Qualitative Heterogeneous Control of Higher Order Systems*

Subramanian Ramamoorthy[1] and Benjamin Kuipers[2]

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

Abstract. This paper presents the qualitative heterogeneous control framework, a methodology for the design of a controlled hybrid system based on attractors and transitions between them. This framework designs a robust controller that can accommodate bounded amounts of parametric and structural uncertainty. This framework provides a number of advantages over other similar techniques. The local models used in the design process are qualitative, allowing the use of partial knowledge about system structure, and nonlinear, allowing regions and transitions to be defined in terms of dynamical attractors. In addition, we define boundaries between local models in a natural manner, appealing to intrinsic properties of the system. We demonstrate the use of this framework by designing a novel control algorithm for the cart-pole system. In addition, we illustrate how traditional algorithms, such as linear quadratic regulators, can be incorporated within this framework. The design is validated by experiments with a physical system.

1 Introduction

Multiple model approaches to control are useful for complex dynamical systems because the local models can be simple and intuitive, and because global behavior can be concisely described as a finite graph of transitions among models. Hybrid systems are often constructed with the local models being linear and their operating regions being polygonal. Qualitative models add the ability to express incomplete knowledge of the dynamical system and of the controller, describing a family of controllers and systems and predicting the behaviors. A qualitative model can often give a completely accurate (though imprecise) description of a nonlinear system over a larger and more naturally defined local region than can be usefully approximated by a local linear model.

Qualitative heterogeneous control (QHC) is an approach to designing controllers for complex nonlinear systems [1, 2]. It works by defining a hybrid system consisting of a set of qualitatively described control laws, each with its own operating region. The local controllers and their operating regions are designed so that any fully specified system

* 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.

and controller satisfying the given set of qualitative constraints is guaranteed to exhibit the desired qualitative behavior within each local region and at its boundaries. The local behaviors are designed to abstract to a global transition graph with the desired global behavior. The qualitative constraints are a set of weak sufficient conditions that guarantee the desired global behavior. The remaining degrees of freedom on the way to a concrete design are available to the designer for optimization according to any desired criterion, since the global qualitative behavior is already guaranteed. QHC therefore provides a separation of concerns between qualitative correctness and quantitative optimization.

In recent work [2], we demonstrated the design of a controller for pumping up and balancing a free pendulum, controlling the torque applied at the pivot. The control laws transform the natural dynamics of the pendulum to match different instances of the same generic behavior model: the qualitative damped oscillator, $\ddot{x} + f(\dot{x}) + g(x) = 0$, with positive or negative damping. In the current paper, we extend QHC to a more complex, but still very familiar system. The cart-pole version of the pendulum (Figure 2) has the same goal of pumping the pendulum up and keeping it balanced, but now we must bring the cart to the center of its track and keep it away from the endpoints, and we want it to recover gracefully from large disturbances. We demonstrate a method using time-scale abstraction to decompose the fourth-order cart-pole system into two weakly interacting second-order systems: the pole system that can be controlled by a modified version of the pivot-torque controller, and the cart system that can be controlled in a similar way. The multiple-model structure allows us to handle the interactions between systems effectively at the different model boundaries. We also demonstrate the use of traditional linear controller design methods such as LQR for a local controller for the *Balance* region.

Furthermore, we implement our QHC control law on a physical implementation of the cart-pole system and show that our framework accommodates simple solutions to aspects of the physical dynamics such as static friction that are often omitted or handled in ad hoc ways. The solution to this problem demonstrates the use of a behavior model other than the damped oscillator, the Lienard equation, in order to produce a limit cycle with desired properties.

More generally, QHC exploits the robustness of structurally stable orbits in dynamical systems. From dynamical systems theory, we know that if there exists a connected phase space volume that maps into itself under the forward evolution then the flow is globally contracting onto an attractor [3, 4]. In the simplest cases, these attractors will be fixed points. For our purposes, we are interested in implementing controllers that enforce residence of the flow within a finite subset of phase space. In simple systems, fixed points and limit cycles possess this property. Complex systems that exhibit chaotic behavior possess strange attractors that satisfy the same property of residence of the flow in a finite volume in phase space. So we consider stable fixed-points, limit cycles, and chaotic attractors all to be examples of controlled flows. This general notion of viewing the stability question in control design as one of defining an appropriate contraction property or residence of flows in a finite volume has been explored in recent literature, e.g., [4]. The QHC framework makes it possible to utilize stable attractors, periodic orbits, strange attractors and even divergent flows to synthesize global trajectories from appropriately defined local regions and transitions between them. An approach to the composition of global behaviors from such local dynamical models and orbits is seen in [5, 6]. The advantage that QHC brings over this existing work is derived from the fact that local models defined in terms of QDEs enable the design of controllers for a *class* of systems, with the guarantee that any numerical instance of the specified QDEs will possess the specified dynamical property.

Similar approaches have been explored in the context of controlling chaotic systems. The idea of chaos control originated with [7], popularly referred to as the OGY technique. The technique consists of waiting for a natural passage of the chaotic orbit close to a desired periodic orbit, and then applying a small perturbation designed to stabilize the periodic dynamics, with flexibility in switching from one behavior to another. [8] contains an extensive review of recent work in this area. These methods have largely focused on algorithms that allow a particular orbit or dynamical state to be achieved due to small control actions. [9] describes the Perfect Moment algorithm that is similar in spirit to the QHC framework. This algorithm autonomously explores and maps the phase space of a chaotic system to identify useful dynamical orbits. Then, utilizing the property of sensitive dependence on initial conditions, it synthesizes a trajectory based on these identified dynamical orbits to achieve a desired dynamical state starting from a specified initial condition. The advantage provided by QHC over these approaches is derived from the use of a variety of *qualitatively* defined dynamical orbits, and a structured synthesis technique with guarantees on behaviors and transitions that apply to classes of systems sharing the specified property.

