Qualitative Modeling and Heterogeneous Control of Global System Behavior*

Benjamin Kuipers[1] and Subramanian Ramamoorthy[2]

[1] Computer Science Department, University of Texas at Austin, Austin, Texas 78712 USA.
kuipers@cs.utexas.edu
[2] Electrical and Computer Engineering Department, University of Texas at Austin, Austin, Texas 78712, and National Instruments Corp., 11500 N. Mopac Expwy, Building B, Austin, Texas 78759 USA. s.ramamoorthy@ni.com

Abstract. Multiple model approaches to the control of complex dynamical systems are attractive because the local models can be simple and intuitive, and global behavior can be analyzed in terms of transitions among local operating regions. In this paper, we argue that the use of qualitative models further improves the strengths of the multiple model approach by allowing each local model to describe a large class of useful non-linear dynamical systems. In addition, reasoning with qualitative models naturally identifies weak sufficient conditions adequate to prove qualitative properties such as stability. We demonstrate our approach by building a global controller for the free pendulum. We specify and validate local controllers by matching their structures to simple generic qualitative models. This process identifies qualitative constraints on the controller designs, sufficient to guarantee the desired local properties and to determine the possible transitions between local regions. This, in turn, allows the continuous phase portrait to be abstracted to a simple transition graph. The degrees of freedom in the design that are unconstrained by the qualitative description remain available for optimization by the designer for any other purpose.

1 Introduction

Multiple model approaches to the control of complex dynamical systems are attractive because the local models can be simple and intuitive, and global behavior can be analyzed in terms of transitions among local operating regimes [1].

In this paper, we argue that the use of qualitative models further improves the strengths of the multiple model approach by allowing each local model to describe a large class of useful non-linear dynamical systems [2]. In addition, reasoning with qualitative models naturally identifies weak sufficient conditions adequate to prove qualitative properties such as stability. Since a qualitative model only constrains certain aspects of a real system, the remaining degrees of freedom are available for optimization according to any criterion the designer chooses.

* This work has taken place in the Intelligent Robotics Lab at the Artificial Intelligence Laboratory, The University of Texas at Austin. Research of the Intelligent Robotics lab is supported in part by NSF grants IRI-9504138 and CDA 9617327, and by funding from Tivoli Corporation.

We use the QSIM framework for representing qualitative differential equations (QDEs) and doing qualitative simulation to predict the set of all possible behaviors of a QDE and initial state [2]. A QDE is a qualitative abstraction of a set of ODEs, in which the domain of each variable is described in terms of a finite, totally ordered set of *landmark values*, and an unknown function may be described in terms of regions of monotonic behavior and tuples of corresponding landmark values it passes through. Qualitative simulation predicts a transition graph of qualitative states guaranteed to describe all solutions to all ODE models consistent with the given QDE. By querying QSIM output with a temporal logic model-checker, we can prove universal statements in temporal logic as theorems about sets of dynamical systems described by the QDE [3].

Because it is consistent with nonlinear models, a simple and intuitive QDE model can cover a larger region of the state space than would be possible for a linear ODE. Because a QDE model can express incomplete knowledge, it can be formulated even when the model is not fully specified, and it can express sufficient conditions for a desired guarantee while leaving other degrees of freedom unspecified. These properties are helpful in abstracting the continuous state space of the system to a compact and useful transition graph.

1.1 Abstraction from Continuous to Discrete States

The discrete transition-graph representation is important for reasoning about large-scale hybrid systems, because it allows the analyst to focus on which large-granularity state the system is in rather than on its detailed dynamics. The representation facilitates analysis of the system using temporal logic and automata theory [4], and building hierarchical representations for knowledge of dynamics [5].

We decompose the state space into a set of regions with disjoint interiors, though boundary points may be shared. To be useful, the description of the dynamical system, restricted to each region, should be significantly simpler than the description of the global system. Each region is then abstracted to a node in the transition-graph model.

A transition from one node to another represents the existence of a trajectory between the corresponding regions through their common boundary in the continuous state space. Consider the set of continuous trajectories with initial states in the region. If all of those trajectories stay within the region, then the abstracted node has no outgoing transitions. If some trajectories cross the region's boundary and pass into other regions, then the abstracted model includes transitions to each of the corresponding nodes.

QSIM predicts all possible behaviors of a system, given a QDE model and a qualitative description of its initial state. Therefore, if the region can be characterized by a qualitative description, and if the dynamical system restricted to that region can be described by a QDE, then qualitative simulation can infer the corresponding transitions.

Qualitative modeling and simulation is not a "magic bullet" for proving properties of arbitrary nonlinear and heterogeneous systems. However, it does provide a much more expressive language for describing the qualitative and semi-quantitative properties of classes of non-linear dynamical systems, and inferring properties of the sets of all possible behaviors of those systems. It provides more flexibility and power for a designer to specify intended properties of a dynamical system. It also provides tools for proving that a qualitatively specified design achieves its desired goals.

