References for 15-869 Feb 11, 2014

The overall topic in class today was PD control.   If you are still looking for a final project, it is very straightforward to experiment with different controllers using one of the off the shelf simulators such as ODE.   You might design a project to test some of the pros and cons of the ideas presented in the papers below.

Proportional-Derivative (PD) controllers are feedback controllers. One of the first things to realize is that they work by creating a response to errors and constraint violations.   In this regard, they are somewhat analogous to penalty based methods for collision and contact.   As such, if they are used to respond to constraint violations (e.g., joint limits) or to closely track captured motion data, they have many of the same problems as penalty based contact response methods.   In particular, avoiding large errors means cranking up the stiffness, which in turn means running your simulations with small timesteps.

PD controllers can be used very effectively, however, and without tiny timesteps, if they are thought of not for managing constraints or eliminating errors, but instead thought of as guidance, suggestions, or handles for creating a desired motion.   In this latter case, the user designs a controller with the understanding that the setpoints used to create the motion are not actual poses of the motion itself, but instead are control handles that give the direction that is needed to get the character to where it needs to go.   Excellent examples of this idea are found here:

Jessica K. Hodgins, Wayne L. Wooten, David C. Brogan, and James F. O'Brien. **Animating human athletics**. In *Proceedings of SIGGRAPH 95,* pages 71–78, August 1995. http://graphics.cs.cmu.edu/projects/animating_human_athletics/animatinghumanatheltics/Animating%20Human%20Athletics/Animating%20Human%20Athletics.html

Yin, KangKang, Kevin Loken, and Michiel van de Panne. "**Simbicon: Simple biped locomotion control**." In *ACM Transactions on Graphics (TOG)*, vol. 26, no. 3, p. 105. ACM, 2007. https://www.cs.ubc.ca/~van/papers/Simbicon.htm

Sok, Kwang Won, Manmyung Kim, and Jehee Lee. "**Simulating biped behaviors from human motion data**." In *ACM Transactions on Graphics (TOG)*, vol. 26, no. 3, p. 107. ACM, 2007. http://mrl.snu.ac.kr/research/ProjectSimulBiped/SimulBiped.html

The strategy pursued by the first two of these papers is reminiscent of the Equilibrium Point Hypothesis in motor control:

Gomi, Hiroaki, and Mitsuo Kawato. "**Equilibrium-point control hypothesis examined by measured arm stiffness during multijoint movement**." *Science* 272, no. 5258 (1996): 117-120. http://e.guigon.free.fr/rsc/article/GomiKawato96.pdf

This quote from Wikipedia gives a nice brief introduction to the idea behind the Equilibrium Point Hypothesis:

> In the Equilibrium Point hypothesis, all movements are generated by the nervous system through a gradual transition of equilibrium points along a desired trajectory. "Equilibrium point" in this sense is taken to mean a state where a field has zero force, meaning opposing muscles are in a state of balance with each other, like two rubber bands pulling the joint to a stable position. Equilibrium point control is also called "threshold control" because signals sent from the CNS to the periphery are thought to modulate the threshold length of each muscle. In this theory, motor neurons send commands to muscles, which changes the force–length relation within a muscle, resulting in a shift of the system's equilibrium point. The nervous system would not need to directly estimate limb dynamics, but rather muscles and spinal reflexes would provide all the necessary information about the system's state.
> http://en.wikipedia.org/wiki/Degrees_of_freedom_problem#Equilibrium_point_hypothesis_and_threshold_control

The strategy pursued by the third of the animation papers above (Sok, Won, Kim, and Lee 2007) is reminiscent of Feedback Error Learning, developed in the neuroscience community as an explanation of how the brain may learn to compensate for different effects. The actual process of feedback error learning is simple. You take your feedback force and torque terms obtained from one run of a repetitive task, scale them by a discount factor, and incorporate them as anticipatory feedforward signals for the next run. Many repetitions of this process should lead to a controller with substantially less error in accomplishing the desired task.

Kawato, Mitsuo, and Hiroaki Gomi. "**A computational model of four regions of the cerebellum based on feedback-error learning**." *Biological cybernetics* 68, no. 2 (1992): 95-103. http://link.springer.com/content/pdf/10.1007/BF00201431.pdf

Anticipatory feedforward terms do not have to be computed from feedback error learning, however. Instead, they can often be computed directly from the input motion without the learning process. The following nice paper describes how this works, and how the feedforward terms allow the user to separate out the accuracy of tracking from the overall stiffness of the character in responding to perturbations.

