Human Inspired Control of a Small Humanoid Robot in Highly Dynamic Environments or Jimmy Darwin Rocks the Bongo Board

Jacky Baltes, Chris Iverach-Brereton[✉], and John Anderson

University of Manitoba, Winnipeg, MB R3T2N2, Canada
chrisib@cs.umanitoba.ca
http://aalab.cs.umanitoba.ca/

Abstract. This paper describes three human-inspired approaches to balancing in highly dynamic environments. In this particular work, we focus on balancing on a bongo board - a common device used for human balance and coordination training - as an example of a highly dynamic environment. The three approaches were developed to overcome limitations in robot hardware. Starting with an approach based around a simple PD controller for the centre of gravity, we then move to a hybrid control mechanism that uses a predictive control scheme to overcome limitation in sensor sensitivity, noise, latency, and jitter. Our third control approach attempts to maintain a dynamically stable limit cycle rather than a static equilibrium point, in order to overcome limitations in the speed of the actuators. The humanoid robot Jimmy is now able to balance for several seconds and can compensate for external disturbances (e.g., the bongo board hitting the table). A video of the robot Jimmy balancing on the bongo board can be found at http://youtu.be/ia2ZYqqF-lw.

1 Introduction

This paper describes our research on active balancing reflexes for humanoid robots. Rapid progress in both hardware and software in recent years has led to impressive improvements in the performance of humanoid robots. For example, the soccer playing robots participating in the RoboCup competition [1] can walk and turn quickly, as well as stand up rapidly after falling. In the multi-event HuroCup competition [2], the world record in the sprint event (3 m walking forward followed by 3 m walking backward) has improved from 01:07.50 s. in 2009 to 00:25.50 s. in 2013. Similarly, the world record times in the marathon, which is traditionally held outdoors, improved from 37:30.00 over 42.195 m in 2007 to 13:24.39 over 120 m in 2013. Today, most humanoid robots have little difficulty traversing flat and even surfaces with sufficient friction.

The problem of traversing an irregular and potentially unstable surface, on the other hand, is still extremely difficult and remains without a general solution. Today's robots do not have sufficiently powerful actuators, nor enough sensors to be able to move over a rubble pile or similar environment.

In recent years, we have therefore focused on balancing in challenging, yet achievable environments. Examples are our robot Tao-Pie-Pie [3], which active balanced over a uneven balance field, and our ice and inline skating humanoid robot Jennifer [4], which demonstrated gaits that were stable on moving wheels and on ice. In winning this year's FIRA Hurocup [5] in the kid-size division, our robots demonstrated a broad range of achievements in adaptive humanoid motion, scoring highly in weightlifting, climbing, and sprinting, as well as soccer, while using the same unaltered robot in these and other events.

Balancing skills are central to all of these, as well as most other humanoid movement. In this paper, we describe our work toward balancing on a bongo board using a small humanoid robot (Jimmy, a DARwin-OP robot made by Robotis). A bongo board is a device commonly used in human training for balance and coordination, and consists of a small board that is placed on top of a cylindrical fulcrum. Figure 1 shows our humanoid robot Jimmy on top of the bongo board used in this work. Jimmy is a Robotis DARwIn-OP robot [6], standing approximately 45 cm tall. He has 22 Degrees Of Freedom (DOFs) and has a six-axis gyroscope/accelerometer located in the torso. Jimmy has been modified from the manufacturer's stock configuration with two single-DOF hands and four-point force (FSR) sensors in the feet.
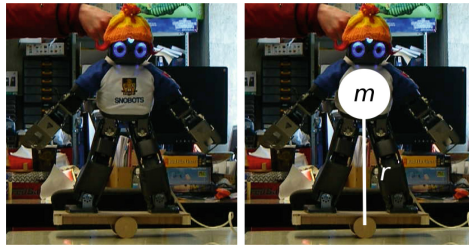

Fig. 1. Jimmy on the bongo board (left). The figure on the right shows how the bongo board corresponds to an inverted pendulum problem.

