD, 21218, USA*
[d]*Department of Biology, Neuroscience and Cognitive Science, University of Maryland, College Park, MD, 20742, USA*

Abstract

The rhythmic locomotion of animals, such as walking, swimming, and flying, is controlled by groups of neurons called central pattern generators (CPGs). CPGs can autonomously produce rhythmic output, but under normal biological conditions make extensive use of peripheral sensory feedback. Models of CPGs have been used to control robot locomotion, but none of these models have incorporated sensory feedback adaptation. We have constructed an adaptive CPG in an analog VLSI chip, and have used the chip to control a running robot leg. We show that adaptation based on sensory feedback permits a stable gait even in an underactuated condition: the leg can be driven using a hip actuator alone while the knee is purely passive.

Keywords: Neuromorphic engineering; Locomotion; Central pattern generator; Sensorimotor integration; Adaptive control

1. Introduction

A central pattern generator (CPG) is most generally defined as a biological circuit capable of autonomously generating sustained oscillations [2,7]. In particular, locomotor CPGs produce a rhythmic pattern of neuronal discharge that can drive muscles in a fashion similar to that seen during normal locomotion. Locomotor CPGs are

* Corresponding author.
E-mail addresses: tlewis@iguana-robotics.com (M. Anthony Lewis), hartmann@brain.jpl.nasa.gov (M.J. Hartmann), etienne@ece.jhu.edu (R. Etienne-Cummings), ac61@umail.umd.edu (A.H. Cohen).

0925-2312/01/$ - see front matter
PII: S0925-2312(01)00506-9

autonomous in the sense that they can operate without input from higher centers or from sensors. Under normal conditions, however, they make extensive use of sensory feedback from the muscles and skin, as well as descending input [3].

In the early 1980s Cohen and colleagues [4] analyzed a CPG model using a system of phase-coupled oscillators. Lewis et al. [9,12,10] developed a model of the CPG called adaptive ring rules (ARR). Fundamentally, ARR is based on phase-coupled oscillators, but is highly elaborated. The elaboration includes the use of plastic output functions and phase-coupling functions, and allows modeling of plasticity and learning in CPG elements. ARR has sufficient complexity to drive robotic devices [9,10]. In hardware, there have been several VLSI implementations of CPGs used to control robot movement [16,14,15]. In particular, Still and Schölkopf [16] developed a supervised-learning method for tuning a CPG based on knowledge of desired phase and duty cycle.

Here we present an adaptive VLSI chip, based on established principles of locomotor-control circuits in the nervous system, that mimics many features of a biological CPG. The circuit can successfully control a robot leg running on a treadmill. It uses real-time sensory feedback from position sensors to stabilize running rhythmicity, and a competitive mechanism to establish the value for an adaptive parameter. Unlike Still and Schölkopf [16] we do not use supervised learning, and thus view the present work as complementary to that effort.

In the presence of sensory feedback, the adaptive properties of the circuit can produce stable running even when the robot leg is "underactuated," meaning that not all of its moving joints are actively driven. The leg can be driven stably with a hip actuator alone, allowing the knee to be passive. We show that the sensory feedback serves to entrain the CPG, such that the gait can recover from large asymmetries in the backward and forward swings of the leg. Selective "lesions" of this sensory feedback resulted in a deterioration of the gait, but the gait recovered when sensory feedback was restored. We also show that sensory feedback is critical for compensating for momentary external perturbations.

It should be noted that we have, in essence, the minimal system that can produce running behavior. This system has three essential elements: (1) A driven hip and passive knee, (2) a CPG capable of entrainment and (3) a CPG capable of output amplitude modulation. If any essential element is removed, the system will not work. The mechanisms of adjustment are simple and straight-forward. Our approach represents an extreme in the spectrum of modeling. We tend towards an absolutely minimal system, removing obfuscating detail. From this starting point we can make our model increasingly complex to incorporate additional biological details.

2. A novel CPG chip and circuits

The circuits described below were implemented on a VLSI chip which is not shown here for reasons of space. The circuits consume less than $1\,\mu W$ of power and occupy less than $0.4\,mm^2$ in area. Those interested in more details should refer to [11].

2.1. The integrate-and-fire neuron

All of the neurons in our circuit are modified "integrate-and-fire" neurons, as shown in Fig. 1 [13,8]. The capacitor $\mathbf{C}_{\mathrm{mem}}$ represents the *membrane capacitance* of a biological neuron, and the charge on $\mathbf{C}_{\mathrm{mem}}$ is thus directly proportional to the *membrane potential* ($\mathbf{V}_{\mathrm{mem}}$) of such a cell. $\mathbf{C}_{\mathrm{mem}}$ integrates the impinging charge from both excitatory and inhibitory spikes. Each spike puts charge onto the capacitor, but exactly how much charge is transferred per spike is determined by the excitatory and inhibitory synaptic weights, which are set by external voltages (bias voltages).

