

Sequential Composition for Control of Underactuated Systems

George Kantor

Alfred A. Rizzi

CMU-RI-TR-03-23

November 2003

Robotics Institute
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213
{kantor,arizzi}@ri.cmu.edu

© Carnegie Mellon University

Abstract

We present a new approach to developing hybrid feedback policies for the control of systems with nonholonomic constraints. We extend the idea of *sequential composition* and use it to switch between controllers in a state based manner, resulting in a globally convergent, pure feedback policy. Individual controllers in the palette are inspired by *variable constraint control*. We solve a vision guided unicycle navigation problem and verify the result through simulation and experiment. This technical report presents a detailed description of results published in [10].

1 Introduction

This paper proposes a framework for the construction of globally convergent purely feedback based controllers for underactuated kinematic systems. To accomplish this, we combine two ideas that have recently been presented in the literature. We use the idea of *variable constraint control (VCC)* [8, 9] to divide the problem into a natural sequence of subproblems that can be solved via feedback control. We then use *sequential composition* [3] to connect a palette of controllers and automatically orchestrate their switching to generate a globally convergent feedback policy.

To demonstrate this framework, we present a controller that solves the navigation problem for a unicycle-like robot operating in a plane. Feedback for the controller is provided by a robot-mounted camera that measures the relative pose between the robot and a fixed engineered landmark. The resulting system respects the constraint that the landmark remain within the field of view of the robot-mounted camera at all times.

This work builds on ideas developed in the large body of work that exists in the field of controlling underactuated systems. The core of that work focuses on the development of open loop steering strategies that generate desired motions by creating sequences of piecewise constant inputs [12] and by using oscillatory inputs with carefully chosen phase and frequency relationships [13, 16]. Brockett showed, early on, that it is not possible to stabilize an underactuated system with a smooth, time invariant feedback controller [2]. However, a number of authors have found ways to work around this fact by using time varying feedback [20], non-smooth feedback [21], exploiting singularities in the workspace [1, 6], adding to the system state [15], and, as in the case of the previously mentioned variable constraint control, a hybrid of continuous and discrete feedback [8, 17].

This paper is also related to recent work in visual servoing for mobile robot control. The idea of using sequential composition for visual servoing tasks was introduced in [4] and further described in [5]. Examples that employ visual servoing in the presence of nonholonomic constraints include [18, 7, 22]. Of particular relevance to this paper, [17] describes a hybrid control strategy for a nonholonomic robot subject to additional constraints from limited camera field of view.

The controller presented here exhibits a number of features which, taken together, distinguish it from the previous work in the area. It is a pure feedback policy; the action taken by the robot at any time is based solely on the current robot pose. No path is planned *a priori* and the controller requires no additional state. Both holonomic and nonholonomic constraints are explicitly considered in the controller design. Most importantly, the use of sequential composition provides a switching policy that, in the case where there is no sensor noise, guarantees an absence of undesirable behaviors such as limit cycles, chattering, and oscillatory switching patterns.

2 Previous Work

2.1 Variable Constraint Control

The control strategy presented in this paper borrows conceptually from a recently introduced method of feedback control for underactuated kinematic systems known as *variable constraint control (VCC)* [8]. In VCC, a submanifold of the state space is found that contains the goal state and on which one of the nonintegrable constraints can be replaced by an integrable constraint. The integrable constraint then induces a foliation of the submanifold with the leaves of the foliation being the integral curves of the (now holonomic) constraint. The idea is then to identify the leaf on which the goal state lies and then steer the system to that leaf. Once the leaf is reached, any control that holds the state on the submanifold will also implicitly hold the state on the leaf, effectively reducing the dimension of the control problem.

This concept can then be iterated, successively driving to a sequence of manifolds of progressively lower dimension until the goal state is reached.

While the idea of VCC is applicable to systems of general dimension, we focus on the case of small time locally controllable (STLC) systems with three states and two inputs, both because VCC is easy to explain concretely for such systems and because the unicycle robot addressed in this paper meets this description. Consider the STLC drift free affine control system

$$\dot{q} = g_1(q)u_1 + g_2(q)u_2 \triangleq Bu,$$

with state $q \in \mathbb{R}^3$ and input $u = [u_1 \ u_2]^T \in \mathbb{R}^2$. Denote the goal state by q_f . Since the system is underactuated, it has an associated constraint that can be written in Pfaffian form as $A(q)\dot{q} = 0$. Because the system is STLC, the constraint is guaranteed to be nonintegrable, i.e. there does not exist any function $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ such that $\frac{\partial f}{\partial q} = A(q)$.