The bongo board's fulcrum can freely move left and right, forcing the robot to balance in those directions to keep the board from touching the ground on either side. Balancing on the bongo board is a non-trivial task even for humans. Moreover, because the fulcrum can move, shifting the centre of mass can allow the board to remain balanced and off the ground while shifting the fulcrum from side to side, and this and other tricks are used by human acrobats for entertainment purposes.

The remainder of this paper is organized as follows. Section 2 presents an analysis of the dynamics of the bongo board and shows the relationship to other inverted pendulum problems. Section 3 describes the design and implementation of our three control strategies for the bongo board. Section 3 also describes several challenges imposed by the robot hardware and how we overcame them. Section 4 provides a brief numerical evaluation of the three control strategies

we developed. Additional discussion appears in with Sect. 5, which also provides directions for future work.

2   Analysis and Related Work

This section gives a brief introduction to the dynamics of an inverted pendulum [7].

2.1   Dynamics of the Inverted Pendulum Problem

The dynamics of the inverted pendulum problem are well-studied and well understood and form the basis of many motion control algorithms for bipedal humanoid walking robots [8].

There has been a lot of theoretical work in the area of highly dynamic balancing [9–11], but practical implementations are still lacking. Anderson et al. describe an adaptive torque based approach [12] that is able to balance a humanoid robot on a simple see saw. In simulation, their approach is also able to balance a humanoid robot on the more challenging bongo board.

A similar system is described by Hyon [13] is able to balance a robot on a see saw in the presence of unknown disturbances.

2.2   Dynamics of the Bongo Board

The problem of balancing on a bongo board is similar to the cart and rod problem, as can be seen in Fig. 1. The robot can be modelled as a single point mass balancing on top of the board, and the goal is for the robot and board to balance without touching the ground or the robot falling off the board. In other words, the inverted pendulum system formed by the robot and the bongo board should balance.

The difference between the bongo board and the card and rod problem is that when balancing on a bongo board, (a) the pivot point of the robot will rotate along the circumference of the wheel, and (b) the position of the pivot point cannot be controlled directly - only indirectly by controlling the motion of the humanoid robot balancing on the board. Figure 2 shows the how the robot manipulates its limbs in order to control the inclination of the board.

The robot uses its legs and hips to provide coarse control over the inclination, elevation, and lateral position of the CoM with respect to the deck of the bongo board. Moving the shoulders provides fine control over the lateral position of the CoM, as well as applying torque to the system which can be used to provide angular control over the system. Note that from the perspective of the rider d – the distance between the point of contact between the deck and the wheel and the lateral position of the rider's CoM – is unknown.

As with a more traditional inverted pendulum problem, the inclination and angular velocity of the mass (in this case the rider and the deck) must be controlled in order for the system to remain in a stable position.
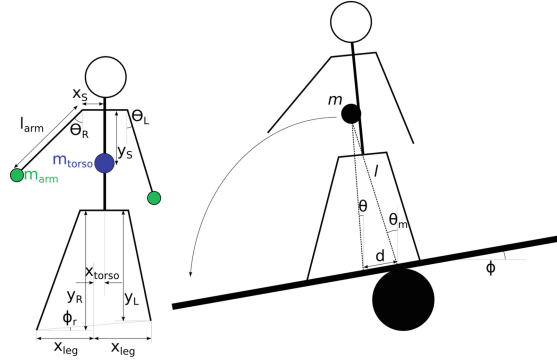
Fig. 2. A bongo board with the rider in an unstable position. The rider's CoM, m, is the weighted average of the CoMs of the torso and the arms. The robot's arms and legs can be adjusted independently to offset the torso laterally (x_{torso}) and vertically (y_L, y_R, ϕ_r). The angle between the point of contact between the bongo board's deck and wheel and m is given by θ_m. The distance between the point of contact and robot's CoM measured parallel with the deck is given by d. The CoM is located at height l above the deck.

3   Design and Implementation

Based on the analysis in the previous section, we began examining how people balance on a bongo board, and what considerations had to be made to adapt humanoid robot balancing to this task. Through the experimentation with a simple control regime, it became clear that significant complications arise with current robotic technology that are easily taken for granted in simple balancing tasks in humans. In particular, sensor noise, sensor latency, and actuator latency are major problems which required the development of more sophisticated control approaches. The three approaches that we moved through in our work are presented in the following subsections.