1.2 Example: the Free Pendulum

The free pendulum (Figure 1) is a simple but non-trivial non-linear dynamical system. The task of balancing the pendulum in the upright position is widely used as a textbook exercise in control, and as a target for machine learning methods that learn dynamical control laws. The inverted pendulum is also an important practical model for tasks ranging from robot walking to missile launching.

We demonstrate our approach by building a global controller for the free pendulum. We specify and validate local controllers by matching their structures to simple generic qualitative models. The qualitative framework of QSIM allows us to generalize simple familiar systems like the damped harmonic oscillator ("damped spring"), by replacing linear terms with monotonic functions. Either by using QSIM or analytically (as we do in this paper), it is not difficult to prove useful qualitative properties of the damped spring and important variants such as the spring with negative damping.

There is an open-ended set of local models that have desirable properties to be incorporated into a heterogeneous hybrid model. We explore some simple but useful examples here. The set of useful transitions among local models is also currently open-ended, but may turn out in the end to be finite, at least under qualitative description. We provide some useful examples here, but no suggestion yet about the limits of such a set.

This process identifies qualitative constraints on the controller designs, adequate to guarantee the desired local properties and to determine the possible transitions between local regions. This, in turn, allows the continuous phase portrait to be abstracted to a simple transition graph.

2 Qualitative Properties of Damped Oscillators

Before addressing the pendulum, we need to prove a couple of useful lemmas about the properties of two generic qualitative models: the spring with damping friction and the spring with negative damping.

Consider the familiar mass-spring system. The key fact about springs is Hooke's Law, which says that the restoring force exerted by a spring is proportional to its displacement from its rest position. If x represents the spring's displacement from rest, then

$$F = ma = m\ddot{x} = -k_1 x.$$

We add a damping friction force to the linear model by adding a term proportional to $\dot{x}$ and opposite in direction. (Real damping friction is often non-linear.)

$$F = ma = m\ddot{x} = -k_1 x - k_2 \dot{x}.$$

Rearranging and renaming the constants, we get a linear model of the damped spring:

$$\ddot{x} + b\dot{x} + cx = 0. \qquad (1)$$

The linear model is easy to solve, but it embodies simplifying assumptions that are often unrealistic. By generalizing linear terms in equation (1) to monotonic functions, and allowing the functions to be described qualitatively rather than specified precisely, we get a model

$$\ddot{x} + f(\dot{x}) + g(x) = 0$$

that encompasses a large number of precise ODE models, including ones that are much more realistic descriptions of the world.

To make qualitative simulation possible, we must restrict our attention to "reasonable" functions, which are defined below along with some useful concepts for expressing qualitative models.

Definition 1. *Where $[a, b] \subseteq \Re^*$, the function $f : [a, b] \rightarrow \Re^*$ is a* reasonable function *over $[a, b]$ if*

1. *f is continuous on $[a, b]$,*
2. *f is continuously differentiable on (a, b),*
3. *f has only finitely many critical points in any bounded interval,*
4. *The one-sided limits $\lim_{t \rightarrow a^+} f'(t)$ and $\lim_{t \rightarrow b^-} f'(t)$ exist in $\Re^*$. Define $f'(a)$ and $f'(b)$ to be equal to these limits.*

Definition 2. *M^+ is the set of reasonable functions $f : [a, b] \rightarrow \Re^*$ such that $f' > 0$ over (a, b).*

Definition 3. *M_0^+ is the set of $f \in M^+$ such that $f(0) = 0$.*

Definition 4. $[x]_0 = sign(x) \in \{+, 0, -\}$.

Here we establish the important qualitative properties of the monotonic "damped spring" model.

Lemma 1. *Let $A \subseteq \Re^2$ include $(0, 0)$ in its interior, and let S be a system governed by the QDE*

$$\ddot{x} + f(\dot{x}) + g(x) = 0 \qquad (2)$$

for every $(x, \dot{x}) \in A$, where f and g are reasonable functions such that $f \in M_0^+$ and $[g(x)]_0 = [x]_0$. Then for any trajectory $(x(t), \dot{x}(t))$ of S that lies entirely within A,

$$\lim_{t \rightarrow \infty} (x(t), \dot{x}(t)) = (0, 0).$$

Proof: We rewrite equation (2) as

$$\begin{aligned} \dot{x_1} &= f_1(x_1, x_2) = x_2 \\ \dot{x_2} &= f_2(x_1, x_2) = -f(x_2) - g(x_1) \end{aligned} \qquad (3)$$

Because g is a reasonable function, we know that $g'(0)$ is defined. Since $[g(x)]_0 = [x]_0$, we conclude that $g(0) = 0$ and $g'(0) > 0$. Any fixed-point of equation (3) must satisfy $\dot{x_1} = \dot{x_2} = 0$, which implies that the only fixed point is at $x_1 = x_2 = 0$.