The first step in the VCC synthesis process is to find a function $H : \mathbb{R}^3 \rightarrow \mathbb{R}$ that satisfies two properties:

1. $H(q_f) = 0$,
2. there exists a set $\mathcal{H} = \{q \in \mathbb{R}^3 | H(q) = 0\}$ and an integrable constraint $\hat{A}(q)$ such that

$$A(q)|_{\mathcal{H}} = \hat{A}(q)$$

Given a suitable H , the next step requires creating a second function $F : \mathbb{R}^3 \rightarrow \mathbb{R}$ that satisfies two properties:

1. $F(q_f) = 0$,
2. F is an integral of $\hat{A}(q)$, i.e. $\frac{\partial F}{\partial q} = \hat{A}(q)$.

Finally construct a third function $N : \mathbb{R}^3 \rightarrow \mathbb{R}$ that also satisfies two conditions:

1. $N(q_f) = 0$
2. $N(q)$ must be independent of H and F , i.e. the gradient $\frac{\partial N}{\partial q}$ is linearly independent of $\frac{\partial H}{\partial q}$ and $\frac{\partial F}{\partial q}$ on the set $\{q \in \mathbb{R}^3 | F(q) = H(q) = 0\}$.

Note that because of the independence condition on N , q_f is the unique point where $H(q) = F(q) = N(q) = 0$. With H, F , and N in hand, the procedure for driving to the goal state is straightforward. First steer to the 2-dimensional manifold defined by $\{q \in \mathbb{R}^3 | F(q) = 0\}$. Practically, this can be achieved with a closed loop controller that forces F toward zero. Then, holding F at zero, steer to the 1-dimensional manifold $\{q \in \mathbb{R}^3 | F(q) = H(q) = 0\}$. This can be achieved via feedback control that forces both H and F toward zero simultaneously. Finally, hold H at zero, and steer to $N(q) = 0$ with a closed loop controller that forces H and N toward zero. Note that in the last step F is not explicitly controlled. However, since F is an integral of $\hat{A}(q)$, $F(q)$ will remain zero as long as $H(q)$ is held at zero. Thus the goal state is reached when N reaches zero.

2.2 Sequential Composition

In contrast to traditional approaches to navigation problems of this class, we focus on the development of globally (or at least essentially globally) convergent feedback strategies to accomplish the task. Since our task centers on regulation of a system that exhibits a nonholonomic constraint we know that there does not exist such a policy that is globally smooth [2], and we are resigned to using strategies that incorporate switching behavior. However, the policy we propose utilizes a deterministic (state based) switching strategy to *sequentially compose* the individual controllers and guarantees global convergence as a direct consequence of the stability properties of the underlying behaviors.

Specifically we rely on a generalization of the notion of *sequential composition* first proposed in [3] and applied to more general systems in [11, 18, 19]. The notion underlying this scheme is to specify the control system in terms of a collection of strategies and an associated selection (switching) scheme selects the active strategy based only on the current state of the system. The result is a hybrid on-line control policy that, if properly deployed, makes use of the entire collection of available controllers to systematically make progress toward a goal based on state.

More formally, given a set of controllers, $\mathcal{U} = \{\Phi_1, \dots, \Phi_N\}$, each with an associated goal, $\mathcal{G}(\Phi_i)$, and domain, $\mathcal{D}(\Phi_i)$. It is presumed that under the action of Φ_i any state that starts in $\mathcal{D}(\Phi_i)$ will be taken to $\mathcal{G}(\Phi_i)$ without leaving the set $\mathcal{D}(\Phi_i)$. We then say that controller Φ_1 prepares controller Φ_2 , denoted $\Phi_1 \succeq \Phi_2$, if its goal lies within the domain of the second, $\mathcal{G}(\Phi_1) \subset \mathcal{D}(\Phi_2)$ [3]. For an appropriately parameterized set of controllers, \mathcal{U} , this relation induces a generally cyclic directed graph. We assume that the overall goal \mathcal{G} is contained in the goal of at least one controller, i.e. $\mathcal{G} \subset \mathcal{G}(\Phi_n)$. Then by starting with Φ_n and recursively tracing the prepares relation backwards through the corresponding graph, one arrives at the set of all controllers from whose domains the overall goal might be eventually reached by switching between control policies. We denote this set $\mathcal{U}_{\mathcal{G}} \subset \mathcal{U}$. This backwards trace induces a partial order over the set of controllers and their domains that can be searched to produce a state based controller selection strategy. In practice this is accomplished by searching the available controllers from highest to lowest priority until a controller whose domain contains the state is found.

The component control strategies presented in Section 3.5 depart from this approach in that their domains are not strictly speaking invariant and that their goal sets are no longer point attractors. Rather, we generalize the approach to utilize goal sets that are full dimensional sets, and consider their domain to be the set of states that lead to entry of the goal set. That is to say if $x \in \mathcal{D}(\Phi_i)$, then under the influence of Φ_i , x will remain in $\mathcal{D}(\Phi_i)$ until $x \in \mathcal{G}(\Phi_i)$, giving us a notion of conditional positive invariance until satisfaction of the goal condition.