Previous research by Wang [14] has shown that Proportional-Derivative (PD) controllers are effective at controlling the angular velocity and inclination of an inverted pendulum. Because the bongo board is a similar problem to more traditional inverted pendulum problems we chose to use a PD controller as the basis of our approach.

3.1   Stiff-Upper-Lip Policy and Sensor Fusion

The Stiff-Upper-Lip control policy is directly inspired by the behaviour observed in humans balancing on a bongo board. The goal is to maintain the torso in an upright position and to compensate for the motions of the bongo board by moving the legs only and thus moving the centre of gravity. Figure 3 show stills from a video of a human using the stiff upper lip policy.

When using the Stiff-Upper-Lip policy a PD controller based on the robot's current inclination and angular velocity, as recorded by the robot's on-board

Fig. 3. Bongo Board: Stiff-Upper-Lip Policy. Note how the torso of the player is almost stationary and the legs compensate for the motion of the board.

sensors, is used to control the angular velocity and inclination of the torso, using the following control law:

$$\theta_{Torso} = K_p(\theta_{Torso}) + K_d(\dot{\theta}_{Torso})$$

The robot keeps its torso at a constant height relative to the deck and controls the inclination by extending and contracting its legs (y_L and y_R in Fig. 2).

The first problem with adapting this approach to current robot technology can be seen in Fig. 3 itself: it requires bending of the torso. The robot used in this work does not have the necessary DOF in the torso to execute this motion. Therefore, the necessary control can only be approximated by raising and lowering of the individual hip joints, as described above.

The second problem uncovered was that the gyroscope on the robot was not sensitive enough to detect any angular velocity until the robot was already moving at approximately $3°/s$, as can be shown in Fig. 4. To compensate for this problem we use the robot's three-axis accelerometer to measure the inclination. The angular velocity is then estimated using the difference between the present and previous inclination divided by the time between readings. The time between readings is governed by the serial connection between the robot's main processor and its sub-controller and is approximately 8 ms.

3.2   Do-The-Shake Policy and Predictive Control

In spite of overcoming the sensitivity issue of the gyroscope, the other two problems still remain: (a) latency and (b) jitter in the control. To deal with these, we added a one time-step prediction for the PD controller. The result and the error of the prediction of the inclination angle and the angular velocity can be seen in Figs. 5 and 6. The average error in the prediction of the inclination was $-0.002°$ with a standard deviation of 2.437. The average error in the prediction of the angular velocity was $0.0006°/s$ with a standard deviation of 2.931.

The prediction greatly improved the performance of Jimmy's balancing, but it was still limited by the slow speed of the actuators. Furthermore, the lack of a servo in the torso resulted in only a limited range of motion. However, shifting
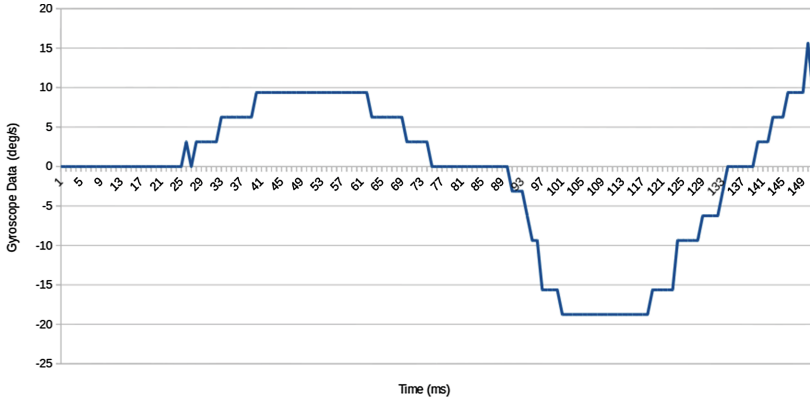
Fig. 4. Sensor Readings from the Y Plane Gyroscope Using the Stiff-Upper-Lip Policy. The gyroscope measures in approximately $3°/s$ increments, which was too coarse for balancing.
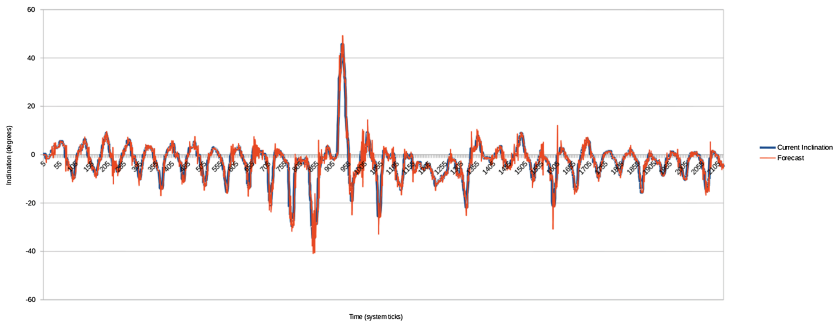