2 A Qualitative Behavior Model: The Damped Oscillator

To design a local control law, we select a well-understood qualitative behavior model, and define the control law u so that the natural behavior of the system is transformed into that of the model. A model with attractive properties is the linear damped oscillator, $\ddot{x}+a\dot{x}+bx=0$. It is straight-forward to generalize this to a nonlinear model $\ddot{x}+f(\dot{x})+g(x)=0$ where f and g are monotonic functions. As it happens, we can generalize f and g even further to the sign-equality constraints that appear in Lemmas 1 and 2. Lemma 1 tells us that any damped oscillator matching these requirements converges to a stable fixed-point. Lemma 2 tells us that the same oscillator, but with negative damping, necessarily diverges. We will use these models in several different ways.

Definition 1. *Where $[a,b]\subseteq\Re^*$, the function $f:[a,b]\rightarrow\Re^*$is a reasonable function over $[a,b]$ if f is continuous on $[a,b]$, continuously differentiable on (a,b), has only finitely many critical points in any bounded interval, and the one-sided limits $\lim_{t\rightarrow a^+} f'(t)$ and $\lim_{t\rightarrow b^-} f'(t)$ exist in $\Re^*$. $f'(a)$ and $f'(b)$ are defined to be equal to these limits.*

Notation. M^+ is the set of reasonable functions $f:[a,b]\rightarrow\Re^*$ such that $f'>0$ over (a,b). M_0^+ is the set of $f\in M^+$ such that $f(0)=0$. $[x]_0=sign(x)\in\{+,0,-\}$.

Lemma 1. *Let $A\subseteq\Re^2$ include (0,0) in its interior, and let S be a system governed by the QDE $\ddot{x}+f(\dot{x})+g(x)=0$ for every $(x,\dot{x})\in A$, where f and g are reasonable functions such that $[f(\dot{x})]_0=[\dot{x}]_0$ and $[g(x)]_0=[x]_0$. Then for any trajectory $(x(t),\dot{x}(t))$ of S that lies entirely within A, $\lim_{t\rightarrow\infty}(x(t),\dot{x}(t))=(0,0)$*

Lemma 2. *Let $A \subseteq \Re^2$ include (0,0) in its interior, and let S be a system governed by the QDE $\ddot{x} - f(\dot{x}) + g(x) = 0$ for every $(x, \dot{x}) \in A$, where f and g are reasonable functions such that $[f(\dot{x})]_0 = [\dot{x}]_0$ 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.*

3 The Pivot-Torque Pendulum Controller

In recent work [2], we demonstrated the design of a controller for pumping up and balancing a free pendulum, controlling the torque applied at the pivot: $\ddot{\theta} + f(\dot{\theta}) + k\cos\theta - u(\theta, \dot{\theta}) = 0$. The resulting QHC controller has three regions: *Pump* (to raise the pendulum upward from its downward position), *Spin* (to slow down a rapidly-spinning pendulum), and *Balance* (to maintain the pendulum in its upward position). All three local controllers were designed by specifying control laws u that would transform the natural dynamics of the pendulum to match different instances of the same generic behavior model: the qualitative damped oscillator, $\ddot{x} + f(\dot{x}) + g(x) = 0$. (In the case of *Pump*, the damping is negative.) The boundary of the *Balance* region is determined by the maximum torque that can be applied, and the boundary separating the *Pump* and *Spin* regions is defined to be a sliding mode controller leading directly into the *Balance* region.

The pivot-torque pendulum controller can be summarized by the following equations. To be able to write each local control law with its fixed point at (0,0), we use $\theta = 0$ to refer to the fixed-point with the pendulum pointing downward, and $\phi = 0 = \theta + \pi$ to refer to the fixed-point with the pendulum pointing upward.

Given a model,

$$\ddot{\phi} + f(\dot{\phi}) - k\sin\phi + u(\phi, \dot{\phi}) = 0 \tag{1}$$

or equivalently, $\ddot{\theta} + f(\dot{\theta}) + k\sin\theta - u(\theta, \dot{\theta}) = 0$, we apply one of the following control laws,

$$\begin{aligned}
\textit{Balance}: u(\phi, \dot{\phi}) &= g(\phi) + h(\dot{\phi}), \text{ such that } [g(\phi) - k\sin\phi]_0 = [\phi]_0\,,\ [h(\dot{\phi})]_0 = [\dot{\phi}]_0 \\
\textit{Pump}: u(\phi, \dot{\phi}) &= -h(\dot{\phi}), \text{ such that } [(h - f)(\dot{\phi})]_0 = [\dot{\phi}]_0 \\
\textit{Spin}: u(\phi, \dot{\phi}) &= f_d(\dot{\phi}), \text{ such that } [f_d(\dot{\phi})]_0 = [\dot{\phi}]_0
\end{aligned} \tag{2}$$

The selection of the control region depends on the values of two parameters:

$$\alpha = \frac{\dot{\phi}^2}{\dot{\phi}^2_{\max}} + \frac{\int_0^{\phi} g(\phi) - k\sin\phi\, d\phi}{\int_0^{\phi_{\max}} g(\phi) - k\sin\phi\, d\phi} \text{ and } s(\phi, \dot{\phi}) = \tfrac{1}{2}\dot{\phi}^2 - k(1 - \cos\phi). \tag{3}$$

$\alpha \le 1$ describes the region of applicability of the *Balance* controller based on physical limitations ($\phi_{\max}$ and $\dot{\phi}_{\max}$) and the requirement that the system should not exit the *Balance* region due to control actions, once it has entered it. $s(\phi, \dot{\phi})$ represents the energy of the system, with the *separatrix* of the pendulum (the locus of points where $s(\phi, \dot{\phi}) = 0$) serving as the boundary between the *Pump* and *Spin* regions. The rule for selecting control mode is thus:

$$\begin{array}{r} \text{if } \alpha \le 1 \; then \; \textit{Balance} \\ \text{else if } s < 0 \; then \; \textit{Pump} \\ else \; \textit{Spin} \end{array} \tag{4}$$

The operation of the heterogeneous pendulum controller can be summarized in a discrete transition graph, as shown in Figure 1.