The values for $\mathbf{I}_{\mathrm{spon}}$ and $\mathbf{I}_{\mathrm{dis}}$ are set by carefully adjusting bias voltages at the start of an experiment. Once set, these parameters are never manually adjusted again. However, for spike frequency adaptation, the value of $\mathbf{I}_{\mathrm{spon}}$ is automatically adjusted based on sensory inputs and a competitive adaptation rule, which is described below (Section 2.4).

The hysteretic comparator has an upper and lower threshold, both of which are set at design time and can not be changed after fabrication. When the membrane potential rises to exceed the *upper* threshold, the output of the comparator goes high. The output will stay high until the membrane potential falls below the *lower* threshold. From a theoretical point of view, the comparator may be thought of as providing the second state variable needed for the system to oscillate.

The output of the comparator would stay high forever if we did not have a mechanism to bring the membrane potential back below (the lower) threshold. This is accomplished by feeding back the spike train into a transistor at the input to the comparator ("Spike Reset"). This resets the membrane potential because the feedback opens a path between the membrane capacitor and ground. Because the size of the discharging current ($\mathbf{I}_{\mathrm{dis}}$), is larger than the charging current ($\mathbf{I}_{\mathrm{spon}}$), charge is forced away from the capacitor. Hence, the membrane potential decreases until it is below the lower threshold, and the output of the comparator goes to zero.

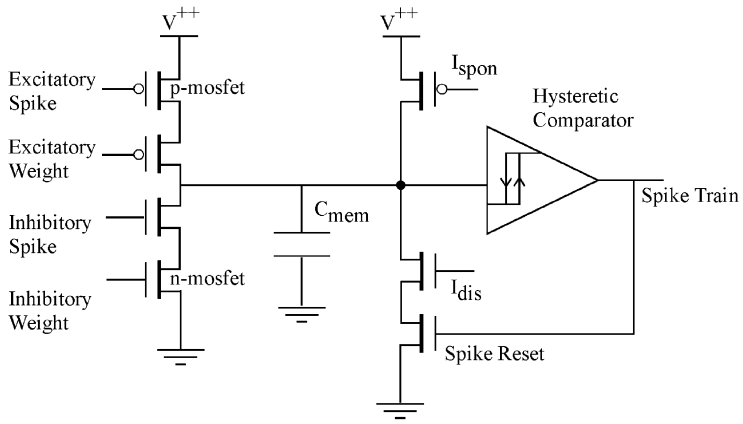

Fig. 1. The integrate-and-fire neuron. $\mathbf{I}_{\mathrm{spon}}$ and $\mathbf{I}_{\mathrm{dis}}$ set the inter-spike interval and the spike width, respectively. Transistors with "bubbles" indicate p-channel mosfets.

Note that in this circuit, the *interspike interval* is determined by the magnitude of $\mathbf{I}_{spon}$, since it determines the rate at which the membrane potential charges. In contrast, the *spike width* is set by the strength of the reset current, $\mathbf{I}_{dis}$. The spike height is fixed and is equal to V^{++} (the supply voltage). The behavior of the neuron can be varied depending on the ratio of $\mathbf{I}_{dis}/(2.0*\mathbf{I}_{spon})$. If this ratio is set to 1, the neuron will oscillate with a 50% duty cycle. If the ratio is very high (e.g., > 25), the neuron will fire a continuous series of narrow spikes.

2.2. Linking together several integrate-and-fire neurons to create the CPG

As illustrated in Fig. 2a, we created the basic CPG circuit by linking together three of the integrate-and-fire neurons shown in Fig. 1. The circuit consists of a "pacemaker" neuron and two "motoneurons." For the pacemaker neuron, we set $\mathbf{I}_{spon}$ and $\mathbf{I}_{dis}$ to establish a low spike rate (~ 0.5 to 3 Hz), mimicking the basic central rhythm in biological systems [1]. In contrast, for the motoneurons we set $\mathbf{I}_{spon}$ and $\mathbf{I}_{dis}$ so that the spike rate was more typical of biological multiunit motoneuron activity (~ 500 Hz).