Fig. 5. Comparison between predicted and actual inclination angle
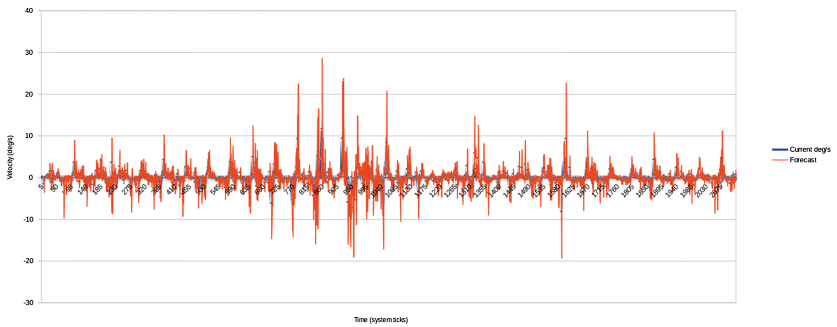

Fig. 6. Comparison between predictaed and actual angular velocity

the torso is not the only way for the robot to move its Center of Gravity (CoG). Figure 3 clearly shows that a human can also use his or her arms to balance on the board. We therefore extended the Stiff-Upper-Lip policy into a hybrid control scheme that moves the hips for coarse corrections and the arms for fine corrections to the CoG of the robot.

The hybrid controller was implemented by applying a correction to both arms and hips only when the error in angular velocity or inclination angle was above a threshold. In this case, the gain of the hip control was significantly larger than that of the arm controller. If the error in angular velocity or inclination angle was small, only the PD control for moving the arms side to side was used.

The following control law was used to calculate the torso angle θ_{Torso} and the displacement of the arms from the neutral position d_{Arms}.

$$\theta'_{Torso} = predicted(\theta_{Torso}, \dot{\theta}_{Torso})$$

Case 1 (Small inclination and angular velocity error):

$$d_{Arms} = Ka_p(\theta'_{Torso}) + Ka_d(\dot{\theta}'_{Torso})$$

Case 2 (Large inclination or angular velocity error):

$$\theta_{Torso} = Kh_p(\theta'_{Torso}) + Kh_d(\dot{\theta}'_{Torso})$$
$$d_{Arms} = Ka_p(\theta'_{Torso}) + Ka_d(\dot{\theta}'_{Torso})$$

The performance of the Do-The-Shake policy was better than that of the Stiff-Upper-Lip policy, but the robot was still not able to balance on its own continuously. The latency and jitter as well as the delay in execution of the correction commands was limiting the performance of the balancing of the robot.

3.3   The Lets-Sway Policy - Dynamically Stable Balancing

The latency in the system meant that it was impossible for Jimmy to correct for tilting of the bongo board quickly enough. By watching humans on the bongo board it became apparent that this is also a problem for humans. Instead of trying to maintain the board in a statically stable position, humans appear to enter a dynamically stable limit cycle, continuously swaying left to right.