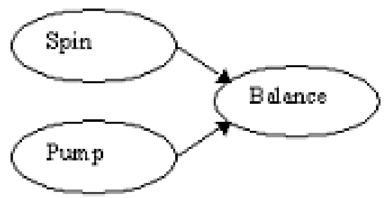

Fig. 1. Transition graph structure of the heterogeneous controller for the pivot-torque pendulum

4 The Cart-Pole System

The cart-pole system is a common benchmark problem in the control systems literature. Early work in linear control and stabilization of unstable systems focused on the basic stabilization problem for this system. The system is still being used in the current literature, e.g., [10] illustrates the idea of energy based control of pendulum swing-up, [11] presents a hybrid control algorithm that globally stabilizes a cart-pole system and [12] presents an approach to the control of a periodic orbit in a nonlinear system using the cart-pole as an example. In a different context, the inverted pendulum model has been used as an abstraction for many physically meaningful phenomena. A hypothesis in the biomechanics community is that a model known as the Spring Loaded Inverted Pendulum is the control target for the musculoskeletal system. This hypothesis has been explored and experimentally supported in [13, 14]. Pratt et. al. [15] have implemented successful walking robots based on this principle and suggest that an intuitive control algorithm designed from task specifications would be of value to many communities, such as robotics and biomechanics.

The cart-pole system considered in this paper is seen in Figure 2. It consists of a cart that moves on a horizontal track of finite length. The pole is represented by a point mass attached to the end of a massless thin rod of length l that is attached to the cart at a pivot capable of unconstrained (360^o) rotation. The primary control objective is to stabilize the system at $[x, \dot{x}, \phi, \dot{\phi}] = [0, 0, 0, 0]$, starting from $[0, 0, \pi, 0]$.

The cart-pole system is a commonly seen demonstration in many control laboratories. While this system has been stabilized by a wide variety of control algorithms, most of them suffer from a number of failure modes when people interact with these systems. For instance, by hitting the pole with a large velocity, one may cause the control action to become large and the cart may hit the end of the track in an attempt to regain control. We recognize that these sudden disturbances take the form of instantaneous, non-smooth displacement of the system in phase space. Our intent is to map a suitable control action to all regions in the phase space of the physical system in order to improve the robustness of the system.

The dynamic model for the cart-pole system is given by,

$$(M + m)\ddot{x} + ml\cos\phi\ddot{\phi} - ml\sin\phi\dot{\phi}^2 = F - f_c(\dot{x}) \tag{5}$$

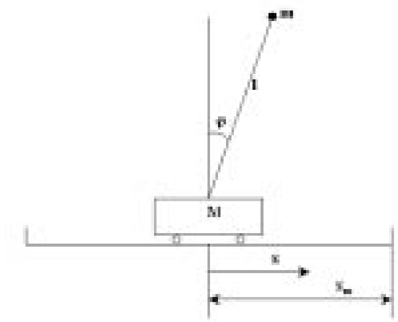

Fig. 2. The Cart-pole system

$$ml\cos\phi\ddot{x} + ml^2\ddot{\phi} - mgl\sin\phi = -f_p(\dot{\phi}) \tag{6}$$

where x, ϕ represent the cart position and pole angle respectively. The following simplified equations represent the dynamics of the pole, where we control the state of the system by applying $\ddot{x}$. The dynamics of the cart are defined by the time evolution of $\ddot{x}$. Equation (7) describes the system around the upper fixed-point $\phi = 0$, while equation (8) describes it around the lower fixed-point $\theta = 0$.

$$\ddot{\phi} + f(\dot{\phi}) - k\sin\phi + \ddot{x}\cos\phi = 0 \tag{7}$$

$$\ddot{\theta} + f(\dot{\theta}) + k\sin\theta - \ddot{x}\cos\theta = 0 \tag{8}$$

5 A Qualitative Heterogeneous Controller based on Time Scale Separation

The cart-pole system as described above includes a pendulum as a sub-system. However, the pivot-torque pendulum is controlled directly by u, while in the cart-actuated version, the effect of the applied acceleration $u = \ddot{x}$ on angular acceleration $\ddot{\phi}$ is scaled by $\cos\phi$. Furthermore, in the cart-pole system, we also have the control objective to stabilize the cart at $x = 0, \dot{x} = 0$ while the pole is stabilized vertically by the pendulum controller.

It is possible to control the position of the cart by a damped spring, similar to the *Balance* control action for the pole (2). The combined control action can be described as,

$$\ddot{x} = varsat\left\{-\frac{f_1(\dot{x})}{\cos\phi} - \frac{g_1(x)}{\cos\phi} + \frac{u(\phi,\dot{\phi})}{\cos\phi}\right\} \tag{9}$$

where $f_1(\dot{x}), g_1(x)$ represent any reasonable functions satisfying Lemma 1, used to regulate the position and velocity of the cart. The term $u(\phi,\dot{\phi})$ represents the control law (2) for the three control modes of the pole. $varsat$ refers to a saturation action whose magnitude depends on the controller mode. This allows the designer to set local control law parameters in such a way that the system never saturates in *Balance* while the *Pump* and *Spin* control actions can be restricted in strength.

$$varsat(x) = \begin{cases} sat(x, x_{\max\, bal}), controller = Balance \\ sat(x, x_{\max\, pump}), controller = Pump \\ sat(x, x_{\max\, spin}), controller = Spin \end{cases}$$

$$sat(x, x_{\max\, i}) = \begin{cases} -x_{\max\, i}, x < -x_{\max\, i} \\ x, -x_{\max\, i} \le x \le x_{\max\, i} \\ x_{\max\, i}, x > x_{\max\, i} \end{cases}$$

There are two independent requirements on the cart controller. When the pole controller is in the *Pump* and *Spin* modes, the addition of the cart stabilizing term to the pole control action should not affect the existence of the sliding mode between the *Pump* and *Spin* regions. On the other hand, when the pole controller is in the *Balance* mode, the sliding mode is not a concern and the primary requirement is that the composite system defined in terms of $[x, \dot{x}, \phi, \dot{\phi}]$ should be stable.