The output of the pacemaker is connected to the "inhibitory spike" input of the motoneurons, and thus *inhibits* motoneuron firing: the motoneurons can not fire when the pacemaker output is high. In principle, we could have used two pacemakers with an arbitrary phase relationship to inhibit the two motoneurons independently. For the experiments described in this paper, however, we always wanted the motoneurons to burst exactly 180° out of phase with each other. We therefore used only a single

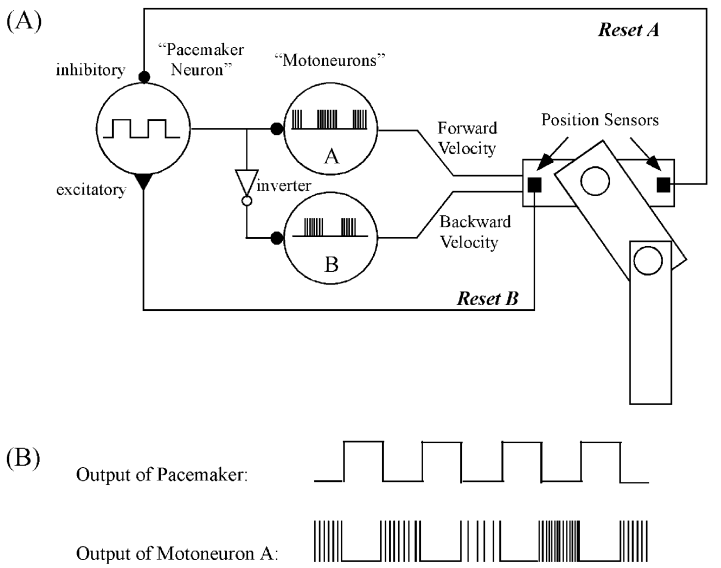

Fig. 2. Several integrate-and-fire neurons are linked together to form the CPG. (A) Using sensory feedback to entrain the pacemaker neuron. (B) Output of the pacemaker and motoneurons.

pacemaker and an inverter, as shown in Fig. 2a, to produce two motoneuron spike trains that were 180 degrees out of phase. These spike trains are then converted to a signal that can actually drive the servo-motors; this conversion is a technical detail and is not illustrated.

For clarity, we have schematized the output of each stage of the circuit in Fig. 2b. The pacemaker outputs square waves with a 50% duty cycle. The output of motoneuron A consists of trains of spike bursts. The burst duration and inter-burst interval are determined by the pacemaker, while the number of spikes per burst is determined by the spiking rate of the motoneuron. If there are many spikes in a particular burst, the final amplitude of the limb swing will be large, while if there are few spikes in the burst, the amplitude of the limb swing will be small. Thus the number of spikes per burst essentially indicates the motor velocity.

2.3. Entraining the CPG with sensory feedback

When an animal's leg is moved to an extreme position, sensory receptors send a signal to the animal's locomotor CPG, causing the phase to reset [1]. In other words, during natural locomotion, the CPG can be entrained with sensory feedback. This entrainment effectively results in a centered stride, so that the leg moves forwards and backwards through approximately equal angles. Similarly, we entrained our CPG with sensory feedback to adaptively recenter the swing of the robot leg.

Fig. 2a illustrates our method for entraining the CPG with sensory feedback to center the swing of the robot leg. Specifically, the pacemaker neuron can be stopped or started with direct inhibitory and excitatory inputs from sensors on the leg. If the inputs to the pacemaker are strongly inhibitory, the membrane capacitor ($\mathbf{C}_{\mathrm{mem}}$) will be shunted and discharged completely (cf. Fig. 1). The output of the pacemaker will go low, and will remain in this state until the inhibition is released, at which point the normal dynamics of the pacemaker will resume. In contrast, if the input to the pacemaker is strongly excitatory, the pacemaker will be driven high, and resume its normal dynamics only when the excitation is released. Thus, if the sensory inputs are periodic, the pacemaker output is driven such that it is phase locked to the inputs, and the pacemaker is entrained to the dynamics of the system under control.

As shown in Fig. 2a, we used position sensors to signal the extremes of hip motion. The sensors send a signal to **Reset A** or **Reset B** to cause the pacemaker circuit to reset, and thus a hip joint velocity reversal. Reversing the trajectory of motion at the end points allows even a very asymmetric forward and backward stride (intentionally introduced) to be adaptively re-centered, as will be shown in the results.

2.4. Adjusting swing amplitude using a competitive adaptation mechanism

Fig. 3 shows a schematic for using a competitive adaptation mechanism to adapt the motoneuron spiking frequency to adjust the swing amplitude of the limb. $\mathbf{V}_{\mathrm{learn}}$, the voltage across $\mathbf{C}_{\mathrm{learn}}$, directly regulates the current $\mathbf{I}_{\mathrm{spon}}$ for both motoneurons (cf. Fig. 1). This mechanism is analogous to modulating cell-membrane conductances to alter the excitability of a cell. In the absence of training inputs, $\mathbf{V}_{\mathrm{learn}}$ decays at