The Lets-Sway control is similar to the Do-The-Shake Policy, but instead of attempting to maintain an inclination of zero degrees and an angular velocity of zero degress, the controller is tracking a sine curve of the inclination angle. That is, the robot Jimmy continuously moves the CoG by swaying with the hips. Even though each position along the path is statically unstable, the resulting limit cycle results in dynamically stable behaviour. Dynamically stable limit cycles have been used previously when trying to stabilize a humanoid robot [15].

A small PD controller with only moderate gain is controlling this movement. Similar to the Do-The-Shake Policy, the arms provide fine corrections for the centre of gravity.
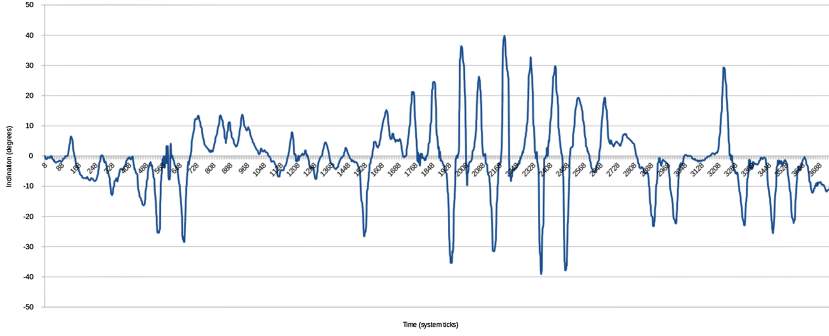
Fig. 7. Inclination Angle of the Do-The-Shake Policy. Using both arms and legs to control the CoG results in smoother balancing than the Stiff-Upper-Lip Policy.

$$\theta'_{Torso} = predicted(\theta_{Torso}, \dot{\theta}_{Torso})$$
$$\theta_{Desired} = sin(\omega t)$$
$$\theta_{Torso} = Kh_p(\theta'_{Board} - \theta_{Desired}) + Kh_d(\dot{\theta}_{Board} - \dot{\theta}'_{Desired})$$
$$d_{Arms} = Ka_p(\theta'_{Board} - \theta_{Desired}) + Ka_d(\dot{\theta}_{Board} - \dot{\theta}'_{Desired})$$

The Lets-Sway policy led to much better performance as can be seen when comparing the accelerometer data from Figs. 7 and 8. The resulting motion is more stable and regular as compared to that of the Do-The-Shake policy.

This was also apparent when watching the performance of the robot. The robot is able to balance for several cycles without help and can compensate if the board hits the table. A video of Jimmy rocking the bongo board using the Lets-Sway policy can be found on youtube (http://www.youtube.com/watch?v=ia2ZYqqF-lw).

4   Evaluation

To compare all three methods we developed we used the robot's average inclination and angular velocity. Under ideal circumstances the system should preserve an angular velocity and inclination near zero. Figures 9 and 10 show the inclination and angular velocity recorded when using each of the three control policies. Table 1 shows the average incliation and angular velocity as well as the standard deviations for each control policy.

From Fig. 9 we can see that the let's-sway policy exhibits cyclic spikes in its inclination; peaks and valleys occur at roughly regular intervals and are generally similar in magnitude. Comparitively the other two policies both exhibit much more extreme inclinations and show less regular cycles. Similarly, from Fig. 10 we
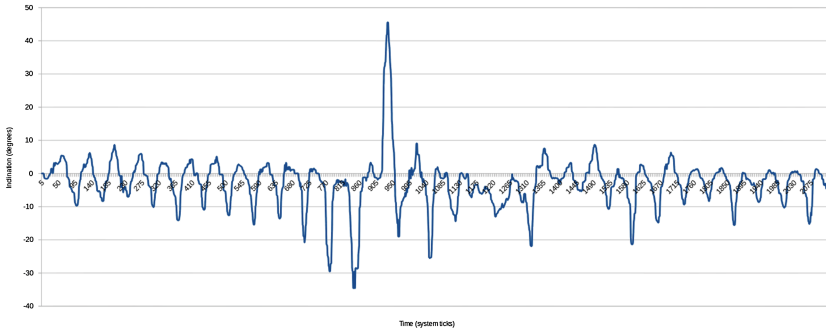
Fig. 8. Lets-Sway Policy. The robot attempts to maintain a dynamically stable limit cycle by moving its hips side to side. The arms are used for fine grained corrections.
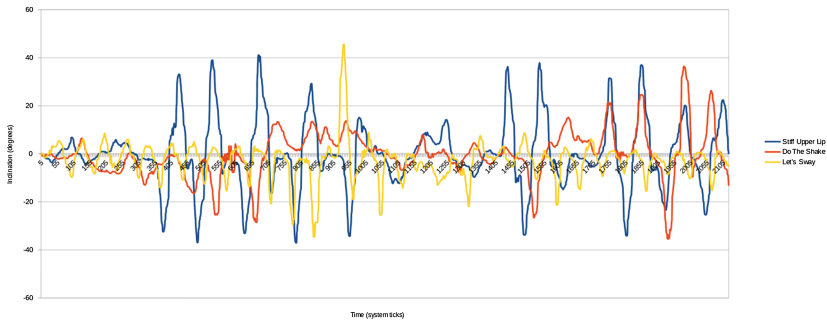