We now derive the constraints on cart control, corresponding to *Pump* and *Spin* modes of the pole controller. Taking the derivative of the expression for s, expressed in terms of variables ϕ and $\dot{\phi}$,

$$\dot{s}(\phi, \dot{\phi}) = -\dot{\phi} f(\dot{\phi}) + \dot{\phi} f_1(\dot{x}) + \dot{\phi} g_1(x) - \dot{\phi} u(\phi, \dot{\phi}) \tag{10}$$

Substitute the *Pump* control law,

$$\dot{s} = \dot{\phi}(h - f)(\dot{\phi}) + \dot{\phi} f_1(\dot{x}) + \dot{\phi} g_1(x) \tag{11}$$

In order to have $\dot{s} \ge 0$,

$$[(h - f)(\dot{\phi}) + f_1(\dot{x}) + g_1(x)]_0 = [\dot{\phi}]_0 \tag{12}$$

If we substitute the *Spin* control law,

$$\dot{s} = -\dot{\phi}(f + f_d)(\dot{\phi}) + \dot{\phi} f_1(\dot{x}) + \dot{\phi} g_1(x) \tag{13}$$

In order to have $\dot{s} \le 0$

$$[(f + f_d)(\dot{\phi}) - f_1(\dot{x}) - g_1(x)]_0 = [\dot{\phi}]_0 \tag{14}$$

We know that $[(h - f)(\dot{\phi})]_0 = [\dot{\phi}]_0$ and $[(f + f_d)(\dot{\phi})]_0 = [\dot{\phi}]_0$. From equations (12) and (14), we see that $f_1(\dot{x}) + g_1(x)$ need to be sufficiently small with respect to $(h - f)(\dot{\phi}), (f + f_d)(\dot{\phi})$ for the sign equality to be preserved (even if these terms are opposed to each other). It is conceivable that for small values $\dot{\phi} \to 0$, these constraints may be violated for $f_1(\dot{x}) + g_1(x) \ne 0$. Fortunately, the sliding mode constrains the trajectory to be on the separatrix where it is bounded away from $\dot{\phi} = 0$ by the width of the *Balance* region (equation 3). Therefore, it is possible to select f_1 and g_1 to be sufficiently small to satisfy equations (12) and (14). This amounts to requiring a sufficient separation between the time-scales of the pole and the cart controllers.

We then need constraints on cart control, corresponding to the *Balance* mode of the pole controller. The general principle behind this analysis is summarized below.

The uncontrolled cart pole system can be linearized and written as,

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} = \begin{bmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & k & -c \end{bmatrix} \begin{bmatrix} x \\ \dot{x} \\ \phi \\ \dot{\phi} \end{bmatrix} + \begin{bmatrix} 0 \\ 1 \\ 0 \\ -1 \end{bmatrix} \ddot{x}, \tag{15}$$

where $\mathbf{x} = \begin{bmatrix} x \ \dot{x} \ \phi \ \dot{\phi} \end{bmatrix}^T$, $\mathbf{u} = [\ddot{x}]$, and $f(\dot{\phi}) = c\dot{\phi}$. It is easy to verify that $(\mathbf{A}, \mathbf{B})$ is controllable (i.e., the controllability matrix has full rank). This implies that there exists some feedback control action that can place the closed loop eigenvalues of the controlled system at any point in the left half of the complex plane.

Now, the controlled, nonlinear equation of the higher order system, $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$, can be similarly linearized to yield the model, $\dot{\mathbf{x}} = \tilde{\mathbf{A}}\mathbf{x}$ where $\tilde{\mathbf{A}}$ refers to the Jacobian linearization of the controlled system, $\tilde{\mathbf{A}} = \partial \mathbf{f}/\partial \mathbf{x}|_{\mathbf{x}=\mathbf{0}}$. Applying Lyapunov's Linearization method [16, 3], we know that if the linearized system has $\tilde{\mathbf{A}}$ Hurwitz (i.e., if all eigenvalues of $\tilde{\mathbf{A}}$ are strictly in the left-half complex plane), then the equilibrium point of the nonlinear system, $\dot{\mathbf{x}} = \mathbf{f}(\mathbf{x})$, is asymptotically stable. We place the closed loop eigenvalues by appropriate selection of controller parameters to make $\tilde{\mathbf{A}}$ Hurwitz.

Consider an instance of the controller defined by the QDE in equation (9), with a linear pole *Balance* controller,

$$\ddot{x} = -\frac{d_1}{\cos\phi}x - \frac{d_2}{\cos\phi}\dot{x} + \frac{(c_{11}+k)\phi + c_{12}\dot{\phi}}{\cos\phi} \tag{16}$$

The controlled pole equation can be written as,

$$\ddot{\phi} = -c\dot{\phi} + k\sin\phi + d_1 x + d_2\dot{x} - (c_{11}+k)\phi - c_{12}\dot{\phi} \tag{17}$$

From this, the Jacobian linearization $\dot{\mathbf{x}} = \tilde{\mathbf{A}}\mathbf{x}$ becomes,

$$\dot{\mathbf{x}} = \begin{vmatrix} 0 & 1 & 0 & 0 \\ -d_1 & -d_2 & c_{11}+k & c_{12} \\ 0 & 0 & 0 & 1 \\ d_1 & d_2 & -c_{11} & -c-c_{12} \end{vmatrix} \mathbf{x} = \tilde{\mathbf{A}}\mathbf{x} \tag{18}$$

We select parameters that make $\tilde{\mathbf{A}}$ Hurwitz. In many physical systems, the magnitudes and/or direction of x and $\dot{\phi}$ may be mismatched. In that case, the appropriate signs and scaling factors would need to be incorporated into the above analysis. The last remaining issue is that of determining the region of attraction for this stable equilibrium. Here one would use a converse Lyapunov argument based on the fact that $\tilde{\mathbf{A}}$ is Hurwitz if and only if, for any given symmetric positive definite $\mathbf{Q}$, there exists a unique symmetric, positive definite $\mathbf{P}$ which is a solution to the Lyapunov equation:

$$\mathbf{PA} + \mathbf{A^T P} + \mathbf{Q} = \mathbf{0}$$

This matrix $\mathbf{P}$ is the basis for the definition of a Lyapunov function, $\mathbf{V}(\mathbf{x}) = \mathbf{x^T P x}$. The boundary of the *Balance* region is defined by a level curve of $\mathbf{V}(\mathbf{x})$. In practice, the Lyapunov equation may be difficult to solve for **P**, in which case it may be possible to determine the domain of attraction experimentally or by numerical simulations (see [5]).