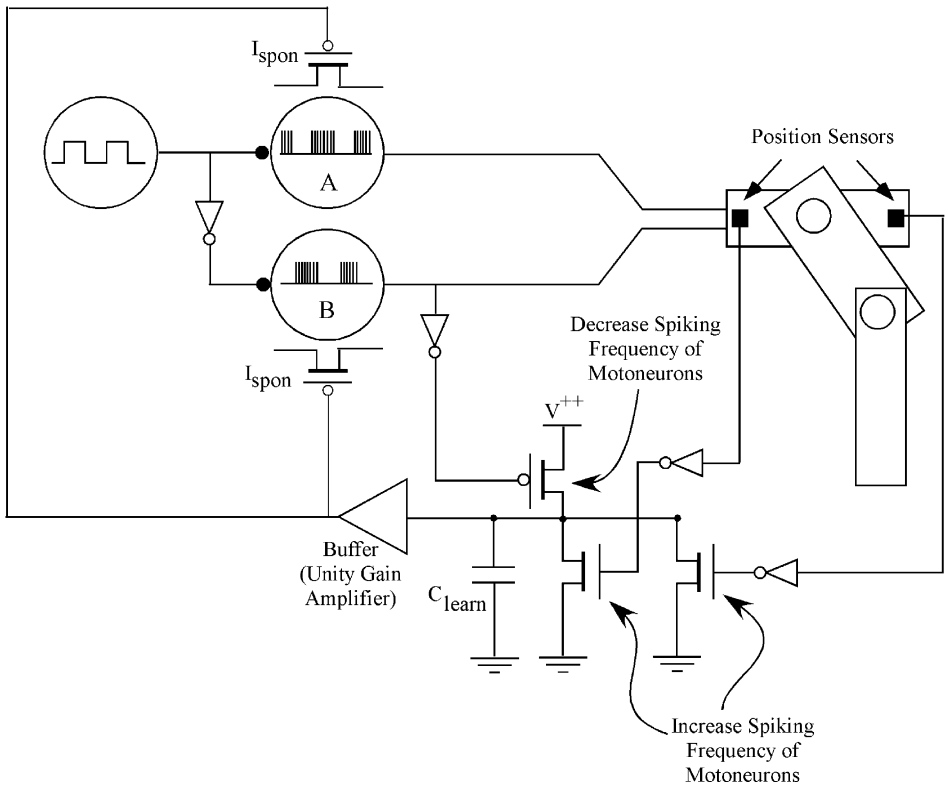

Fig. 3. Spike-frequency adaptation. Competition between signals from the position sensors and from motoneuron B maintain a constant voltage across $\mathbf{C}_{\text{learn}}$, thus adjusting both motorneurons' spiking frequency. Note that the motoneuron spikes must be inverted to provide the correct polarity signal for the p-channel mosfet.

approximately 0.1 V/s. Training inputs come directly from motoneuron B and from the two position sensors. Spikes from the motoneuron increase $\mathbf{V}_{\text{learn}}$. Increasing $\mathbf{V}_{\text{learn}}$ decreases $\mathbf{I}_{\text{spon}}$, which in turn increases the interspike interval, and thus decreases the spiking frequency of the motoneurons. As the spiking frequency decreases, so does the swing amplitude of the limb (see Fig. 2). If allowed to operate in isolation, this part of the circuit would ensure that over time the limb would completely stop swinging.

A competitive mechanism is used to counterbalance this tendency of the motoneuron to shut itself down. Specifically, if the limb does not reach the extreme positions, the position sensors will not be activated. This decreases $\mathbf{V}_{\text{learn}}$, which increases $\mathbf{I}_{\text{spon}}$, which in turn increases the spiking frequency of the motoneurons, and thus the stride length of the limb swing. Thus via competition, a motoneuron spike rate is found that drives the leg sufficiently to trigger the joint feedback sensors, but not so large as to overdrive the limb.

In practice, the stride length increases until the limb hits the extremes and the position sensors are activated. At that point, the input to the "increase frequency" transistors disappears, and the "decrease frequency" signal from the motoneuron dominates. As the spiking frequency decreases, the stride length decreases until the limb no longer hits the extremes. This in turn reactivates the "increase frequency" transistors, and so on. Thus a constant push-pull relationship is established between the "decrease frequency" signal from the motoneuron and the "increase frequency" signal from the position detectors, ensuring that $\mathbf{V}_{\text{learn}}$ will jitter around a stable value. The relative strengths of the "increase frequency" signals and the "decrease frequency" signal will change the amplitude of that jitter.

3. Methods for using the CPG chip

3.1. The robot leg and sensors

The robotic leg is a small (10-cm height) two-joint mechanism. For all experiments described, only the hip joint was driven, and the "knee" remained completely passive. The knee rotated on a low-friction ball bearing joint, and was prevented from hyperextension with a hard mechanical stop. The leg ran on a drum that was free to rotate under the contact forces of the leg.