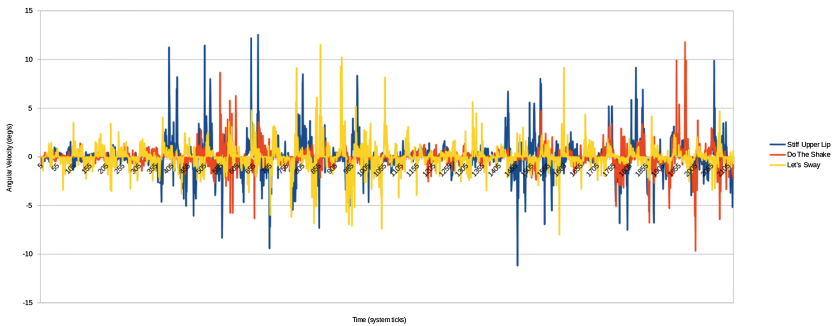**Fig. 9.** The robot's recorded inclination when balancing using all three control policies.

Fig. 10. The robot's recorded angular velocity when balancing using all three control policies.

Table 1. The average inclination and angular velocity recorded when balancing with each control policy

	Inclination	σ^2	Angular velocity	σ^2
Stiff Upper Lip	0.00415	13.9	−4.17E-18	1.57
Do The Shake	−0.178	9.45	0.00612	0.995
Let's Sway	−2.82	7.47	0.00244	1.23

can see that the let's-sway policy demonstrates fewer extreme peaks and valleys in its angular velocity than the other policies.

Of the three policies the stiff-upper-lip is the least stable; the robot's inclination oscillates wildly back and forth, frequently striking the ground on each side. The robot's angular velocity is frequently very high; more than 5°/s.

The introduction of regular oscillations in the let's-sway policy appears to reduce the angular velocity recorded in the robot's torso, though this reduction is not statistically significant. When using the let's-sway policy the robot's inclination is maintained within a narrower range, suggesting that the introduction of a dynamically stable oscillation does improve the robot's ability to balance on the bongo board. This corresponds with our qualitative observations while testing the policies; the let's-sway policy appeared to be the most stable.

5   Conclusions and Future Work

The research described in this paper is still work in progress. The robot is currently able to balance for several seconds, but the board will often hit the table. This is due to the relatively small diameter of the supporting wheel, which means that the robot has very little time to correct and reverse the motion before the board hits the table. We are countering this by increasing the diameter of the supporting wheel by 1 cm.

We are currently in the process of evaluating the performance of our control approach to deal with unknown external disturbances. The experiments will include perturbation of the robot while balancing on the bongo board as well as sudden pushes to the robot while walking on a flat and even surfaces.

There are many possible directions for future research. We plan on adding visual feedback of the optical flow in the image to improve the robot's estimation of its inclination angle and angular velocity.

Furthermore, both the inverted pendulum and the cart and rod problem are textbook examples for applying machine learning techniques to solve control problems. In particular, reinforcement learning is able to solve these types of problem efficiently. We plan to apply reinforcement learning to the bongo board problem.