VCC is distinguished from other hybrid control synthesis techniques in that it guarantees “nice” switching behaviors. The invariance property of the individual control policies ensures that (in the absence of disturbances) that progress through the induced graph will be monotonic towards the goal. This eliminates the possibility of chattering, limit cycles, and other oscillatory switching patterns often associated with hybrid control systems. This monotonic progression through the graph also makes the convergence properties of the composite controller easy to understand: its domain will be the union of the domains of all of the individual controllers: $\bigcup_{\Phi \in \mathcal{U}_{\mathcal{G}}} \mathcal{D}(\Phi)$. Thus we have available an “automatic” method by which to guide the system from any state in this union of domains to the goal.

3 Unicycle Navigation

3.1 Measuring Robot Pose

To support feedback performance of the positioning task described in Section 1 we must provide a means for the robot to measure its pose relative to some task frame. We accomplish this by endowing the robot with a camera system and providing a fixed landmark (beacon) that can be observed from the robot. The specific beacon we utilize is a known sized cube, with each face painted a different color. Utilizing a color camera the robot is able to measure the relative pose between itself and this beacon from a single observation. The physical instantiation of this system has been developed in our laboratory [14] and can effectively measure the six dimensional, $SE(3)$, relative pose between the cube and robot from a single observation. For the purposes of this paper, the robot motion is restricted to the plane, so we project the measurement into $SE(2)$.




Figure 1: Unicycle coordinate system and camera configuration.

Note that there are some limitations on this type of sensor. Obviously, the landmark must be within the field of view of the camera. Furthermore, the projection of the landmark onto the image plane must be large enough to enable an

accurate estimate of pose and small enough to fit within the image. These constraints induce an upper and lower bound on the distance the robot can be from the landmark, denoted d_{\max} and d_{\min} , respectively. Accordingly, it is convenient to represent the pose of the robot in a kind of polar $SE(2)$ coordinates with the origin located at the landmark, as shown in Figure 1. The robot pose is denoted by $q = [d \ \varphi \ \eta]^T$, where φ is the bearing to the landmark from the robot, and η is the bearing to the robot from the landmark, and

$$d = r - d_{\text{mid}} \triangleq r - \frac{d_{\min} + d_{\max}}{2},$$

where r is the distance between robot and landmark. The set of all such q is denoted \mathcal{Q} .

If the camera is mounted so that the focal plane normal is aligned with x -axis in the robot frame, then the set of robot poses from which pose can be measured, denoted Ω_S , can be easily expressed as

$$\Omega_S = \{q \in \mathcal{Q} \mid |\varphi| < \varphi_c, |d| < D\},$$

where $\varphi_c < \frac{\pi}{4}$ is the field of view of the camera and $D \triangleq \frac{d_{\max} - d_{\min}}{2}$. We note that this formulation of both pose coordinates and the visible set Ω_S is very similar to the formulation proposed in [5].

3.2 Unicycle Kinematics

The kinematics of the unicycle robot in the coordinate systems described above are given by

$$\frac{d}{dt} \begin{bmatrix} d \\ \varphi \\ \eta \end{bmatrix} = \begin{bmatrix} -\cos\varphi \\ \frac{\sin\varphi}{d+d_{\text{mid}}} \\ -\frac{\sin\varphi}{d+d_{\text{mid}}} \end{bmatrix} u_1 + \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} u_2, \quad (1)$$

where the real valued control inputs u_1 and u_2 represent the robot forward and angular velocities, respectively.

The underlying no-sideways-slipping constraint can be written in Pfaffian form as $A(q)\dot{q} = 0$, where

$$A(q) = \begin{bmatrix} -\frac{\tan\varphi}{d+d_{\text{mid}}} & 0 & 1 \end{bmatrix}.$$

Now the problem this paper addresses can be clearly stated:

Problem 1 (Vision-Based Unicycle Navigation) *Given a goal state $q_f = [d_f \ 0 \ 0] \in \Omega_S$ and a tolerance ϵ , find a closed loop control policy $u : \Omega_S \rightarrow \mathbb{R}^2$ so that for any initial condition $q(0) = q_0 \in \Omega_S$ and for some finite time t_f , the solution of Equation 1 under u satisfies*

1. For all $t > t_f$, $q(t) \in B_\epsilon(q_f) \triangleq \{q \in \mathcal{Q} \mid \|q - q_f\|_\infty < \epsilon\}$.
2. $q(t) \in \Omega_S$ for all $t > 0$.

The assumption that φ_f and η_f in q_f are both zero is made for clarity of presentation. The technique presented in this paper can be used for general q_f , however the resulting complex expressions that result add a distracting level of tedium. In practice, the assumption on q_f does not reduce the generality of the problem anyway. Since the unicycle kinematics do not rely explicitly on η , goal states that have non-zero η_f can be accomplished by defining a new, rotated landmark frame in which the corresponding goal η is zero. And since the unicycle can turn in place, goal states that have non-zero φ_f can be achieved by independently driving φ to φ_f after the goals for d and η have been achieved.