By the stable manifold theorem [6], the qualitative behavior of the nonlinear system (3) around the fixed point at $(0,0)$ is the same as that of its local linearization:

$$\begin{aligned} \dot{x_1} &= x_2 \\ \dot{x_2} &= -f'(0)x_2 - g'(0)x_1 \end{aligned} \tag{4}$$

The eigenvalues of (4) are

$$\lambda_{1,2} = \frac{1}{2}\left[-f'(0) \pm \sqrt{f'(0)^2 - 4g'(0)}\right].$$

Because $f'(0), g'(0) > 0$, the eigenvalues have negative real parts, so $(0,0)$ is an asymptotically stable fixed point. When the "friction force" term f' is small relative to the "spring force" term g', the eigenvalues will be complex, in which case $(0,0)$ will be a spiral attractor.

Because $[g(x)]_0 = [x]_0$, the "spring force" is always a restoring force, so we can define a Lyapunov function

$$V(x,\dot{x}) = \frac{1}{2}\dot{x}^2 + \int_0^x g(x)\ dx \tag{5}$$

and show that $V(x,\dot{x}) \geq 0$, that $V = 0$ only at $(0,0)$, and that $\frac{d}{dt}V \leq 0$. This means that S is asymptotically stable at $(0,0)$, and that A can contain no limit cycles.

Together, this tells us that any trajectory $(x(t),\dot{x}(t))$ that enters A eventually terminates at $(0,0)$, for any reasonable functions f and g such that $f \in M_0^+$ and $[g(x)]_0 = [x]_0$. □

Now we establish similar properties for another monotonic generalization of the "damped spring", but with negative damping.

Lemma 2. *Let $A \subseteq \Re^2$ include $(0,0)$ in its interior, and let S be a system governed by the QDE*

$$\ddot{x} - f(\dot{x}) + g(x) = 0 \tag{6}$$

for every $(x,\dot{x}) \in A$, where f and g are reasonable functions such that $f \in M_0^+$ and $[g(x)]_0 = [x]_0$. Then $(0,0)$ is the only fixed point of S in A, and it is unstable. Furthermore, A cannot contain a limit cycle.

Proof: The proof of this Lemma is very similar to the previous one. We rewrite equation (6) as

$$\begin{aligned} \dot{x_1} &= f_1(x_1,x_2) = x_2 \\ \dot{x_2} &= f_2(x_1,x_2) = f(x_2) - g(x_1) \end{aligned} \tag{7}$$

As before, because g is a reasonable function, we know that $g'(0)$ is defined. Since $[g(x)]_0 = [x]_0$, we conclude that $g(0) = 0$ and $g'(0) > 0$.

Any fixed-point of equation (7) must satisfy $\dot{x_1} = \dot{x_2} = 0$, so the only fixed point is at $x_1 = x_2 = 0$.

As before, the qualitative behavior of the nonlinear system (7) around the fixed point at $(0,0)$ is the same as that of its local linearization:

$$\begin{aligned} \dot{x_1} &= x_2 \\ \dot{x_2} &= f'(0)x_2 - g'(0)x_1 \end{aligned} \tag{8}$$

The eigenvalues of (8) are

$$\lambda_{1,2} = \frac{1}{2}\left[f'(0) \pm \sqrt{f'(0)^2 - 4g'(0)}\right].$$

In this case, since $f'(0) > 0$, the eigenvalues have positive real parts, and $(0,0)$ is an unstable fixed point. If the “friction force” term f' is small relative to the “spring force” term g', then the eigenvalues will be complex, so $(0,0)$ will be a spiral repellor.

By the Bendixon negative criterion [6], there can be no periodic orbits contained in A because

$$\frac{\partial f_1}{\partial x_1} + \frac{\partial f_2}{\partial x_2} = f'(x_2)$$

is always positive over A. That is, A cannot contain a limit cycle.

Therefore, except for the unstable fixed-point at $(0,0)$ itself, any trajectory $(x(t), \dot{x}(t))$ that starts in A, eventually leaves A, for any reasonable functions f and g such that $f \in M_0^+$ and $[g(x)]_0 = [x]_0$. □

In this context, like that of Lemma 1, we can interpret $V(x, \dot{x})$ from equation (5) as representing the total energy of the system, but here we can show that energy is increasing steadily except at isolated points.

2.1 Proof by Qualitative Simulation

These Lemmas establishing the properties of the monotonic spring models were proved by hand to make this paper self-contained, and because these models are generic and useful.

It is also possible to generate proofs of these and similar statements automatically from the QSIM QDE models. The Guaranteed Coverage Theorem states that every real behavior of every model described by the QDE is predicted by QSIM [7, 2]. Then we can use a temporal logic model-checker to establish whether the predicted behavior tree is a model of a specified statement in temporal logic. For universal statements, the completeness of the model-checker and QSIM Guaranteed Coverage combine to show that a positive response from the model-checker implies that the temporal logic statement is a theorem for all behaviors of all dynamical systems consistent with the given QDE [3].

This method of deriving the necessary lemmas using QSIM makes it possible to generalize this approach to more complex models as in [8].

