Mystery World

One of the most successful techniques to arise from modern control theory works in the following way. Using current measurements of states of the system, control values are determined that make the output of the system follow a desired trajectory over a specified interval of time into the future. Assuming sufficient computational power and a sufficiently accurate model of the system, numerical algorithms are used to compute these control values. The computations are usually generated by an optimization method and can take into account constraints that must be satisfied, for example, on the states and controls. Of the computed control values, only the first value is implemented. The procedure is repeated at the next time instant and so forth. This method is known as either receding horizon control or model predictive control.

A receding horizon controller does not necessarily resemble a PID control law or even an LQG or H_∞ control law. In effect, the control law is an extremely nonlinear feedback control law, and there may be no way to represent it by an explicit formula. Nevertheless, in 1988, Elmer Gilbert and Sathiya Keerthi showed that, under fairly mild assumptions, receding horizon control can stabilize a system; it is, after all, a closed-loop control strategy. Consequently,

Digital Object Identifier 10.1109/MCS.2011.940722
Date of publication: 16 May 2011

Contributors

Yang Wang (right) with his parents in Wu Xi, China.

Stephen Boyd (left) and Jacob Mattingley (right) in Tokyo, Japan.

Nalin Chaturvedi and his wife Devashree enjoying La Jolla beach in San Diego.

Amit and Minakshi Sanyal.

receding horizon control rests on a solid foundation.

A disadvantage of receding horizon control is that it cannot easily account for model uncertainty, and thus it shifts the paradigm of control from “the reason for feedback is uncertainty” to “the reason for model-based optimization is constraints.”

The feature article by Jacob Mattingley, Yang Wang, and Stephen Boyd focuses on optimization methods for receding horizon control that have improved efficiency. Their approach, based on convex optimization and facilitated by automatic code generation, can provide dramatic improvements in execution times, from, say, seconds to milliseconds. For embedded systems, this speedup can open the door to new applications.

The next feature of this issue focuses on the age-old problem of modeling rotations. If you’ve ever studied the kinematics of a car, ship, aircraft, or spacecraft, you know that rotational motion is a central feature of the model. The most obvious approach is to use 3×3 matrices that are orthogonal and have determinant equal to one. These matrices comprise the Lie group SO(3). But you may also know that rotations can be modeled by Euler angles, for example, yaw, pitch, and roll. Yet another representation of rotation is given by the Rodrigues formula, which depends on the axis of rotation (the eigenaxis) and the angle about that axis (the eigenangle). Scaling the eigenaxis by a trigonometric function of the eigenangle gives rise to alternative representations. For example, the Euler parameters are obtained by scaling the eigenaxis by the sine of half the eigenangle and appending the cosine of half the eigenangle; the resulting four-vector is a representation of the unit quaternions. Analogously, scaling the eigenaxis by the tangent of half the eigenangle yields the Gibbs parameters, which are also called Rodrigues parameters.

The choice of representation is often hotly debated. First, there is the question of nonlinearity. We think of position as the integral of velocity under translational motion. For rotational motion, however, the relationship between attitude and angular velocity is much more complicated.

YuMing and his wife Hong celebrating the graduation of their son Chuan from MIT.

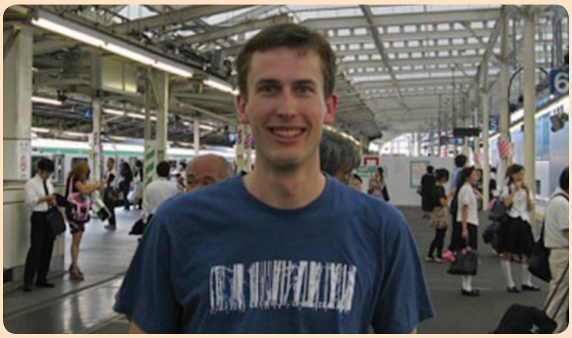

Jacob Mattingley at the Yokohama Station in Japan.

Margaret and Harris McClamroch.

For rotation matrices, Poisson's equation relates the angular velocity vector to the rotation matrix. Using Euler angles, Euler parameters, or Gibbs parameters, the angular velocity vector can be expressed in terms of the derivatives of these quantities, and a differential equation can be integrated to relate the rate of change of these quantities to the attitude.