3.3 VCC Functions

Consider the choice $H(q) = \varphi$. Clearly $H(q_f) = 0$, meeting the first requirement of H for VCC. Additionally

$$\hat{A}(q) = A(q)|_{\mathcal{H}} = [0 \quad 0 \quad 1]$$

is integrable, so this choice of H is suitable. The function $F(q)$ must satisfy both $F(q_f) = 0$ and $\frac{\partial F}{\partial q} = \hat{A}(q)$. Both conditions are satisfied by the choice $F(q) = \eta$. Finally, the function $N(q)$ must be such that $N(q_f) = 0$ and $\frac{\partial N}{\partial q}$ must be linearly independent from $\frac{\partial H}{\partial q} = [0 \ 1 \ 0]$ and $\frac{\partial F}{\partial q} = [0 \ 0 \ 1]$. These conditions result in $N(q) = d - d_f$.

The strategy to drive $q \rightarrow q_f$ is to first drive $F = \eta$ close to zero, then drive $H = \varphi$ close to zero while simultaneously driving F closer to zero, then finally to drive $N = d - d_f$ towards zero while simultaneously driving H closer to zero. Recall that, because of the integrability of $\hat{A}(q)$, F should stay close to zero in the third step even though it is not explicitly being controlled. In the following sections, controllers to execute each of these steps are derived and then tied together using sequential composition to produce a solution to Problem 1. In the sequel, we will not use the names H , F , and N and will refer to them in terms of φ , η , and $d - d_f$ directly. We keep in mind though that the procedure that follows applies to more general VCC functions as well.

3.4 Controller Preliminaries

Some of the following controllers use an input transformation to simplify their expression. Specifically, a new input v is defined by

$$v = \begin{bmatrix} -\cos\varphi & 0 \\ \frac{\sin\varphi}{d+d_{\text{mid}}} & 1 \end{bmatrix} u \triangleq B_1 u. \quad (2)$$

Note that the matrix B_1 is invertible for $q \in \Omega_S$. In terms of v , the unicycle kinematics become

$$\begin{bmatrix} \dot{d} \\ \dot{\varphi} \\ \dot{\eta} \end{bmatrix} = \begin{bmatrix} 1 \\ 0 \\ \frac{\tan\varphi}{d+d_{\text{mid}}} \end{bmatrix} v_1 + \begin{bmatrix} 0 \\ 1 \\ 0 \end{bmatrix} v_2.$$

Hence the transformed inputs v_1 and v_2 directly control \dot{d} and $\dot{\varphi}$, respectively.

We also define a norm-like function M that will be used to measure the magnitudes of d and φ relative to the constraints D and φ_c :

$$M(q) = \frac{d^2}{D^2} + \frac{\varphi^2}{\varphi_c^2} + \epsilon_M \quad (3)$$

for some $\epsilon_M \in (0, \frac{1}{2})$. With this definition of M , we are guaranteed that $q \in \Omega_S$ whenever $M(q) < 1 + \epsilon_M$.

Finally, we define the sign function:

$$\text{sign}(x) = \begin{cases} -1 & \text{if } x < 0, \\ 0 & \text{if } x = 0, \\ 1 & \text{if } x > 0. \end{cases}$$

3.5 Palette of Controllers

The controllers that will be used to solve Problem 1 are presented here. The behaviors that they produce when applied to the unicycle kinematics of Equation 1 are stated and proven. In each of these declarations, the symbol t_0 represents the time at which that controller is invoked, and a subscript zero attached to a variable denotes the value of the variable at $t = t_0$, eg. $d_0 = d(t_0)$.

Controller $\Phi 1$ (drive φ away from zero) Suppose that $\varphi_0 = 0$. Then the control

$$u = \begin{bmatrix} 0 \\ 1 \end{bmatrix}$$

drives φ away from zero, i.e. for arbitrarily small time $t > 0$, $\varphi(t_0 + t) > 0$.

Controller $\Phi 2$ (drive $M \searrow 1$) Consider the control

$$u = \begin{bmatrix} 0 \\ -\text{sign}\varphi \end{bmatrix}.$$

If $q_0 = q(t_0) \in \Omega_S$ satisfies $\varphi_0 \neq 0$ and $M(q_0) > 1$, then there exists some time $T > t_0$ such that for any $\epsilon > 0$

1. $|M(q(T)) - 1| < \epsilon$
2. $q(t) \in \Omega_S$ for all $t \in [t_0, T]$.