3 A Controller for the Pendulum

By appealing to the qualitative properties of solutions to these very general models, we can give a simple and natural derivation for a controller for the pendulum, able to pump it up and stabilize it in the inverted position.

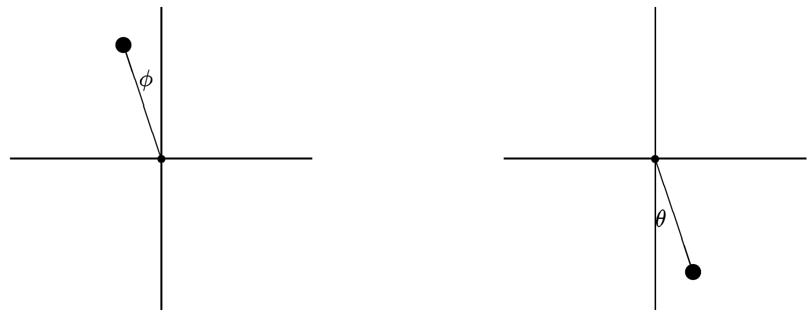

Fig. 1. Local models of the pendulum (where $\phi = \theta - \pi$): (a) $\phi = 0$ at the unstable fixed-point, and (b) $\theta = 0$ at the stable fixed-point.

3.1 Stabilizing the Inverted Pendulum

The pendulum is a mass on a rigid, massless rod, attached to a fixed pivot. The variable ϕ measures the angular position counter-clockwise from the vertical (Figure 1(a)). We consider only $\phi \in (-\pi/2, +\pi/2)$.[1]

The angular acceleration due to gravity is $k \sin \phi$, and there is a small amount of damping friction $-f(\dot{\phi})$, where $f \in M_0^+$. A control action $u(\phi, \dot{\phi})$ exerts angular acceleration at the pivot. The resulting model of the pendulum is:

$$\ddot{\phi} + f(\dot{\phi}) - k \sin \phi + u(\phi, \dot{\phi}) = 0. \tag{9}$$

Our goal is to design $u(\phi, \dot{\phi})$ so that the system is asymptotically stable at $(\phi, \dot{\phi}) = (0, 0)$.

Lemma 1 provides a simple sufficient condition: make the pendulum behave like a monotonic damped spring. We define the controller for the **Balance** region to be:

$$u(\phi, \dot{\phi}) = g(\phi) \text{ such that } [g(\phi) - k \sin \phi]_0 = [\phi]_0. \tag{10}$$

Since $k \sin \phi$ increases monotonically with ϕ over $(-\pi/2, +\pi/2)$, $g(\phi)$ must increase at least as fast in order to ensure that $[g(\phi) - k \sin \phi]_0 = [\phi]_0$.

We can get faster convergence by augmenting the natural damping $f(\dot{\phi})$ with a damping term $h(\dot{\phi})$ included in the control law, giving us

$$u(\phi, \dot{\phi}) = g(\phi) + h(\dot{\phi}) \text{ where } [g(\phi) - k \sin \phi]_0 = [\phi]_0 \text{ and } h \in M_0^+(\dot{\phi}). \tag{11}$$

If there is a bound u_{max} on the control action u, then the limiting angle ϕ_{max} beyond which the controller cannot restore the pendulum to $\phi = 0$ is given by the constraint

$$u_{max} = k \sin \phi_{max}. \tag{12}$$

[1] The derivation here applies over the larger interval $(-\pi, +\pi)$, but the maximum control force is required at $\phi = \pm\pi/2$. The controller design problem is less interesting if the controller is powerful enough to lift the pendulum directly to $\phi = 0$ from any value of ϕ.

The maximum velocity $\dot{\phi}_{max}$ that the **Balance** controller can tolerate at $\phi = 0$ is then determined by the constraint

$$\frac{1}{2}\dot{\phi}_{max}^2 = \int_0^{\phi_{max}} g(\phi) - k \sin\phi \, d\phi \tag{13}$$

which represents the conversion of the kinetic energy of the system (9) at $(0, \dot{\phi}_{max})$ into potential energy at $(\phi_{max}, 0)$.

Therefore, we define the region of applicability for the **Balance** controller by the ellipse

$$\frac{\phi^2}{\phi_{max}^2} + \frac{\dot{\phi}^2}{\dot{\phi}_{max}^2} \leq 1. \tag{14}$$

Note that the shapes of the non-linear functions g and h are only very weakly constrained. The qualitative constraints in (11) provide weak sufficient conditions guaranteeing the stability of the inverted pendulum controller. However, there is plenty of freedom available to the designer to select the properties of g and h to optimize any desired criterion.

3.2 Pumping Up the Hanging Pendulum

With no input, the stable state of the pendulum is hanging straight down. We use the variable θ to measure the angular position counter-clockwise from straight down (Figure 1(b)). The goal is to pump energy into the pendulum, swinging it progressively higher, until it reaches the region where the inverted pendulum controller can balance it in the upright position.