Another issue concerns the "global" nature of each representation. Euler angles are fairly easy to visualize, and they can be used to represent all attitudes. However, at certain attitudes, two of the rotation axes are aligned, and this loss of a degree of freedom translates into the inability to represent rotation about another axis. In effect, the "virtual gimbals" are locked due to the fact that the mathematical representation of the rotation matrix is defective. One way around this difficulty is to use Euler parameters, which comprise four parameters confined to the unit sphere in four-dimensional space. This representation is lower dimensional—and thus more computationally efficient—than the rotation matrix itself, which is parameterized by nine parameters satisfying six constraints. The drawback of Euler parameters is the fact that they double cover SO(3), that is, for each matrix in SO(3), there is a pair of Euler parameters that represent the same attitude. This pair is located at antipodal, that is, diametrically opposite, points on the hypersphere of Euler parameters embedded in four-dimensional Euclidean vector space.

Double covering would seem to be innocuous, but it can cause problems. The most immediate issue is the fact that a control law based on Euler parameters may give rise to two different control values for two different vectors of Euler parameters that, in fact, represent the same *physical* attitude, that is, the same matrix in SO(3). This lack of consistency is not a problem per se but is rather an inconvenient artifact of the representation. A more serious problem arises from the fact that the desired equilibrium attitude of a rigid body may be specified in terms of a given set of Euler parameters. Therefore, if the current attitude is close to the desired attitude, but the corresponding Euler parameters are close to the antipodal set of Euler parameters in the opposite hemisphere, then the control law may rotate the body through almost 360° to reach a physical attitude that was, in fact, close to the desired orientation. This phenomenon of needless rotation through large angles is called *unwinding*.

Engineers know all this, and the solution in practice is to confine the Euler parameters to one hemisphere. But here's the rub: A switching rule is needed at the boundary of the hemispheres, and that introduces a discontinuity. Now, the closed-loop dynamics are discontinuous, and discontinuities lead to additional difficulties, such as chattering due to noise.

The feature article by Nalin Chaturvedi, Amit Sanyal, and Harris McClamroch asks the following question: What kind of closed-loop behavior can be achieved by insisting on continuous control based on rotation matrices? In short, by avoiding Euler and Gibbs parameters, they take a "retro" approach to attitude control based on SO(3).

SO(3) is difficult to visualize; in fact, an entire book has been written on visualizing the quaternions (*Visualizing Quaternions* by A.J. Hanson, Elsevier, 2006). In the more complicated case of SO(3), one approach is to begin with the closed unit solid ball in three-space. With the identity matrix located at the center of the ball, the opposite (that is, antipodal) points on the surface of the ball are "identified," which essentially means that all pairs of opposite points are glued together. I haven't seen an image of the result of this procedure; we'll just have to let this month's cover capture the mysterious world of SO(3).

What the feature article on SO(3) in this issue shows is that every continuous closed-loop system on SO(3) has at least four equilibria, three of which are unstable and one of which is asymptotically stable. This is the price paid for avoiding the possibility of unwinding without resorting to discontinuous switching. Fortunately, the additional equilibria are saddle points, which means that, in practice, the trajectory does not remain there but may move slowly in the vicinity as it passes by on the way to the stable equilibrium.

In addition to these feature articles, this issue brings you a complete suite of articles by the CSS leadership—the "President's Message," "CSS News," "Membership Activities," "Technical Committee Activities," and a new department called "Publication Activities." We also bring you an "Ask the Experts" column on automated welding, two "People in Control" interviews, lots of book announcements in "Bookshelf," two conference reports, and previews of both the 2010 MSC and the 2010 CDC/ECC. This year's CDC is, in fact, the 50th CDC—an event that you definitely do not want to miss. We also publish an obituary of Howard Rosenbrock, a giant in our field. And we end with some observations on a basic systems concept.

By the time you receive this issue, the ACC will be coming up shortly. This year we're headed to the west coast to partake in the usual midyear get-together to hear what everyone's been up to research-wise. Can't wait!

Dennis S. Bernstein