The leg has three sensors on it. Two inductive (LVDT) sensors monitor the position of the knee and hip joints, and a pressure sensor on the foot monitors ground forces. The LVDT sensor measuring hip position produced a signal that was monotonically increasing with hip angle over the swing range used in our experiments. We were therefore able to create position sensors simply by appropriately thresholding the output of the hip sensor. In other words, the "position sensors" indicated in Figs. 2a and 3 were just the extreme values of the hip sensor output.

3.2. Circuitry used

As discussed in Section 2, we have two separate mechanisms to change two independent parameters of the limb swing. By using the outputs of the position sensors to entrain the pacemaker, we can change the *phase*[1] of the motoneuron spike bursts, and thus ensure a symmetric swing (Section 2.3). By using the outputs of the position sensors and motoneuron as inputs to an adaptive circuit, we can change the *frequency* of the motoneuron spiking within a burst, and thus ensure an acceptable stride length (Section 2.4). In the results shown below, we combined the circuits shown in Figs. 2a and 3, so as to enable both burst duration adaptation and spike frequency adaptation, and thus stabilize the running gait of the leg.

[1] Note that for a cycle of a given (fixed) period, varying the phase will not in and of itself change the amplitude of the swing. The swing-amplitude changes because appropriate phase resetting changes the total amount of time that there are spikes.

We performed several experiments on the system's ability to recenter an asymmetric stride (adaptive stride recentering, Section 2.3). To mimic the effect of "sensory lesions," we simply turned off the feedback from the position detectors to the pacemakers (disconnected Resets A and B). The competitive adaptation circuit, which adjusted swing amplitude, was left intact.

3.3. Chip bias

To demonstrate that the system could in fact use feedback from the position sensors to compensate for an asymmetric swing, we had to first generate an asymmetric gait. We therefore introduced a bias in the chip that tended to make the leg swing backwards more than it did forwards.

4. Results

The foremost result is that the circuit adapts such that the passive knee joint has the correct dynamics to enable running. Fig. 4 shows a phase plot of the hip and knee

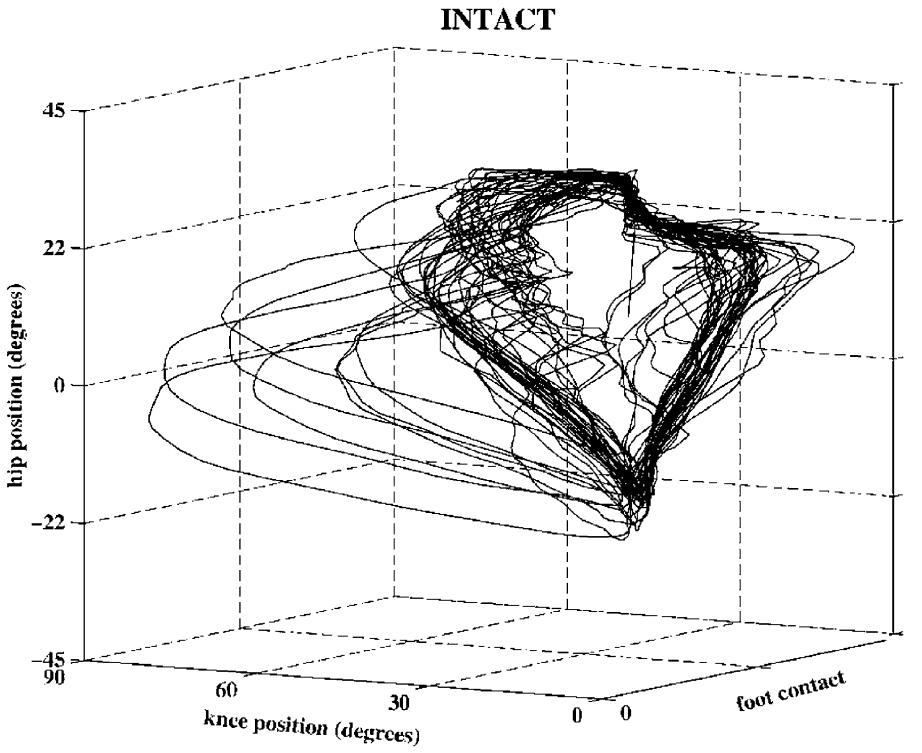

Fig. 4. Hip, knee and foot-contact phase diagram. Most of the trajectory is in a tight bundle, while the outlying trajectories represent perturbations.