Proof: Under this control, φ is strictly monotonically decreasing and is driven to zero in finite time T_1 and d and η do not change, which proves 2. Since $q(T_1) \in \Omega_S$, then $M(q(T_1)) < 1 + \epsilon_M$. If $M(q(T_1)) > 1 - \epsilon_M$, then $|M(q(T_1)) - 1| < \epsilon_M$ and $T = T_1$. Else $M(q(T_1)) < 1$, and since $M(q_0) > 1$ there must be some $T < T_1$ where 1. is satisfied.

Controller $\Phi 3$ (drive $M \nearrow 1$) Consider the control

$$u = B_1^{-1}v = B_1^{-1} \begin{bmatrix} -\text{sign}(\eta\varphi) \\ \text{sign}(\varphi) \end{bmatrix}.$$

If $M(q_0) < 1$ and $\varphi_0 \neq 0$, then there exists some time $T \geq t_0$ such that for any $\epsilon \in (0, \epsilon_M)$

1. $q(t) \in \Omega_S$ for all $t \in [t_0, T]$.
2. Either $|\eta(T)| < \epsilon$ or $|M(q(T)) - 1| < \epsilon$.

To prove 1., simply note that

$$\{q | M(q) < 1 + \epsilon_M\} \subset \Omega_S.$$

To prove 2., first note that

$$\dot{\varphi} = v_1 = \text{sign}\varphi,$$

which means that the magnitude of φ is monotonically increasing for $t > t_0$. Now

$$\dot{\eta} = \frac{\tan\varphi}{d + d_{\text{mid}}}v_2 = -\text{sign}(\eta\varphi)\frac{\tan\varphi}{d + d_{\text{mid}}}$$

The quantity $\text{sign}(\varphi)\tan\varphi$ is greater than zero and strictly increasing and $d + d_{\text{mid}} > 0$, so there exists some time $T_1 \geq t_0$ such that

$$|\eta(T_1)| < \epsilon.$$

To evaluate the second possibility in 2. it is only necessary to consider $t \in [t_0, T_1]$ since $t > T_1$ implies that the first possibility has already happened. Consider

$$\begin{aligned} \dot{M} &= \begin{bmatrix} \frac{\partial M}{\partial d} & \frac{\partial M}{\partial \varphi} \end{bmatrix} \begin{bmatrix} \dot{d} \\ \dot{\varphi} \end{bmatrix} \\ &= \begin{bmatrix} \frac{2d}{D^2} & \frac{2\varphi}{\varphi_c^2} \end{bmatrix} v \\ &= \frac{-2d\text{sign}(\eta\varphi)}{D^2} + \frac{2\varphi\text{sign}(\varphi)}{\varphi_c^2}. \end{aligned} \tag{4}$$

The second term on the right hand side is always positive and growing. To analyze the first term, first note that for $t \in [t_0, T_1]$, the signs of η and φ do not change. Since d is changing in the direction given by $-\text{sign}(\eta\varphi)$, d is always moving in the same direction. Assume for now that T_1 is sufficiently large so that the sign of has time to become the same as $-\text{sign}(\eta\varphi)$, and denote the time at which this happens by T_2 . For $t > T_2$, the first term becomes positive and

growing, guaranteeing the \dot{M} is positive and growing. Again assuming that T_1 is sufficiently large, there is then some time $T_3 \geq t_0$ such that $M(q(T_3)) > 1 - \epsilon$. Since $T_3 < T_1$, then $T = T_3$ and $|M(q(T)) - 1| < \epsilon$. By the definition of T_3 , $M(q(t)) < 1 + \epsilon_M$ for all $t \in [t_0, T]$, which implies that $q(t) \in \Omega_S$ for all such t .

If the assumption that T_1 is sufficiently large for M to reach one fails, then $T = T_1$, and $M = 1 - \epsilon$ was never reached, which implies $q(t) \in \Omega_S$ for all $t \in [t_0, T]$.

Controller $\Phi 4$ (drive $\eta \rightarrow 0, M \rightarrow 1$) Consider the control

$$u = B_1^{-1} \begin{bmatrix} -\lambda(M-1) & -\text{sign}(\eta) \frac{D}{\varphi_c} \\ \text{sign}(\eta) \frac{\varphi_c}{D} & -\lambda(M-1) \end{bmatrix} \begin{bmatrix} d \\ \varphi \end{bmatrix}.$$

There exists $\epsilon \in (0, \epsilon_M)$ such that if $|M(q_0) - 1| < \epsilon$, then

1. There exists $T > t_0$ such that $|\eta(T)| < \delta$ for any $\delta > 0$.
2. $q(t) \in \Omega_S$ for all $t \in [t_0, T]$.

