

Human Inspired Control of a Small Humanoid Robot in Highly Dynamic Environments

or

Jimmy DARwIn Rocks the Bongo Board

Jacky Baltes
Dept. of Computer Science
University of Manitoba
Winnipeg, MB, Canada R3T-2N2
Email: jacky@cs.umanitoba.ca

Chris Iverach-Brereton
Dept. of Computer Science
University of Manitoba
Winnipeg, MB, Canada R3T-2N2

John Anderson
Dept. of Computer Science
University of Manitoba
Winnipeg, MB, Canada R3T-2N2

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://www.youtube.com/watch?v=ia2ZYqqF-lw>.

I. 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 meters walking forward followed by 3m walking backward) has improved from 01:07.50 sec. in 2009 to 00:25.50 sec. in 2013. Similarly, the world record times in the marathon, which is traditionally held outdoors, improved from 37:30.00 over 42.195m in 2007 to 13:24.39 over 120m 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.

The 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.

Jimmy is a Robotis DARwIn OP [6], a 45cm tall humanoid robot that weighs about 4kg. Robotis MX-28 servo motors power 20 degrees of freedom (DOF). Higher level processing is implemented on a FitPc2 processor board which features an 1.6GHz Intel Atom processor and 1GB of RAM. For active balancing, the robot includes a three axis gyroscope and a three axis accelerometer in the torso. Other sensors include a camera and two microphones. We extended the basic DARwIn OP with two force sensors (FSR) sensors in the feet and replaced




Fig. 1. Jimmy on the Bongo Board

the hands with two grippers.

The DARwIn OP uses a two-tiered distributed control architecture for position control of the joints. The higher level processing system sends position and movement commands via a serial link to a low level servo controller based on an ARM Cortex M3 processor running at 72 MHz.

The challenges for a humanoid robot presented by a Bongo board are very similar to those in humans: sensors must detect in real time when changes in motion are necessary to preserve or change the board's positions, and the computation necessary to process sensor values and calculate necessary movement changes must be done quickly enough, in combination with the time needed for actuators to physically move the robot, to maintain the balance of both the robot and the board. This is complicated by limitations in sensor accuracy, noise in the environment or sensors, and latency between the state of the world and sensor and actuator values, and the speed of the actuators themselves. The main focus of this paper is the presentation of three approaches to balancing on a device such as this bongo board while dealing with these complications. These control strategies were inspired by observing humans balancing on the board. We begin with a simple PD controller for the centre of gravity, and then move to a hybrid controller using a predictive approach to overcome limitations in sensor sensitivity, noise, latency, and jitter. We then further move to an approach that attempts to maintain a dynamically stable limit cycle as opposed to a static equilibrium point, to overcome limitations in the speed of the actuators.

The remainder of this paper is organized as follows. Section II presents an analysis of the dynamics of the Bongo

board and shows the relationship to other inverted pendulum problems. Section III describes the design and implementation of our three control strategies for the Bongo board. Section III also describes several challenges imposed by the robot hardware and how we overcame them. Additional discussion appears in with Section IV, which also provides directions for future work.

II. ANALYSIS AND RELATED WORK

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

A. 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].

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.

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.

III. 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.




Fig. 2. 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.

A. Stiff-Upper-Lip Policy and Sensor Fusion

Stiff-Upper-Lip is a control policy that is similar to 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 2 shows a human using the Stiff-Upper-Lip policy.

The PD control of the angular velocity and inclination has previously shown good performance in inverted pendulum problems [14]. Therefore, we implemented the Stiff-Upper-Lip Policy using a simple proportional derivative (PD) controller based on the current inclination and angular velocity of the robot. The robot maintains the torso at a constant height above the board and calculates a desired torso angle using the following control law:

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

The first problem with adapting this approach to current robot technology can be seen in Figure 2 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.

Another problem became readily apparent when watching the robot. The robot reacted correctly to a disturbance, but its reaction was too late. Further investigation revealed that as can be seen in Fig. 3, the gyroscope on Jimmy is not sensitive enough to register an angular velocity when the robot is starting to fall to one side or another. By the time the gyroscope registered any tilt at all, it was too late to correct for the tilt, as the robot could not move its legs fast enough to prevent the board from striking the table.