Angular acceleration due to gravity is $-k \sin\theta$. As before, damping friction is $-f(\dot{\theta})$, where $f \in M_0^+$, and the control action exerts an angular acceleration $u(\theta, \dot{\theta})$ at the pivot. The resulting model of our system is:

$$\ddot{\theta} + f(\dot{\theta}) + k \sin\theta + u(\theta, \dot{\theta}) = 0. \tag{15}$$

Without control action, since $[\sin\theta]_0 = [\theta]_0$ over $-\pi < \theta < \pi$, the model exactly matches the monotonic damped spring model of Lemma 1, so we know that it is asymptotically stable at $(\theta, \dot{\theta}) = (0, 0)$. Unfortunately, this is not where we want it.

Fortunately, Lemma 2 gives us a sufficient condition to transform the stable attractor at $(0, 0)$ into an unstable repellor. We define the controller for the **Pump** region so that the system is modeled by a spring with negative damping, pumping energy into the system. That is, define

$$u(\theta, \dot{\theta}) = -h(\dot{\theta}) \text{ such that } h - f \in M_0^+ \tag{16}$$

Starting with any perturbation from $(0, 0)$, this controller will pump the pendulum to higher and higher swings. Lemma 2 is sufficient to assure us that there are no limit cycles in the region $-\pi < \theta < \pi$ to prevent the trajectory from approaching $\theta = \pi$ so the **Balance** control law can stabilize it in the inverted position.

3.3 The Spinning Pendulum

The **Spin** region represents the behavioral mode of the pendulum when it is spinning freely at high speed. In the **Spin** region, a simple qualitative controller augments the natural friction of the system with additional damping, to slow the system down toward the two other regions.

$$u(\theta, \dot{\theta}) = f_2(\dot{\theta}) \text{ such that } f_2 \in M_0^+. \tag{17}$$

3.4 Bounding the Pump and Spin Regions

One might ask whether the **Pump** controller could be so aggressive that the pendulum would overshoot the **Balance** region entirely. Even with augmented damping by the **Spin** controller, it might be possible to get a limit cycle that alternates between the **Pump** and **Spin** regions. (While the analogy is not perfect, this is one aspect of how the van der Pol oscillator works.)

We can avoid this problem by defining a suitable boundary between the **Pump** and **Spin** regions, and showing that the **Pump** and **Spin** controllers together define a sliding mode controller [9], forcing nearby trajectories to converge to the boundary.

A boundary with the desired properties is the separatrix of the same pendulum,

$$\ddot{\theta} + k \sin \theta = 0 \tag{18}$$

without damping friction or control action. It turns out that this boundary will lead straight into the heart of the **Balance** region (Figure 2).

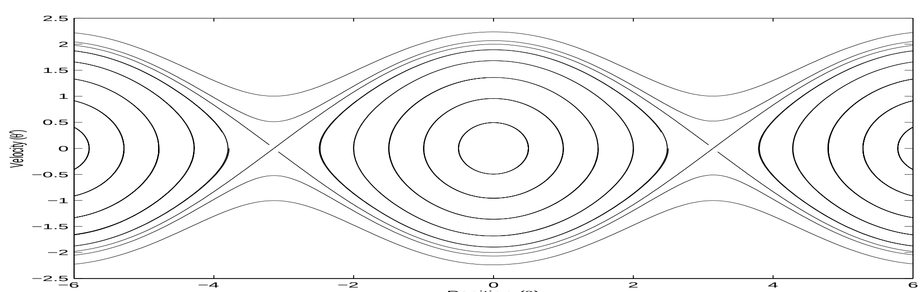

Fig. 2. The $(\theta, \dot{\theta})$ phase portrait of the undamped pendulum. The **Balance** region is an ellipse around the saddle point at $(\pi, 0)$. The **Pump** region is contained within the separatrices, and the **Spin** region is outside the separatrices.

A *separatrix* is a trajectory that starts at an unstable fixed-point of the system and ends at another fixed-point. In the case of the pendulum, the separatrices are the trajectories where the pendulum starts upright and at rest, then swings around once and returns to the upright position, at rest. It is the locus of points $(\theta, \dot{\theta})$ such that the total energy of the system is exactly equal to the potential energy of the motionless pendulum in the upright position.

$$KE + PE = \frac{1}{2}\dot{\theta}^2 + \int_0^{\theta} k \sin\theta \; d\theta = 2k$$

Evaluating the integral and simplifying, we get an equation $s(\theta, \dot{\theta}) = 0$ that defines the separatrix, i.e., the boundary between **Spin** ($s > 0$) and **Pump** ($s < 0$).

$$s(\theta, \dot{\theta}) = \frac{1}{2}\dot{\theta}^2 - k(1 + \cos\theta) = 0. \tag{19}$$

We use the method for defining a sliding mode controller from [9] to ensure that trajectories always approach $s = 0$.