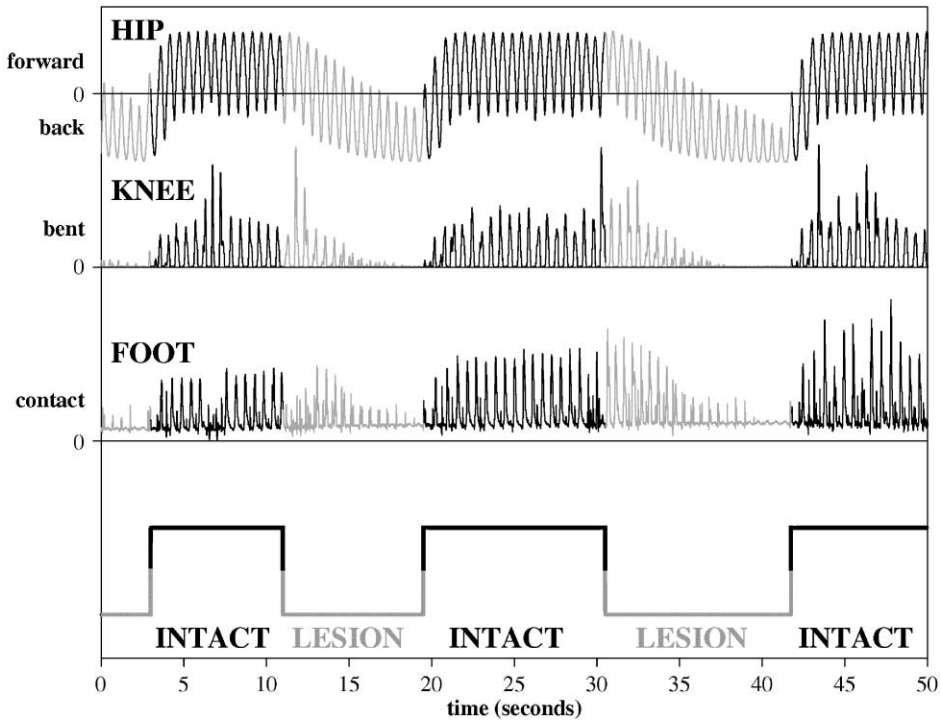

Fig. 5. The effects of lesioning sensory feedback. When the feedback is lesioned (time 11–19 s and 31–42 s), the hip drives backward significantly. As it does the foot begins to lose contact with the ground, and the knee stops moving. When the lesion is reversed at 19 and 42 s, the regularity of the gait is restored.

position, and foot contact force. The bulk of the trajectory describes a tight "spinning top" shape, while the few outlying trajectories are caused by external disturbances. After a disturbance the trajectory quickly returns to its nominal orbit, implying that the system is stable.

We next established that sensory feedback was critical to entrain the pacemaker neuron and thereby ensure a centered stride, by performing selective sensory feedback "lesions." As described above (Section 3.3), we first added a chip bias so that there were significant asymmetries in the forward and backward swing of the limb. Fig. 5 shows the effect of lesioning sensory feedback on the position of the hip and knee joints and on the tactile input to the foot. When feedback is intact, the circuit adjusts for the chip-induced asymmetry, but when feedback is lesioned the leg drifts backwards. The leg returns to a stable gait only after sensory input is restored.

We also found that with sensory feedback, the circuit could compensate for momentary environmental perturbations. As seen above in Fig. 4, environmental perturbations occasionally caused "outlying" trajectories. Fig. 6 shows the trajectory after perturbation in the intact and lesioned cases. In the intact case, a perturbation at cycle '2' leads to outlying trajectories, but the trajectory is quickly restored to the