This sensitivity problem is aggravated by the architecture of the DARwIn-OP. The gyroscope and accelerometer sensors are connected to a small embedded microcontroller (CM-730), but all balancing control is executed on the FitPc2 main processor board. The CM730 and the FitPc2 board are connected via a slow serial connection. According to Robotis specifications, the maximum speed of the serial link is 2 MBps, but in our tests we found that there was too much interference on the

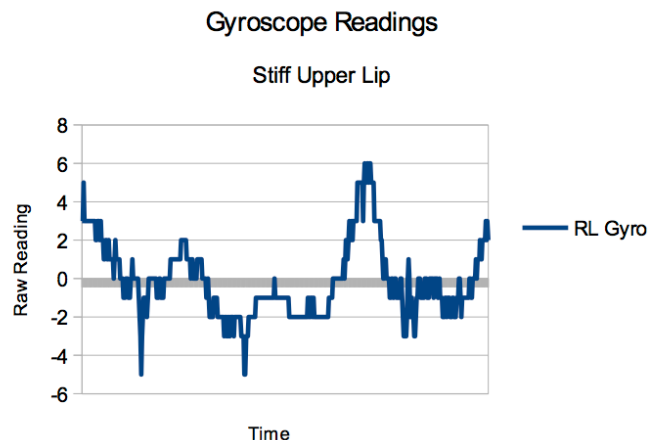


Fig. 3. Sensor Readings from the Y Plane Gyroscope Using the Stiff-Upper-Lip Policy. The gyroscope does not register any movement in the beginning until the angular velocity is already quite high.

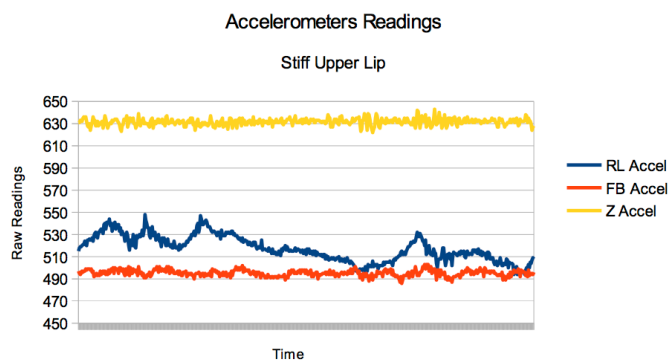


Fig. 4. Sensor Readings from the Accelerometer Using the Stiff-Upper-Lip Policy. The accelerometers react to small changes in inclination angle accurately, but as the angular velocity increases, the measurements become more noisy.

bus at that speed. Therefore, we limited the speed of the serial link to 1 MBps. To be able to react to disturbances, the main processor board needs to request readings from the CM-730 and the CM-730 transmits the sensor reading back to the FitPc2. This introduces a latency of at least 8 ms, but sometimes as high as 16 ms into the system. The jitter makes accurate control for balancing in highly dynamic environments challenging.

Since Jimmy is also equipped with a 3-axis accelerometer, the accelerometer readings were included to estimate inclination angle and angular velocity. As can be seen in Fig. 4, when moving slowly the accelerometers provide a mean of measuring the inclination angle by calculating the angle of gravity. Once the robot starts to rotate quickly, the accelerometer readings become noisy. The accelerometer and gyroscope can compensate for the shortcomings of the other sensor.




Fig. 5. Comparison between Predicted and Actual Inclination Angle




Fig. 6. Comparison between Predicted and Actual Angular Velocity

B. 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 Fig. 5 and Fig. 6.

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 the torso is not the only way for the robot to move its Center of Gravity (CoG). Fig. 2 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} .

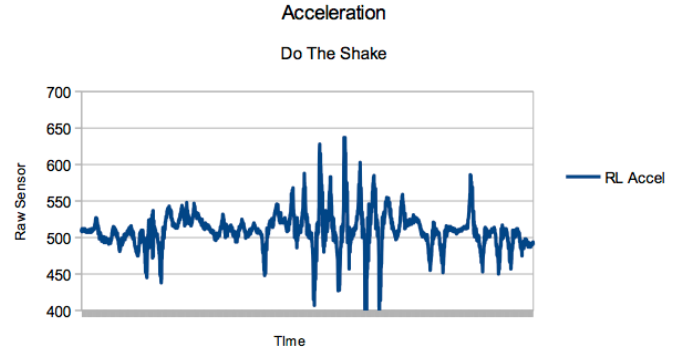


Fig. 7. Inclination Angle of the Do-The-Shake Policy. Using both arms and legs to control the CoG results in smoother balancing than the Stiff-Upper-Lip Policy.

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

Case 1 (Small inclination and angular velocity error):

$$d_{Arms} = K a_p(\theta'_{Torso}) + K a_d(\dot{\theta}'_{Torso})$$

Case 2 (Large inclination or angular velocity error):

$$\theta_{Torso} = K h_p(\theta'_{Torso}) + K h_d(\dot{\theta}'_{Torso})$$

$$d_{Arms} = K a_p(\theta'_{Torso}) + K a_d(\dot{\theta}'_{Torso})$$

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.