Differentiating (19) and substituting for $\ddot{\theta}$, we get:

$$\begin{aligned}
\dot{s} &= \dot{\theta}\;\ddot{\theta} + k \sin\theta\;\dot{\theta} \\
&= \dot{\theta}\;(-f(\dot{\theta}) - k\sin\theta - u(\theta, \dot{\theta})) + k\sin\theta\;\dot{\theta} \\
&= -\dot{\theta} f(\dot{\theta}) - \dot{\theta}\; u(\theta, \dot{\theta})
\end{aligned}$$

Now, examine the **Pump** region, inside the separatrix where $s < 0$, and substitute the **Pump** control law (16) for $u(\theta, \dot{\theta})$.

$$\begin{aligned}
\dot{s}_{pump} &= -\dot{\theta} f(\dot{\theta}) + \dot{\theta}\; h(\dot{\theta}) \text{ where } h - f \in M_0^+ \\
&= \dot{\theta}\;(h - f)(\dot{\theta}) \\
&\geq 0
\end{aligned}$$

Similarly, for the **Spin** region where $s > 0$, substituting its control law (17).

$$\begin{aligned}
\dot{s}_{spin} &= -\dot{\theta} f(\dot{\theta}) - \dot{\theta} f_2(\dot{\theta}) \text{ where } f_2 \in M_0^+ \\
&= -\dot{\theta}(f + f_2)(\dot{\theta}) \\
&\leq 0
\end{aligned}$$

This shows that the **Pump** control law moves the system toward the separatrix from the inside, and the **Spin** control law approaches the separatrix from the outside: the existing control laws define a sliding mode controller with the separatrix $s = 0$ as the attractor. Once the system gets sufficiently close to the boundary, it will follow the separatrix, directly into the **Balance** region. In particular, it is impossible for an aggressive **Pump** controller to overshoot the **Balance** region.

3.5 Heterogeneous Control of the Free Pendulum

We have derived local control laws for the three relevant regions. The region definition for **Balance** takes priority over the defining relations for **Pump** or **Spin**.

- **Balance**: $(\phi, \dot{\phi}) \approx (0, 0)$, more precisely $\phi^2/\phi_{max}^2 + \dot{\phi}^2/\dot{\phi}_{max}^2 \le 1$ from equation (14). Stabilize the unstable saddle by adding a "spring-like" attractive force:

$$u(\phi, \dot{\phi}) = g(\phi) + h(\dot{\phi}) \text{ such that } [g(\phi) - k \sin\phi]_0 = [\phi]_0 \text{ and } h \in M_0^+(\dot{\phi}).$$

- **Pump**: $s(\theta, \dot{\theta}) < 0$, where s is defined in equation (19). Pump the system away from the stable attractor at $(0, 0)$ by adding to the controller a destabilizing "anti-frictional" force:

$$u(\theta, \dot{\theta}) = -h(\dot{\theta}) \text{ such that } h - f \in M_0^+.$$

- **Spin**: $s(\theta, \dot{\theta}) > 0$, where s is defined in equation (19). Slow down a quickly spinning pendulum by augmenting the (small) natural friction of the system with a "friction-like" damping control:

$$u(\theta, \dot{\theta}) = f_2(\dot{\theta}) \text{ such that } f_2 \in M_0^+.$$

We have shown that the qualitative constraints associated with each local law are sufficient to guarantee that its local performance is as desired. We need to demonstrate that the continuous behavior of the controlled pendulum can be abstracted to the discrete transition model consisting of the operating regions of the controller.

$$\begin{array}{ccccc} \textbf{Pump} & & \not\leftrightarrow & & \textbf{Spin} \\ & \searrow & & \swarrow & \\ & & \textbf{Balance} & & \end{array}$$

- **Pump** $\not\leftrightarrow$ **Spin**. Since the boundary $s = 0$ between **Pump** and **Spin** is the attractor for a sliding mode controller, in theory no trajectory can cross from one side of the boundary to the other. In practice, the trajectory will "chatter" around the boundary. The boundary can be made fuzzy to eliminate discontinuous changes in control action, but in any case, the trajectory can be kept very close to the boundary [9].
- **Pump** $\rightarrow$ **Balance**. The discussion in section 3.4 shows that $s(\theta, \dot{\theta})$ increases throughout **Pump**. Therefore, the maximum amplitude θ_{max} of the pendulum's swings, where $\dot{\theta} = 0$, must increase. Since these values are determined by $s = -k(1 + \cos\theta_{max})$, the value of θ_{max} must increase in absolute value, toward $\pm\pi$. Lemma 2 says that **Pump** contains no fixed point or limit cycle. Therefore, eventually the extremal point $(\theta_{max}, 0)$ will lie within the region of applicability of the **Balance** controller (14), which will capture the trajectory, bringing it to the fixed point at $(\theta, \dot{\theta}) = (\pi, 0)$ (i.e., $(\phi, \dot{\phi}) = (0, 0)$).
- **Spin** $\rightarrow$ **Balance**. Similarly, we have shown that $s(\theta, \dot{\theta})$ decreases throughout **Spin**. Therefore, the minimum velocity $\dot{\theta}_{min}$, which occurs where $\theta = \pi$ (i.e., $\phi = 0$), must also decrease in absolute value. The extremal point $(\pi, \dot{\theta}_{min})$ will eventually fall within the region of applicability of the **Balance** controller (14), which will capture the trajectory and bring it to the desired fixed-point.
- **Balance**. Lemma 1 guarantees that, once the system's trajectory enters **Balance**, it cannot leave. Therefore, there are no outgoing transitions from the **Balance** region.