The above analysis assumed a linear spring model for cart control. One may also use a nonlinear spring controller. The nonlinear spring controller provides the advantage that the cart position can be kept away from the ends of the track, i.e., $|x| < x_m$, by the potential barrier of the spring action. Consider an instance of a nonlinear spring defined by,

$$\ddot{x} = sat\left\{ -\frac{d_1}{\cos\phi} \tanh^{-1}\left(\frac{x}{x_m}\right) - \frac{d_2}{\cos\phi}\dot{x} + \frac{u(\phi, \dot{\phi})}{\cos\phi} \right\} \tag{19}$$

The first term provides the potential barrier necessary to prevent the cart from hitting the ends of the track. Note that $\tanh^{-1}$ (which is in M_0^+) is very near linear over $[-0.5, +0.5]$ and diverges to $\pm\infty$ at ± 1, respectively. As $|x| \to x_m$, the nonlinear cart control action (equation 19) asymptotically approaches infinity and theoretically prevents the cart from reaching the end of the track. In practice, if $\ddot{x}$ saturates, then this potential barrier cannot always prevent the cart from hitting the end of the track. However, by reducing the size of the *Balance* region by reducing ϕ_{max} and hence $\dot{\phi}_{max}$, and by bounding the value of $f_d(\dot{\phi})$ in the *Spin* region, it is possible to constrain the system so that the saturated $\ddot{x}$ is sufficient to keep the cart from hitting the ends of the track.

The controller designed thus far was implemented on a physical system. Figure 3 shows the phase plots of the controlled cart and pole systems. It is seen that the system stabilizes within a very small neighborhood of the point $\phi = 0, \dot{\phi} = 0, x = 0, \dot{x} = 0$. The physical system exhibits a limit cycle about the equilibrium point, due to effects such as static friction that are often omitted from simulations. We address these next.

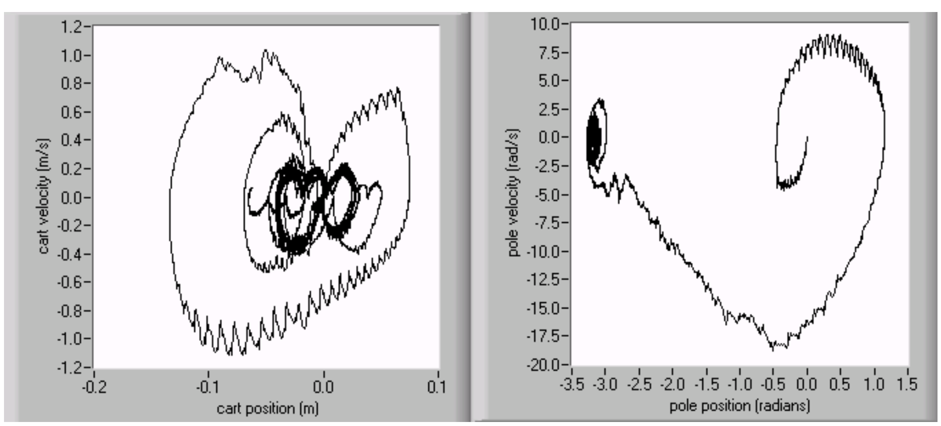

Fig. 3. Phase plots for the cart and pole subsystems. This implementation utilized a qualitative cart-pole controller based on time scale separation.

6 Accommodating real-world effects, the case of static friction

Consider the damped oscillator with static, or Coulomb, friction F_c as well as damping friction $f(\dot{x})$. The resistance of static friction is a constant force $F_c = \eta$ opposing the direction of motion, as long as motion is taking place. Once motion stops, a larger frictional resistance $F_c = \eta + \varepsilon$ must be overcome to get it started again. This transforms the damped oscillator model from $\ddot{x} = -g(x) - f(\dot{x}) - F_c$ to

$$\ddot{x} = \begin{cases} \dot{x} \neq 0 \rightarrow \quad -g(x) - f(\dot{x}) - [\dot{x}]_0 \eta \\ \dot{x} = 0 \rightarrow -[x]_0 \max(|g(x)| - \eta - \varepsilon, 0) \end{cases} \tag{20}$$

To avoid being captured by static friction, we require that the system obey the constraint

$$\dot{x} = 0 \quad \Rightarrow \quad |g(x)| > \eta + \varepsilon \tag{21}$$

This constraint excludes the stable fixed-point, where $x = \dot{x} = \ddot{x} = 0$. Therefore, we will be required to approach a more complex orbit, such as a limit cycle. This arises under two circumstances.

First, in the *Pump* region, the oscillator with negative damping should diverge from its unstable fixed-point at $(\theta, \dot{\theta}) = (0, 0)$, but the initial oscillations in its trajectory are so small as to be absorbed by static friction, so the pendulum never starts pumping. We handle this problem by defining a *Startup* control law, which is only applied when the system state is $[x, \dot{x}, \theta, \dot{\theta}] = [0, 0, 0, 0]$, and which specifies that $\ddot{x}(t) = u_0$ for $t \in [0, \tau]$. The constant cart acceleration $u_0 > \eta + \varepsilon$ and the time interval τ are selected so that the final state $t = \tau$ is contained within the *Pump* region, and is far enough from the origin that an energy argument guarantees that constraint (21) will continue to be satisfied.

Second, in the *Balance* region, as the pendulum approaches the fixed-point at $(\phi, \dot{\phi}) = (0, 0)$, static friction can capture it at a small but perceptibly non-zero value of ϕ, leading to a constant non-zero value of $\ddot{x}$, and hence a runaway cart. We handle this problem by modifying the *Balance* controller to approach a limit cycle rather than a fixed-point, and designing the limit cycle to satisfy constraint (21). In order to create a limit cycle, we require a new qualitative behavior model, since the damped oscillator does not include limit cycles in its behavioral repertoire. However, the Lienard equation describes a set of familiar systems that do exhibit limit cycles, and it makes an excellent target for the QHC methodology [16, 17].

A Lienard system (e.g., the van der Pol oscillator) is defined in terms of the model $\ddot{x} + f(x)\dot{x} + g(x) = 0$ where, broadly speaking, $f(x)$ is positive when $|x|$ is large and negative when $|x|$ is small, and g is such that, in the absence of the damping term $f(x)\dot{x}$ we expect periodic solutions for small x. The property can be summarized by the following lemma. This lemma is taken from the standard literature, see [16] for the proof.