C. 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 degrees, 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.

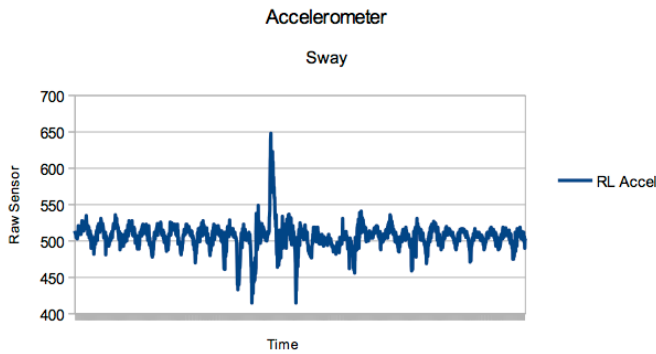


Fig. 8. Lets-Sway Policy. The robot attempts to maintain a dynamically stable limit cycle by moving its hips side to side. The arms are used for fine grained corrections.

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

$$\theta_{Desired} = \sin(\omega t)$$

$$\theta_{Torso} = Kh_p(\theta'_{Board} - \theta_{Desired}) + Kh_d(\dot{\theta}_{Board} - \dot{\theta}'_{Desired})$$

$$d_{Arms} = Ka_p(\theta'_{Board} - \theta_{Desired}) + Ka_d(\dot{\theta}_{Board} - \dot{\theta}'_{Desired})$$

The Lets-Sway policy led to much better performance as can be seen when comparing the accelerometer data from Fig. 7 and Fig. 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>).

IV. 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] X. Chen, P. Stone, L. E. Sucar, and T. van der Zant, Eds., *RoboCup 2012: Robot Soccer World Cup XVI*, ser. Lecture Notes in Computer Science. Springer Verlag, 2013, vol. 7500.
- [2] J. Baltes, K.-Y. Tu, and S. L. Lip, *HuroCup Competition*, FIRA, Sept. 2013. [Online]. Available: <http://www.facebook.com/groups/hurocup>
- [3] S. McGrath, J. Anderson, and J. Baltes, "Model-free active balancing for humanoid robots," in *Proceedings of RoboCup-2008: Robot Soccer World Cup XII*, Suhou, China, L. Iocchi, H. Matsubara, A. Weitzenfeld, and C. Zhou, Eds., 2008.
- [4] C. Iverach-Brereton, A. Winton, and J. Baltes, "Ice skating humanoid robot," in *Advances in Autonomous Robotics*, G. Herrmann, M. Studley, M. Pearson, A. Conn, C. Melhuish, M. Witkowski, J.-H. Kim, and P. Vadakkepat, Eds., vol. 7429. Springer Berlin/Heidelberg, 2012, pp. 209 – 219.
- [5] K. Omar, M. J. Nordin, P. Vadakkepat, A. S. Prabuwo, S. N. H. S. Abdullah, J. Baltes, S. M. A. W. Z. W. Hassan, and M. F. Nasrudin, Eds., *Intelligent Robotics Systems: Inspiring the NEXT*, ser. Communications in Computer and Information Science. Springer Verlag, 2013, vol. 376.
- [6] Robotis Company, "DARwIn-OP," WWW, 09 2013. [Online]. Available: http://support.robotis.com/en/techsupport_eng.htm#product/darwin-op.htm
- [7] J. E. Pratt, "Exploiting inherent robustness and natural dynamics in the control of bipedal walking robots," Ph.D. dissertation, MIT Computer Science, 2000.
- [8] T. Sugihara, Y. Nakamura, and H. Inoue, "Real-time humanoid motion generation through zmp manipulation based on inverted pendulum control," in *Robotics and Automation, 2002. Proceedings. ICRA '02. IEEE International Conference on*, vol. 2, 2002, pp. 1404–1409 vol.2.
- [9] 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*, 2006, pp. 200–207.
- [10] J. Park, J. Haan, and F. Park, "Convex optimization algorithms for active balancing of humanoid robots," *Robotics, IEEE Transactions on*, vol. 23, no. 4, pp. 817–822, 2007.
- [11] S. Hyon and G. Cheng, "Gravity compensation and full-body balancing for humanoid robots," in *Humanoid Robots, 2006 6th IEEE-RAS International Conference on*, 2006, pp. 214–221.
- [12] S. Anderson and J. Hodgins, "Adaptive torque-based control of a humanoid robot on an unstable platform," in *Humanoid Robots (Humanoids), 2010 10th IEEE-RAS International Conference on*, 2010, pp. 511–517.
- [13] S. Hyon, "Compliant terrain adaptation for biped humanoids without measuring ground surface and contact forces," *Robotics, IEEE Transactions on*, vol. 25, no. 1, pp. 171–178, 2009.

- [14] J.-J. Wang, "Simulation studies of inverted pendulum based on PID controllers," *Simulation Modelling Practice and Theory*, vol. 19, no. 1, pp. 440 – 449, 2011, Modeling and Performance Analysis of Networking and Collaborative Systems. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1569190X10001796>
- [15] A. Goswami, B. Espiau, and A. Keramane, "Limit cycles and their stability in a passive bipedal gait," in *Robotics and Automation, 1996. Proceedings., 1996 IEEE International Conference on*, vol. 1. IEEE, 1996, pp. 246–251.