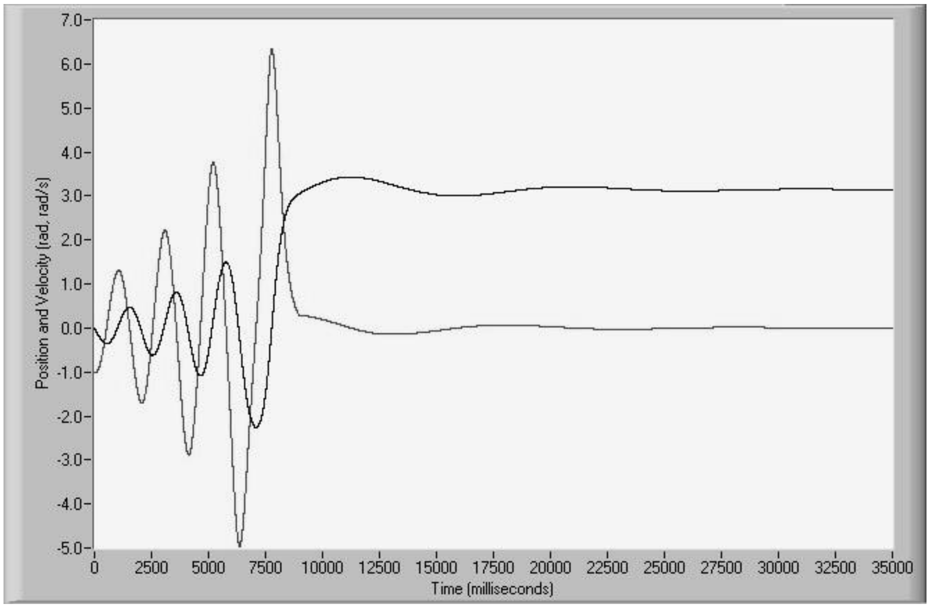

Fig. 3. $\theta(t)$ and $\dot{\theta}(t)$ as the heterogeneous controller pumps a weakly-powered pendulum from $\theta = 0$ to $\theta = \pi$.

Figure 3 shows an example behavior as a very weak controller pumps the pendulum up from $\theta = 0$ and balances it at $\phi = 0$. We define an instance of the pendulum model and the local control laws:

$$
\begin{array}{lll}
\textbf{Plant}: & \ddot{\theta} + c\dot{\theta} + k\sin\theta + u(\theta, \dot{\theta}) = 0 & c = 0.01 \quad k = 10 \quad u_{max} = 4 \\
\textbf{Balance}: & u = (c_{11} + k)(\theta - \pi) + c_{12}\dot{\theta} & c_{11} = 0.4 \; c_{12} = 0.3 \\
\textbf{Spin}: & u = c_2\dot{\theta} & c_2 = 0.5 \\
\textbf{Pump}: & u = -(c + c_3)\dot{\theta} & c_3 = 0.5
\end{array}
$$

The plant model is chosen with normal gravity, slight friction, and a maximum control action too weak to lift the pendulum directly up. The local control laws are all linear for simplicity, though they could be designed to be nonlinear. The controllers are defined so that the desired behavior is guaranteed as long as the parameters c_i are all positive. The specific values for the c_i are chosen to ensure that $u < u_{max}$.

Given the maximum control action u_{max} and the gain $c_{11} + k$ of the **Balance** controller, we can determine the bounds $\phi_{max} = 0.4$ and $\dot{\phi}_{max} = 0.3$ for the **Balance** region from equations (12) and (13), respectively. We define the switching strategy to be

If $\alpha \leq 1$ then **Balance**
else if $s < 0$ then **Pump**
else **Spin**

where

$$\alpha = \frac{\phi^2}{\phi^2_{max}} + \frac{\dot{\phi}^2}{\dot{\phi}^2_{max}} \text{ and } s = \frac{1}{2}\dot{\theta}^2 - k(1 + \cos\theta).$$