Lemma 3. *Let $A \subseteq \Re^2$ include (0,0) in its interior, and let S be a system governed by the QDE $\ddot{x} + f(x)\dot{x} + g(x) = 0$ for every $(x, \dot{x}) \in A$, where f and g are reasonable functions. Define the function,*

$$F(x) = \int_0^x f(u)du$$

S has a unique, stable limit cycle surrounding the origin if
$g(x)$ is an odd function and $g(x) > 0$ for $x > 0$ (i.e., $[g(x)]_0 = [x]_0$),
$f(x)$ is an even function,
$F(x)$ is an odd function, which has exactly one positive zero at $x = a$,
$F(x)$ is negative in $0 < x < a$, is positive and nondecreasing in $x > a$, and $F(x) \to \infty$ as $x \to \infty$.

We utilize this lemma and define a switched system inside *Balance*. Starting with the system

$$\ddot{\phi} + c\dot{\phi} - k\sin\phi + \ddot{x}\cos\phi = 0 \tag{22}$$

In the region, $-a/2 < \phi < a/2$, we define the controller,

$$\ddot{x} = \frac{-(c+p_1)\dot{\phi}}{\cos\phi} + \frac{(k+p_2)\phi}{\cos\phi} \tag{23}$$

so that the closed loop system takes the form (assume $\cos\phi \approx 1, \sin\phi \approx \phi$),

$$\ddot{\phi} - p_1\dot{\phi} + p_2\phi = 0 \tag{24}$$

In the region $(-\phi_{\max} < \phi < -a/2) \cup (a/2 < \phi < \phi_{\max})$, the controlled system is,

$$\ddot{\phi} + (c + c_{12})\dot{\phi} + c_{11}\phi = 0 \tag{25}$$

Setting $(c + c_{12}) = p_1; c_{11} = p_2$ satisfies the conditions in Lemma 3.

Now, Balance was defined such that any trajectory that enters the region will remain inside indefinitely. In order to ensure that the addition of the above controller does not violate this property, the value of a should be chosen such that the energy added in the region $-a/2 < \phi < a/2$ is not greater than what can be dissipated in the region $(-\phi_{\max} < \phi < -a/2) \cup (a/2 < \phi < \phi_{\max})$. This analysis would be identical to that in [2], where energy changes in *Pump* and *Balance* were derived.

The effect of including this region is that the cart-pole system executes a limit cycle as seen in Figure 4. The advantage of this approach lies in the fact that relaxation oscillators retain their structure through a very wide range of parameters p_1, p_2, and overcome practical problems such as static friction [16, 17].

7 A Linear Quadratic Regulator for Balance

A controller is defined by the qualitative constraints it satisfies and its region of applicability. In QHC, it is possible to include local controllers designed by different methodologies, as long as they satisfy the desired qualitative constraints inside the region of applicability. To illustrate this, we designed a linear quadratic regulator (LQR)

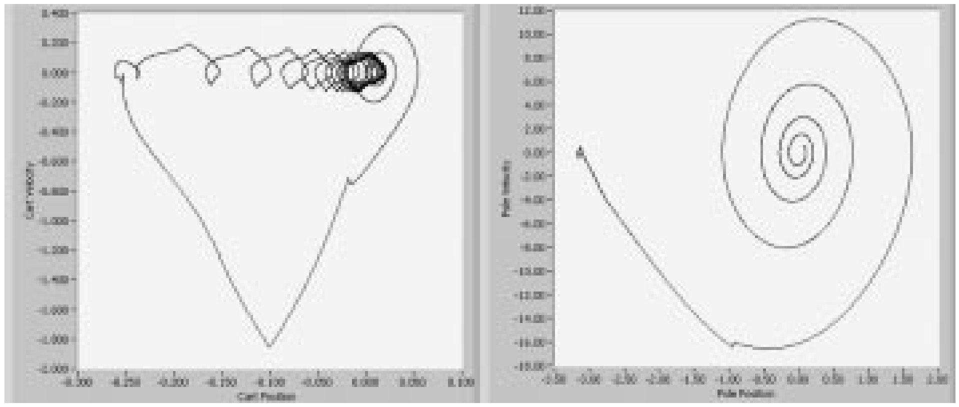

Fig. 4. Phase plots for the cart and pole subsystems, simulation results. These plots illustrate the system executing a limit cycle in both the cart and pole subsystems.

as the *Balance* controller. The primary requirement on the *Balance* controller is that any trajectory that has entered *Balance* should remain inside until perturbed by an external process. LQR possesses this property, as described in [18]. In addition, $\phi_{\max}$ and $\dot{\phi}_{\max}$ can be selected in such a way that the closed loop response of the cart satisfies $-x_m < x < x_m$.

Linear Quadratic Regulators stabilize the plant $\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u}$, $\mathbf{x} = [x \; \dot{x} \; \phi \; \dot{\phi}]^T$ by applying a control, $\mathbf{u} = -\mathbf{G}\mathbf{x}$ where $\mathbf{G}$ is selected so as to minimize a cost function, $J = \int_0^\infty \mathbf{x^T Q x} + \mathbf{u^T R u} dt$. For our example, we select $\mathbf{Q} = diag(4, 0, 4, 0)$ and $\mathbf{R} = [0.5]$ to penalize cart and pole positions equally, control action to a lesser extent, and cart and pole velocities not at all. Other parameter choices do not affect the substance of our argument.

The operation of this controller is explained by the fact that the controller iteratively minimizes the cost function J which is an ‘energy-like’ quadratic function. Further, an analysis of the eigenvalues of the closed loop system indicates that all eigenvalues are negative and the system is stable as required.

Figure 5 shows the experimental results of a heterogeneous controller that utilized an LQR *Balance* controller. As expected, a large part of the heterogeneous behavior is identical to the experimental result in Figure 3. The primary observable difference is in the size of the limit cycle in the cart phase portrait, in the *Balance* mode of pole control.