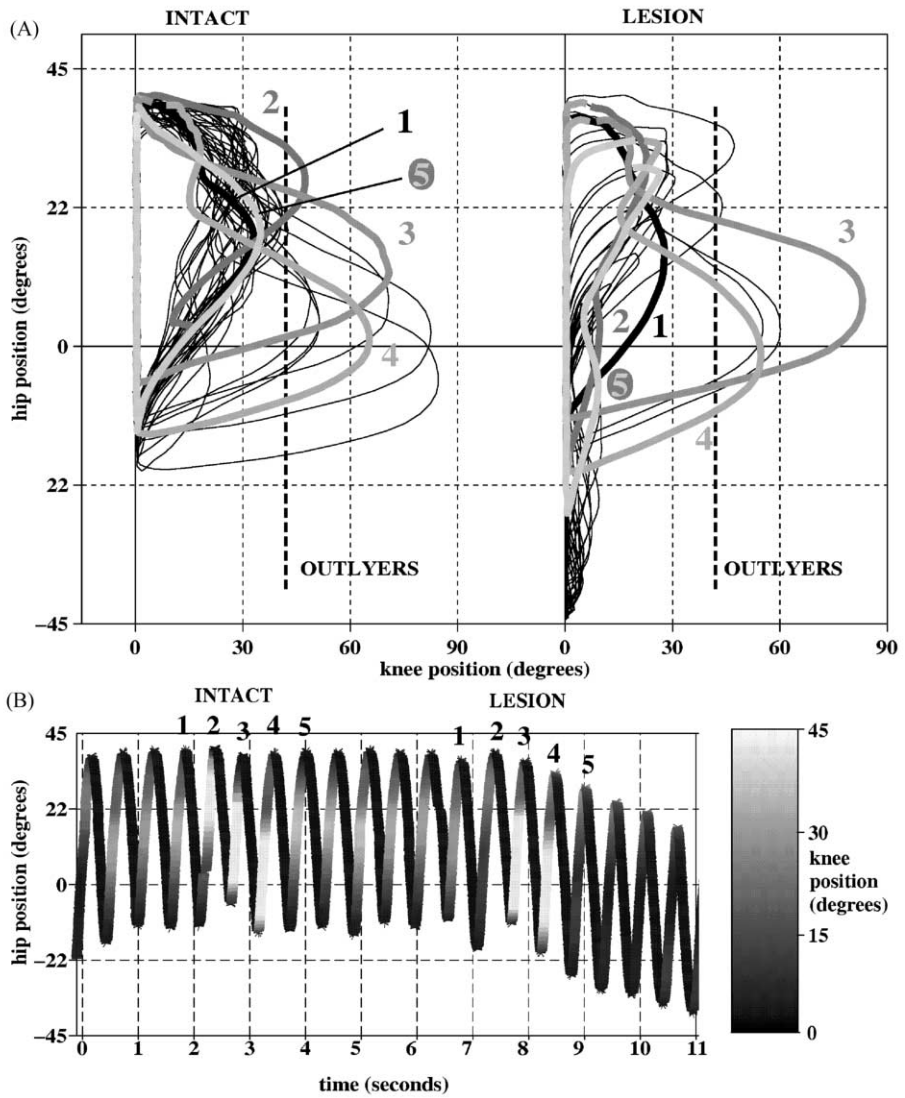

Fig. 6. Effect of perturbations on gait with intact and lesioned sensory feedback. (A) Five sequential trajectories (numbered) in intact and lesioned conditions are represented as ranging between black and light gray. A perturbation at 2 in the intact case leads initially to worse trajectories (3 and 4), but quickly stabilizes to the nominal orbit (5). In the lesioned case, chip bias causes a perturbation at 2 from which the gait can not recover; the hip is forced backward (3, 4, and 5). (B) The same 10 trajectories shown in A presented as hip position through time, with knee position color-coded. Intact sensory feedback permits recovery while lesioning causes drift of both the hip and knee.

nominal orbit. In the lesioned case, removal of sensory feedback causes the chip bias to destroy the trajectory of the leg and the gait quickly deteriorates.

5. Conclusions

In this paper we present the first experimental results of an adaptive aVLSI CPG chip controlling a robotic leg. Our work differs from previous studies in several respects. First, we allow adaptation based on sensory input. Second, our chip has short-term *onboard* memory devices that allow the continuous, real-time adaptation of both center-of-stride and stride amplitude. In addition, we make use of integrate-and-fire neurons for the output motor neurons. Finally, our abstraction is at a higher level than other reported work [14,15], which lends itself to easier implementation of on-chip learning.

Our use of an underactuated leg underscores the point that the natural dynamics of a mechanical system can be used as part of the control, as occurs in real biological systems [5,6] The energy provided to the hip is sufficient to excite an orbital trajectory of the knee. The hip, knee, and foot sensor orbit is remarkably stable in the presence of sensory feedback (Fig. 4).

Finally, it should be emphasized that the system being controlled is non-linear and the chip itself uses non-linear elements to control it. We have made no attempt to linearize the system, and instead take advantage of the non-linearities. Because (1) we do not make use of models or linearization, (2) we adapt principles from biological systems, and (3) these principles can easily be implemented with low-power integrated circuits, we are able to achieve a very compact solution. Further experimentation with this system will allow us to determine whether a robot can be made to walk by coupling together circuits as presented here. The current results are certainly promising.

References

[1] O. Andersson, S. Grillner, Peripheral control of the cat's step cycle II. Entrainment of the central pattern generators for locomotion by sinusoidal hip movements during fictive locomotion, Acta Physiol. Scand. 118 (1983) 229–239.

[2] T. G. Brown, On the nature of the fundamental activity of the nervous centres; together with an analysis of the conditioning of the rhythmic activity in progression, and a theory of the evolution of function in the nervous system, J. Physiol. (Lond). 48 (1914) 18–46.

[3] A.H. Cohen, S. Rossignol, S. Grillner, Neural Control of Rhythmic Movements in Vertebrates, Wiley & Sons, New York, NY, 1988.

[4] A.H. Cohen, P.J. Holmes, R.H. Rand, The nature of the coupling between segmental oscillators of the lamprey spinal generator for locomotion: a mathematical model, J. Math. Biol. 13 (1982) 345–369.