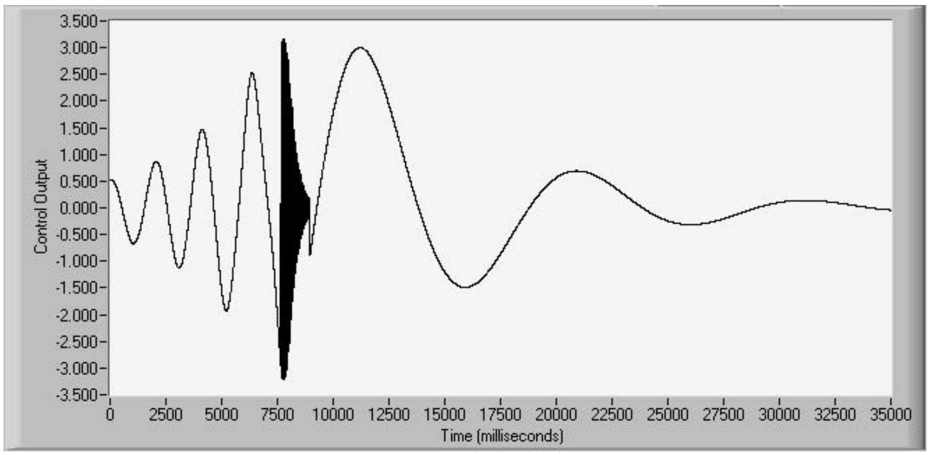

Fig. 4. The control action $u(t)$ shows chattering along the sliding mode.

Because of the sharp transitions among regions, the control action $u(t)$ "chatters" back and forth across the **Spin-Balance** interface (Figure 4). A "dead zone" along the boundary where $u = 0$ produces a virtually identical behavior, but without chatter in the control action. Fuzzy boundaries would presumably have the same effect.

4 Discussion

4.1 Regions with Fuzzy Boundaries

In some cases, it is convenient to have local models and corresponding regions that overlap, or have gradual rather than sharp boundaries. Such regions can be described by fuzzy set membership functions. In the simplest case, the continuous state space can be decomposed into *pure* regions, where only one membership function is non-zero, and *overlap* regions, where two (or perhaps a small finite number of) regions have non-zero membership functions. The dynamical system in an overlap region is the weighted average of the overlapping local models, weighted by the values of the membership functions.

The qualitative QDE formalism is particularly useful for representing overlap regions, since not only the local models, but even more, the shapes of the membership functions in the overlap region may be only partially known or specified. QSIM can establish which properties of the local models, and of the overlapping membership functions, are sufficient to guarantee that trajectories through an overlap region can be abstracted to a transition from one pure region to another. Kuipers and Åström [8] demonstrated this for controllers for a simple water tank and for a highly nonlinear chemical reaction.

4.2 Feedback Linearization

Feedback linearization [10] designs a control law for a system to add a term compensating for the non-linearities in the system, making the sum linear and therefore suitable for well-understood control methods. The problem is that this approach demands precise knowledge of the nonlinear system.

In our qualitative method, we make the much weaker requirement that the sum of the nonlinear system and the controller be monotonic. This may be achievable even with incomplete knowledge of the original system, for example with bounding envelopes around unknown functions. Incomplete knowledge in this form will reduce the remaining degrees of freedom available for optimization, but it will not affect the qualitative guarantee of stability.

4.3 Conclusions

By using qualitative models, we make it possible to express incomplete knowledge of the dynamics of the uncontrolled plant, and to separate the properties of the controller needed to provide qualitative guarantees from the remaining degrees of freedom that can be used for optimization. Qualitative models can also express natural nonlinear models, allowing the use of larger and more natural local models in a multiple-model framework. Furthermore, QSIM can be used to prove the necessary properties of generic qualitative models, or of the specific models that describe the controlled system. These features are illustrated by the design of a heterogeneous controller for the free pendulum.

References

1. Murray-Smith, R., Johansen, T.A.: Multiple Model Approaches to Modelling and Control. Taylor and Francis, UK (1997)
2. Kuipers, B.J.: Qualitative Reasoning: Modeling and Simulation with Incomplete Knowledge. MIT Press, Cambridge, MA (1994)
3. Shults, B., Kuipers, B.: Proving properties of continuous systems: qualitative simulation and temporal logic. Artificial Intelligence **92** (1997) 91–129
4. Alur, R., Courcoubetis, C., Halbwachs, N., Henzinger, T.A., Ho, P.H., Nicollin, X., Olivero, A., Sifakis, J., Yovine, S.: The algorithmic analysis of hybrid systems. Theoretical Computer Science **138** (1995) 3–34
5. Varaiya, P.: A question about hierarchical systems. In Djaferis, T., Schick, I., eds.: System Theory: Modeling, Analysis and Control. Kluwer (2000)
6. Guckenheimer, J., Holmes, P.: Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields. Springer-Verlag, Berlin (1983)
7. Kuipers, B.: Qualitative simulation. Artificial Intelligence **29** (1986) 289–338
8. Kuipers, B.J., Åström, K.: The composition and validation of heterogeneous control laws. Automatica **30** (1994) 233–249
9. Slotine, J.J., Li, W.: Applied Nonlinear Control. Prentice Hall, Englewood Cliffs NJ (1991)
10. Friedland, B.: Advanced Control System Design. Prentice-Hall (1996)