To begin the proof, note that substituting u into the unicycle kinematics yields a harmonic oscillator equation with a nonlinear damping term:

$$\begin{bmatrix} \dot{d} \\ \dot{\varphi} \end{bmatrix} = \begin{bmatrix} -\lambda(M-1) & -\text{sign}(\eta) \frac{D}{\varphi_c} \\ \text{sign}(\eta) \frac{\varphi_c}{D} & -\lambda(M-1) \end{bmatrix} \begin{bmatrix} d \\ \varphi \end{bmatrix}, \quad (5)$$

Define an error term $e = M - 1$. Then

$$\begin{aligned} \dot{e} &= \dot{M} \\ &= \begin{bmatrix} \frac{\partial M}{\partial d} & \frac{\partial M}{\partial \varphi} \end{bmatrix} \begin{bmatrix} \dot{d} \\ \dot{\varphi} \end{bmatrix} \\ &= \begin{bmatrix} \varphi_c^2 d & D^2 \varphi \end{bmatrix} \begin{bmatrix} -(M-1) & -\text{sign}(F) D^2 \\ \text{sign}(F) \varphi_c^2 & -(M-1) \end{bmatrix} \begin{bmatrix} d \\ \varphi \end{bmatrix} \\ &= -(M-1) \left(\frac{2d^2}{D^2} + \frac{2\varphi^2}{\varphi_c^2} \right) \\ &= -e(M - \epsilon_M). \end{aligned}$$

Note that

$$M - \epsilon_M = (M - 1) + 1 - \epsilon_M > 1 - 2\epsilon_M,$$

so $M - \epsilon_M$ is strictly positive provided $\epsilon_M < \frac{1}{2}$. As a result, e converges exponentially to zero, which means that M converges exponentially to one. 2. is then proven by noting that

$$\{q | M(q) < 1 + \epsilon\} \subset \Omega_S,$$

a fact that can be proven easily by contradiction.

Now consider the oscillator phase

$$\theta(t) = \text{atan} \left(\frac{\varphi(t)}{d(t)} \right).$$

Differentiating yields

$$\begin{aligned} \dot{\theta} &= \left(\frac{\dot{\varphi}}{d} - \frac{d\dot{\varphi}}{d^2} \right) \left(1 + \frac{\varphi^2}{d^2} \right)^{-1} \\ &= \text{sign}(\eta) \frac{d^2 \varphi_c^2 + \varphi^2 D^2}{D \varphi_c (d^2 + \varphi^2)}. \end{aligned} \quad (6)$$

If we let $R = \sqrt{\varphi^2 + d^2}$, so that (R, θ) are polar coordinates in the $d - \varphi$ plane, then we can write

$$\begin{aligned}\dot{\theta} &= \frac{\frac{R^2 \cos^2 \theta}{D^2} + \frac{R^2 \sin^2 \theta}{\varphi_c^2}}{R^2} \\ &= \frac{\cos^2 \theta}{D^2} + \frac{\sin^2 \theta}{\varphi_c^2},\end{aligned}$$

so $\dot{\theta}$ is a function of θ alone. Now consider an undamped linear harmonic oscillator

$$\begin{bmatrix} \dot{d}_u \\ \dot{\varphi}_u \end{bmatrix} = \begin{bmatrix} 0 & -\text{sign}(\eta) \frac{D}{\varphi_c} \\ \text{sign}(\eta) \frac{\varphi_c}{D} & 0 \end{bmatrix} \begin{bmatrix} d_u \\ \varphi_u \end{bmatrix}. \quad (7)$$

Following the process used for the damped oscillator, it is not difficult to verify that $\dot{\theta}_u = \dot{\theta}$. So if we integrate the damped and undamped oscillators starting from initial conditions $[d_0 \ \varphi_0]^T$ and $[d_{u0} \ \varphi_{u0}]^T$, respectively, that are in phase, then their solutions $[d(t) \ \varphi(t)]^T$ and $[d_u(t) \ \varphi_u(t)]^T$ remain in phase for all time. This idea, together with the well know fact that the solution to the undamped oscillator Equation 7 starting from initial condition $[d_0 \ 0]^T$ is

$$\begin{aligned}d_u(t) &= d_0 \cos(t - t_0) \\ \varphi_u(t) &= \frac{d_0 \varphi_c}{D} \sin(t - t_0),\end{aligned}$$

yields the following useful facts about the solution to the damped oscillator $\varphi(t)$:

- there is a time τ_0 such that $\varphi(\tau_0) = \varphi(\tau_0 + \pi) = 0$
- under the assumption $|M - 1| < \epsilon < \epsilon_M$,

$$\frac{(D - 2\epsilon)\varphi_c}{D} |\sin(t - \tau_0)| \leq |\varphi(t)| \leq \varphi_c |\sin(t - \tau_0)|.$$

Now to prove 1. compute

$$\begin{aligned}\frac{d}{dt} |\eta| &= \dot{\eta} \text{sign}(\eta) \\ &= -\frac{\tan \varphi}{d + d_{\text{mid}}} \lambda (M - 1) d \text{sign}(\eta) - \frac{D}{\varphi_c} \frac{\varphi \tan \varphi}{d + d_{\text{mid}}} \\ &< D \left(\frac{\lambda \epsilon}{d_{\text{min}}} - \frac{\varphi \tan \varphi}{d_{\text{max}}} \right).\end{aligned}$$