[5] R.J. Full, Integration of individual leg dynamics with whole body movement in arthropod locomotion, in: R.D. Beer, R.E., Ritzmann, T. McKenna (Eds.), Biological Neural Networks in Invertebrate Neuroethology and Robotics, Academic Press, Boston, MA, 1993.

[6] M. Garcia, A. Chatterjee, A. Ruina, M. Coleman, The simplest walking model: stability, complexity, and scaling, ASME J. Biomech. Eng. 120 (1998) 281–288.

[7] S. Grillner, Neurobiological bases of rhythmic motor acts in vertebrates, Science. 228 (1985) 143–148.
[8] F.C. Hoppensteadt, An introduction to the mathematics of neurons, Cambridge Studies in Mathematical Biology, vol. 6, Cambridge University Press, New York, NY, 1986.
[9] M.A. Lewis, Self-organization of locomotory controllers in robots and animals, Ph.D. Thesis in Electrical Engineering, University of Southern California, Los Angeles, CA, 1996.
[10] M.A. Lewis, Gait adaptation in a quadruped robot, Autonomous Robots, (2001) accepted.
[11] M.A. Lewis, R. Etienne-Cummings, A.H. Cohen, M.J. Hartmann, Toward biomorphic control using custom aVLSI CPG chips, International Conference on Robotics and Automation (ICRA), San Francisco, CA, 2000.
[12] M.A. Lewis, L.S. Simo, Elegant stepping: a model of visually triggered gait adaptation, (special issue on adaptive robots), Connect. Sci. 11 (1999) 331–344.
[13] C. Mead, Analog VLSI and Neural Systems, Addison-Wesley, Reading, MA, 1989.
[14] G. N. Patel, J. H. Holleman, S.P. DeWeerth, Analog VLSI model of intersegmental coordination with nearest-neighbor coupling, in: Advances in Neural Information Processing Systems, Cambridge, MA, MIT Press, (1998) pp. 719–725.
[15] M. Simoni, S. DeWeerth, Adaptation in a VLSI Model of a Neuron, Trans. Circuits and Systems II 46 (1999) 967–970.
[16] S. Still, B. Schölkopf, Four-legged walking gait control using a neuromorphic chip interfaced to a support vector learning algorithm, NIPS 2000, Denver, CO, 2000.

M. Anthony Lewis is President and a founder of Iguana Robotics, Inc. He received a B.S. in Cybernetics from the University of California, Los Angeles, and a M.S. and a Ph.D. in Electrical Engineering from the University of Southern California in 1994 and 1996, respectively. His research areas include learning in legged locomotion, visuomotor coordination, and medical applications of biomimetic technology. He began his career in robotics at Hughes Aircraft building snake-like robot manipulators. During graduate school, he worked for two years at the California Institute of Technology's Jet Propulsion Laboratory in the area of anthropomorphic telemanipulators, and he served two years as the director of UCLA Commotion Laboratory working in the area of collective robotics. After graduation, he joined the University of Illinois at Urbana-Champaign and worked in the area of visuomotor control and multi-robot systems.

Ralph Etienne-Cummings received his B.Sc. in Physics, 1988, from Lincoln University, Pennsylvania. He completed his M.S.E.E. and Ph.D. in electrical engineering at the University of Pennsylvania in 1991 and 1994, respectively. Currently, Dr. Etienne-Cummings is an assistant professor of electrical and computer engineering at the Johns Hopkins University, Baltimore, MD. His research interest includes mixed signal VLSI systems, computational sensors, computer vision, neuromorphic engineering, smart structures, mobile robotics and robotics-assisted surgery. He is a recipient of the NSF's CAREER and ONR's YIP Awards.

Mitra J. Hartmann received a B.S. in Applied and Engineering Physics from Cornell University, and a Ph.D. in Integrative Neuroscience from the California Institute of Technology. She is currently a post-doctoral scholar in the Center for Integrated Space Microsystems at Caltech's Jet Propulsion Laboratory. Her research interests include active sensing, neuromorphic sensory processing, sensorimotor control systems, and biomimetic robotics. She is a recipient of Caltech's Everhart Distinguished Graduate Student award.

Avis H. Cohen obtained her Ph.D. from Cornell University, and held post-doctoral positions at the Karolinska Institute in Stockholm with Sten Grillner, and at Washington University, St. Louis with Carl Rovained, developing the lamprey spinal cord preparation for the study of the central pattern generator for locomotion. She then spent 10 years at Cornell in an independent research position continuing her lamprey work, and extending it to mathematical modeling of the intersegmental coordinating system using dynamical systems models. In 1990 she moved to University of Maryland as an Associate Professor in the Biology Department where she has been ever since. She was the director of the fledgling Neuroscience and Cognitive Science Program at UMD, and is a member of the Institute for Systems Research. She is also one of the co-directors of the Telluride Summer Workshop on Neuromorphic Engineering.