8 Discussion

8.1 Robustness in global behaviors

One of the goals motivating this investigation has been the search for control strategies that are visibly robust in the face of environmental disturbances and user interaction. It is well known that algorithms such as linear quadratic regulators can guarantee asymptotic stability only for initial states that are within a finite region of the origin. Nonlinear

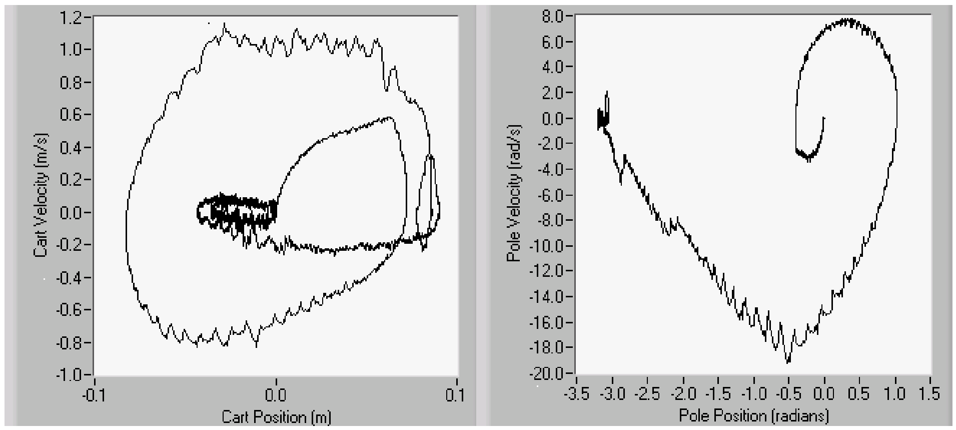

Fig. 5. Phase plots for the cart and pole subsystems. This implementation utilized a qualitative cart-pole controller with LQR in the Balance region.

control design techniques such as feedback linearization may also suffer from lack of robustness [3, 19]. In the QHC methodology, we define local models based on qualitative behavior models. These qualitative behavior models are selected as dynamical systems with stable orbits, such as fixed points and limit cycles that are robust in the face of parametric and structural uncertainties. This makes the global behavior, defined as a composition of these local models, correspondingly robust.

We now present some experimental results to illustrate the results we are able to obtain from the controlled hybrid system. One of our practical goals has been to design a cart-pole controller that can accommodate a large amount of abuse from the user – in terms of state perturbations. The effect of such a perturbation on a traditional control algorithm such as LQR is that the system is pushed outside the region of applicability that results in loss of control or the cart hits the ends of the track, in response to a large control action.

In our design, a large perturbation causes the controller to switch to the *Spin* mode that provides augmented damping until the state variables are within the region of applicability of *Balance*. This behavior is illustrated in Figure 6. It is seen that the system has initially stabilized and then the user imparts a perturbation to the pole velocity, corresponding to an instantaneous change of 30 rads/s. This causes the system to leave *Balance*, to execute *Spin, Spin-Pump Sliding Mode* and eventually to return to *Balance*. In this way, the heterogeneous controller accommodates a wide range of perturbations in state space without losing the global behavior. This is a useful form of robustness for many applications, especially in robotic and biomechanical systems such as those described in [5, 6, 13–15].

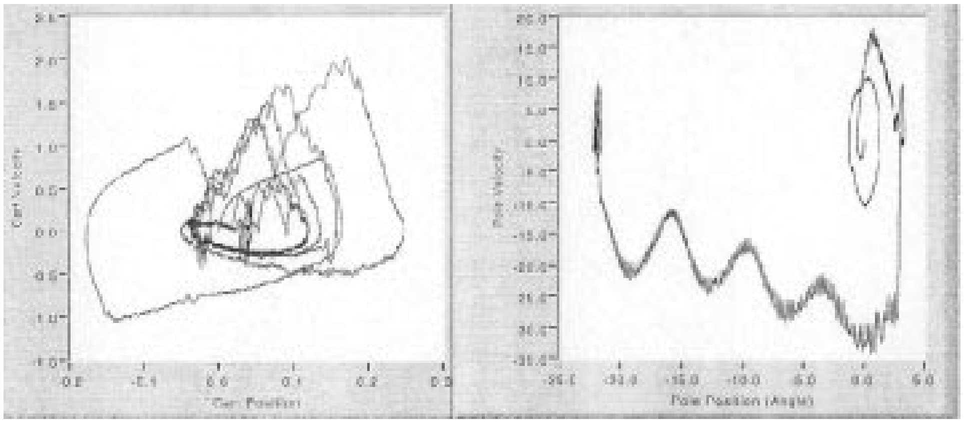

Fig. 6. Phase plots for the cart and pole subsystems. These plots illustrate the system recovering from a hard perturbation in pole velocity.

8.2 Transition graph structure of the heterogeneous controller

The operation of the heterogeneous controller can be described and analyzed in terms of a discrete transition graph, as seen in Figure 7. Some of the more important properties of this graph are summarized below.

1. The graph G1, representing the pole control action, is Figure 1 extended with local regions *Startup* and *Limit Cycle* to handle static friction at very low velocities. The graph G2 represents the cart control action, which has a single mode.
2. The two graphs G1 and G2 execute concurrently. This is the default mode in graph G3. The arguments in this paper provide the basis for stable concurrent operation of G1 and G2.
3. The additional state in G3 provides a way to handle large perturbations. If an excessively high velocity is imparted to the pole, rapidly varying, possibly high amplitude, control signals could be generated by the *Spin* region. To prevent this, the action of the pole controller can be restricted to a subset of the state space, with bounded $\dot{\phi}$. It is clear that this is an invariant set, and any trajectory entering it will converge to the desired setpoint. If the system is perturbed outside this invariant set, the pole controller is turned off (leaving the pole to its natural damping), and damping in the cart controller is augmented. The trajectory must enter the invariant set and stabilize at the desired equilibrium.
4. Energy arguments can be used to determine the residence time of a trajectory in any region, showing that all except *Limit Cycle* have finite residence time. For example, $\dot{s}$ in equation (10) is the instantaneous rate of change of energy in the *Pump* and *Spin* regions, and $s = 0$ defines their shared boundary.