Now

$$\begin{aligned}|\eta(\tau_0 + \pi)| &= |\eta(\tau_0)| + \int_{\tau_0}^{\tau_0 + \pi} \frac{d}{d\sigma} |\eta(\sigma)| d\sigma \\ &< |\eta(\tau_0)| + \int_{\tau_0}^{\tau_0 + \pi} D \left(\frac{\lambda \epsilon}{d_{\text{min}}} - \frac{\varphi(\sigma) \tan \varphi(\sigma)}{d_{\text{max}}} \right) d\sigma \\ &= |\eta(\tau_0)| + \frac{D \lambda \epsilon}{d_{\text{min}}} \pi - \frac{1}{d_{\text{max}}} \int_{\tau_0}^{\tau_0 + \pi} \varphi(\sigma) \tan \varphi(\sigma) d\sigma. \\ &< |\eta(\tau_0)| + \frac{D \lambda \epsilon}{d_{\text{min}}} \pi - \frac{(D - 2\epsilon)\varphi_c}{D d_{\text{max}}} \int_0^\pi \sin(\sigma) \tan \left(\frac{(D - 2\epsilon)\varphi_c}{D d_{\text{max}}} \sin(\sigma) \right) d\sigma. \quad (8)\end{aligned}$$

claim: $\gamma \sin^2(t) < \sin(t) \tan(\gamma \sin(t))$

To prove this draw a right triangle with angle $\gamma \sin t$, adjacent 1, opposite $\tan(\gamma \sin t)$, and hypoteneuse $\sqrt{1 + \tan^2(\gamma \sin t)}$.

The result follows from the definition of \sin . As a result of this claim, we have

$$\begin{aligned}
|\eta(\tau_1)| &< |\eta(\tau_0)| + \frac{D\lambda\epsilon}{d_{\min}}\pi - \frac{(D-2\epsilon)^2\varphi_c^2}{D^2d_{\max}} \int_0^\pi \sin^2(\sigma)d\sigma \\
&= |\eta(\tau_0)| + \frac{D\lambda\epsilon}{d_{\min}}\pi - \frac{(D-2\epsilon)^2\varphi_c^2}{D^2d_{\max}} \frac{\pi}{2} \\
&= |\eta(\tau_0)| + \left(\frac{D\lambda\epsilon}{d_{\min}} - \frac{(D-2\epsilon)^2\varphi_c^2}{2D^2d_{\max}} \right) \pi.
\end{aligned} \tag{9}$$

We can use the quadratic formula to show that $F = \eta$ moves strictly towards zero every π seconds so long as

$$\epsilon < \Gamma + D - \sqrt{\Gamma^2 + 2\Gamma D},$$

where

$$\Gamma = \frac{D^3\lambda d_{\max}}{2d_{\min}\varphi_c}.$$

Controller $\Phi 5$ (drive $\varphi \rightarrow \gamma\eta\text{sign}(d)$) Let $\gamma > 0$ and consider the control law

$$u = \begin{bmatrix} 0 \\ -(\varphi - \gamma\eta\text{sign}(d)) \end{bmatrix}.$$

For any $\epsilon > 0$ and $q_0 \in \Omega_S$, there exists $T > t_0$ such that

1. $|\varphi(T) - \gamma\eta(T)\text{sign}(d(T))| < \epsilon$.
2. If $q(T) \in \Omega_S$, then $q(t) \in \Omega_S$ for all $t \in [t_0, T]$.

Futher, if $\epsilon < \gamma F(q_0)$, then the final state $q(T)$ will be such that any control that pushes η towards zero also pushes d towards zero.

Under u , d and η are constant.

$$\dot{\varphi} = u_2 = -(\varphi - \gamma\eta\text{sign}(d)),$$

so φ converges exponentially to $\gamma\eta\text{sign}(d)$. 1. and 2. then follow directly from the definitions of η and φ .

To reduce η , it is necessary that $\text{sign}(\dot{\eta}) = -\text{sign}(\eta)$. Substituting for $\dot{\eta}$ from the unicycle kinematics yields

$$\begin{aligned}
\text{sign}(\dot{\eta}) &= \text{sign}\left(\frac{-\sin\varphi}{d + d_{\text{mid}}}u_1\right) \\
&= -\text{sign}(\varphi u_1).
\end{aligned}$$

So in order to make $\text{sign}(\dot{\eta}) = -\text{sign}(\eta)$, it is required that $\text{sign}(u_1) = \text{sign}(\varphi\eta)$. Now

$$\begin{aligned}
\text{sign}(\dot{d}) &= \text{sign}(-\cos\varphi u_1) \\
&= -\text{sign}(\varphi\eta).
\end{aligned}$$

So in order to have $\text{sign}(\dot{d}) = -\text{sign}(d)$, it is sufficient that $-\text{sign}(\varphi\eta) = -\text{sign}(d)$, or $\text{sign}(\varphi\eta d) = 1$, which is equivalent to saying that φ and ηd have the same signs. This is guaranteed by the condition on $\epsilon < \gamma\eta_0$.