Yin, KangKang, Michael B. Cline, and Dinesh K. Pai. "**Motion perturbation based on simple neuromotor control models**." In *Computer Graphics and Applications, 2003. Proceedings. 11th Pacific Conference on*, pp. 445-449. IEEE, 2003.
http://www.comp.nus.edu.sg/~kkyin/papers/Yin_PG03.pdf

The incorporation of partial feedforward terms is also a straightforward way to relieve the PD controller of the responsibility of handling things like gravity compensation or coordinated joint motion, which our body knows how to handle very well. Even with only partial feedforward terms added to the PD control equation, stiffness of response to perturbations can be separated from tracking accuracy to a large extent. The sample paper that I used to outline a PD based control algorithm with partial feedforward terms (gravity compensation, wrist motion compensation) is this one:

Pollard, Nancy S., and Victor Brian Zordan. "**Physically based grasping control from example**." In *Proceedings of the 2005 ACM SIGGRAPH/Eurographics symposium on Computer animation*, pp. 311-318. ACM, 2005.
http://www.ri.cmu.edu/publication_view.html?pub_id=5219

For some more insight and general background, the paper below discusses a few misconceptions of PD control. It is a short read and has some interesting points, including making the point that PD control in its simple form may not be a good way to hit specific desired keyframes. The difficulty of hitting precise key poses or constraints using PD control explains why the very first (and very successful) papers mentioned in these notes consider the desired pose and velocity fed to the PD controller not as constraints that must be met, but more as suggestions or guides that are used to drive the resulting motion.

Allen, Brian F., and Petros Faloutsos. "**Misconceptions of PD control in animation**." In *Proceedings of the ACM SIGGRAPH/Eurographics Symposium on Computer Animation*, pp. 231-234. Eurographics Association, 2012.
http://dl.acm.org/citation.cfm?id=2422389

A PD controller that does meet artist designed key poses and velocities can be created, however, in the manner shown in the paper that follows. By giving up the opportunity for the user to define their own stiffness and damping parameters, the authors gain freedoms that can be manipulated to generate a PD controller that will hit particular position and velocity setpoints.

Allen, Brian F., Michael Neff, and Petros Faloutsos. "**Analytic proportional-derivative control for precise and compliant motion**." In *Robotics and Automation (ICRA), 2011 IEEE International Conference on*, pp. 6039-6044. IEEE, 2011.
http://www.brianfosterallen.com/files/ICRA2011_AnalyticPD.pdf

The following paper presents one more point of view. This paper takes the opinion that sometimes you may want PD control with high stiffness. It then goes on to show how an algorithm can manage high stiffnesses and large timesteps simultaneously by looking ahead to the next state of the system, somewhat analogous to implicit integration approaches. This approach can potentially be very interesting not only for high stiffness PD servo applications (e.g., joint limits and motion tracking), but also for collision and contact response, although in the latter case there are still details to be worked out (e.g., don't let the object ever pull on the ground when in contact).

Tan, Jie, Karen Liu, and Greg Turk. "**Stable proportional-derivative controllers**." *Computer Graphics and Applications, IEEE* 31, no. 4 (2011): 34-44.
http://www.cc.gatech.edu/~jtan34/project/spd.html

This paper yet a gives a different perspective in that it takes an impulse based approach to joint control. In addition, it has an insightful introduction that is well worth reading for general perspective on the use of PD control for character animation:

Weinstein, R., Guendelman, E. and Fedkiw, R. *Impulse Based Control of Joints and Muscles*. IEEE Transactions on Visualization and Computer Graphics, vol. 14, no. 1, pp. 37-46, Jan/Feb, 2008
http://rachelpetterson.com/pdcontrol.html

For historical background, this paper brought the control laws used by Marc Raibert and Jessica Hodgins to the graphics community. These algorithms have been further extended by Marc Raibert through Boston Dynamics into controllers for Big Dog and the Atlas robot used in the DARPA Robotics Challenge. We watched the accompanying move "On the Run."

Raibert, Marc H., and Jessica K. Hodgins. "**Animation of dynamic legged locomotion**." In *ACM SIGGRAPH Computer Graphics*, vol. 25, no. 4, pp. 349-358. ACM, 1991.

Among the very first animation papers to discuss physically based control for character animation are:

Armstrong, William W., and Mark W. Green. "**The dynamics of articulated rigid bodies for purposes of animation**." *The Visual Computer* 1, no. 4 (1985): 231-240.
http://link.springer.com/article/10.1007/BF02021812

Wilhelms, Jane P., and Brian A. Barsky. "**Using dynamic analysis to animate articulated bodies such as humans and robots**." In *Computer-Generated Images*, pp. 209-229. Springer Japan, 1985.
http://link.springer.com/chapter/10.1007/978-4-431-68033-8_19