Another direction for future research is team balancing. The goal is for two robots to balance on a single bongo board, one robot to the right and one to the

left of the wheel. Mathematical analysis shows that the combination of the two robots can be viewed as a single system with two separated actuators.

Finally, there are more complicated balancing devices than a bongo board on which these approaches could be adapted. The fulcrum of a bongo board is a cylinder, making banking motion the main focus for balancing, along with translation (sliding the board along the fulcrum). While it is still possible for the robot to fall forward or backward off the bongo board, the board itself is not intended to force movement in these dimension. A Wobble board, on the other hand, allows spherical motion across the fulcrum, making pitching and yawing motions just as important as the banking movements encompassed by a bongo board. On the other hand, a wobble board has a stationary base for its fulcrum, making it still somewhat restricted compared to a device with a free-moving spherical fulcrum.

References

1. Chen, X., Stone, P., Sucar, L.E., van der Zant, T. (eds.): RoboCup 2012. LNCS, vol. 7500. Springer, Heidelberg (2013)
2. Baltes, J., Tu, K.Y., Lip, S.L.: HuroCup Competition. FIRA, September 2013
3. McGrath, S., Anderson, J., Baltes, J.: Model-free active balancing for humanoid robots. In: Iocchi, L., Matsubara, H., Weitzenfeld, A., Zhou, C. (eds.) RoboCup 2008. LNCS, vol. 5399, pp. 544–555. Springer, Heidelberg (2009)
4. Iverach-Brereton, C., Winton, A., Baltes, J.: Ice skating humanoid robot. In: Herrmann, G., Studley, M., Pearson, M., Conn, A., Melhuish, C., Witkowski, M., Kim, J.-H., Vadakkepat, P. (eds.) TAROS-FIRA 2012. LNCS, vol. 7429, pp. 209–219. Springer, Heidelberg (2012)
5. Omar, K. (ed.): FIRA 2013. CCIS, vol. 376. Springer, Heidelberg (2013)
6. Robotis: Darwin-op, 09 2013. http://support.robotis.com/en/techsupport_eng.htm#product/darwin-op.htm
7. Pratt, J.E.: Exploiting Inherent Robustness and Natural Dynamics in the Control of Bipedal Walking Robots. Ph.D. thesis, MIT Computer Science (2000)
8. Sugihara, T., Nakamura, Y., Inoue, H.: Real-time humanoid motion generation through zmp manipulation based on inverted pendulum control. In: Proceedings of the IEEE International Conference on Robotics and Automation ICRA 2002, vol. 2, pp. 1404–1409 (2002)
9. Pratt, J., Carff, J., Drakunov, S., Goswami, A.: Capture point: A step toward humanoid push recovery. In: 2006 6th IEEE-RAS International Conference on Humanoid Robots, pp. 200–207 (2006)
10. Park, J., Haan, J., Park, F.: Convex optimization algorithms for active balancing of humanoid robots. IEEE Trans. Robot. **23**(4), 817–822 (2007)
11. Hyon, S., Cheng, G.: Gravity compensation and full-body balancing for humanoid robots. In: 2006 6th IEEE-RAS International Conference on Humanoid Robots, pp. 214–221 (2006)
12. Anderson, S., Hodgins, J.: Adaptive torque-based control of a humanoid robot on an unstable platform. In: 2010 10th IEEE-RAS International Conference on Humanoid Robots (Humanoids), pp. 511–517 (2010)
13. Hyon, S.: Compliant terrain adaptation for biped humanoids without measuring ground surface and contact forces. IEEE Trans. Robot. **25**(1), 171–178 (2009)

14. Wang, J.J.: Simulation studies of inverted pendulum based on PID controllers. Simul. Modell. Pract. Theor. **19**(1), 440–449 (2011). Modeling and Performance Analysis of Networking and Collaborative Systems
15. Goswami, A., Espiau, B., Keramane, A.: Limit cycles and their stability in a passive bipedal gait. In: Proceedings of the 1996 IEEE International Conference on Robotics and Automation, 1996. vol. 1, pp. 246–251. IEEE (1996)