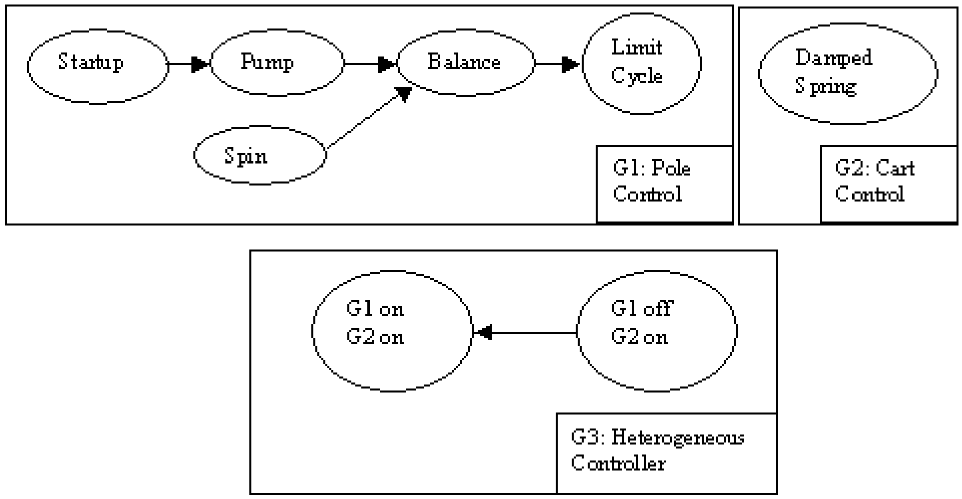

Fig. 7. Transition graph structure of the heterogeneous controller for the cart-pole system

8.3 Conclusions

By applying a control function that transforms the uncontrolled system into one that is an instance of a qualitative behavior model with provable properties, it is possible to design simple, intuitive and robust controllers. QHC applies this concept in a multiple model framework to design global controllers for nonlinear systems. Although the QHC methodology is presented here and in [2] through the simple familiar example of the inverted pendulum, the design method is quite general.

1. Define a transition graph of local regions that guarantees the desired global behavior.
2. For each local region, select a qualitative behavior model (e.g., the damped oscillator) that has the desired qualitative behavior over a region including the intended local region.
3. Select a control policy for each local region that transforms the uncontrolled plant into an instance of the qualitative behavior model over some region containing the intended local model.
4. Define region boundaries with simple reliable descriptions (e.g., energy level curves) such that the local behaviors are guaranteed to cross the region boundaries exactly in the desired ways.

This algorithm is non-deterministic in the sense that earlier choices must be made correctly for later choices to be possible, so backtracking may be necessary. However, if the algorithm terminates, the resulting design is guaranteed to have the desired properties.

Furthermore, since the models are qualitative, the resulting design describes an entire *family* of control laws, all of which are guaranteed to have the desired properties.

This allows a useful separation of concerns between qualitative correctness and optimization. Other advantages of QHC include structured handling of real world effects such as static friction, and robustness of global behavior in the face of parametric and structural uncertainty, and in the face of substantial perturbations.

References

1. Kuipers, B.J., Astrom, K.: The composition and validation of heterogeneous control laws. *Automatica* **30** (1994) 223-249
2. Kuipers, B., Ramamoorthy, S.: Qualitative modeling and heterogeneous control of global system behavior. In C. J. Tomlin and M. R. Greenstreet (Eds.), *Hybrid Systems: Computation and Control*, Lecture Notes in Computer Science, Volume 2289, Springer Verlag, 2002
3. Slotine, J.J.E., Li, W.: *Applied Nonlinear Control*. Prentice Hall, Englewood Cliffs NJ (1991)
4. Lohmiller, W., Slotine, J.J.E.: On contraction analysis for nonlinear systems. *Automatica* **34** (1998) 683-696
5. Burridge, R.R., Rizzi, A.A., Koditschek, D.E.: Sequential composition of dynamically dexterous robot behaviors. *International Journal of Robotics Research* **18** (1999) 534—555
6. Klavins, E., Koditschek, D.E.: Phase regulation of decentralized cyclic robotic systems. *International Journal of Robotics Research* **21** (2002) 257—275
7. Ott, E., Grebogi, C., Yorke, J.A.: Controlling chaos. *Physical Review Letters*. **64** (1990) 1196-1199
8. Boccaletti, S., Grebogi, C., Lai, Y.C., Mancini, H., Maza, D.: The control of chaos: Theory and Applications. *Physics Reports* **329** (2000) 103-197
9. Bradley, E.: Autonomous exploration and control of chaotic systems. *Cybernetics and Systems* **26** (1995) 299-319
10. Astrom, K.J., Furuta, K.: Swinging up a pendulum by energy control. *Automatica* **36** (2000) 287-295
11. Zhao, J., Spong, M.W.: Hybrid control for global stabilization of the cart-pendulum system. *Automatica* **37** (2001) 1941-1951
12. Chung, C.C., Hauser, J.: Nonlinear control of a swinging pendulum. *Automatica* **31** (1995) 851-862
13. Full, R., U. Saranli, M. Buehler, Brown. B., N. Moore, D. Koditschek, and H. Komsuoglu. Evidence for Spring Loaded Inverted Pendulum Running in a Hexapod Robot. *Proc. International Symposium on Experimental Robotics*, Honolulu, Hawaii, 2000.
14. Saranli, U., Schwind, W.J., Koditschek, D.E.: Toward the control of a multi-jointed monoped runner. *Proceedings of the IEEE International Conference on Robotics and Automation*, pp 2676–2682, 1998
15. Pratt, J. and Pratt, G.: Intuitive control of a planar bipedal walking robot. *Proceedings of the IEEE International Conference on Robotics and Automation*, 1998
16. Jordan, D.W., Smith, P.: *Nonlinear Ordinary Differential Equations: An Introduction to Dynamical Systems*. 3^{rd} ed. Oxford University Press, Oxford (1999)
17. Guckenheimer, J., Holmes, P.: *Nonlinear Oscillations, Dynamical Systems, and Bifurcations of Vector Fields*. Springer-Verlag, Berlin (1983)
18. Levine, W.E., ed., *The Control Handbook.* CRC Press (1996)
19. Friedland, B.: *Advanced Control System Design*. Prentice-Hall (1996)