Controller $\Phi 6$ (drive $\varphi \rightarrow 0, \eta \rightarrow 0$) Assume that $\text{sign}(\varphi_0) = \text{sign}(\eta_0 d_0)$, $|\varphi_0| < \frac{1}{2}$, and $|\eta_0| < \frac{1}{2}$. Consider the control

$$u = \begin{bmatrix} \frac{\lambda_2\eta(d+d_{\text{mid}})}{\sin\varphi} \\ \lambda_1\varphi - \lambda_2\eta \end{bmatrix},$$

where $\lambda_1 > 1.7\lambda_2 > 0$.

There exists a time $t_{\text{safe}}(q_0) > t_0$ such that

1. for any $\epsilon_1 > 0$ and $\epsilon_2 > 0$, there exist $\delta_1 > 0$ and $\delta_2 > 0$ such that if $|\varphi_0| < \delta_1$ and $|\eta_0| < \delta_2$ then $|\varphi(t_{\text{safe}})| < \epsilon_1$ and $|\eta(t_{\text{safe}})| < \epsilon_2$.
2. if $q_0 \in \Omega_S$, then $q(t) \in \Omega_S$ for all $t \in [t_0, t_{\text{safe}}(q_0)]$.

Substituting u into the kinematics reveals that this control law is just the controller for the first stage of VCC, meaning

$$\begin{bmatrix} \dot{\eta} \\ \dot{\varphi} \end{bmatrix} = \begin{bmatrix} -\lambda_1 \eta \\ -\lambda_2 \varphi \end{bmatrix}.$$

So η and φ converge to zero exponentially. The proof can then be finished by showing that there is a strictly positive t_{safe} . Under u ,

$$\dot{d} = -\cos\varphi \frac{\lambda_2 \eta (d + d_{\text{mid}})}{\sin\varphi}.$$

The signs of φ and η do not change, and $d + d_{\text{mid}} > d_{\text{min}} > 0$, so d moves in the same direction for the duration of the time that u is applied. The assumption $\text{sign}(\varphi) = \text{sign}(\eta d)$ together with the fact that $\frac{d}{dt} |\eta|_{t=t_0} < 0$ implies that d is moving toward zero at time t_0 , so d can change by D and still have $q \in \Omega_S$. If we can find an upper bound $|\dot{d}|_{\text{max}} > |\dot{d}|$, then we are guaranteed to stay in Ω_S until at least $t_{\text{safe}} = D/|\dot{d}|_{\text{max}}$.

Claim: $\frac{d}{dt} |\dot{d}| < 0$.

Proof of claim:

$$\begin{aligned} \frac{d}{dt} |\dot{d}| &= \ddot{d} \text{sign}(\dot{d}) \\ &= -\lambda_2 \left(\dot{\eta} \frac{(d + d_{\text{mid}})}{\tan\varphi} + \dot{d} \frac{\eta}{\tan\varphi} - \dot{\varphi} \frac{\eta (d + d_{\text{mid}})}{\sin^2\varphi} \right) \text{sign}(-\varphi\eta) \\ &= -\lambda_2 \left(-\lambda_1 \eta \frac{(d + d_{\text{mid}})}{\tan\varphi} - \frac{\lambda_2 \eta (d + d_{\text{mid}})}{\tan\varphi} \frac{\eta}{\tan\varphi} + \lambda_1 \varphi \frac{\eta (d - d_{\text{mid}})}{\sin^2\varphi} \right) \text{sign}(-\varphi\eta) \\ &= -\lambda_1 \lambda_2 (d + d_{\text{mid}}) \left| \frac{\eta}{\tan\varphi} \right| - \lambda_2^2 \eta^2 (d + d_{\text{mid}}) \frac{\text{sign}\eta}{|\tan\varphi|} + \lambda_2^2 (d + d_{\text{mid}}) \left| \frac{\varphi\eta}{\sin^2\varphi} \right| \\ &\leq -\lambda_1 \lambda_2 (d + d_{\text{mid}}) \left| \frac{\eta}{\tan\varphi} \right| + \lambda_2^2 \eta^2 (d + d_{\text{mid}}) \frac{\text{sign}\eta}{|\tan\varphi|} + \lambda_2^2 (d + d_{\text{mid}}) \left| \frac{\varphi\eta}{\sin^2\varphi} \right| \\ &= \frac{\lambda_2 (d + d_{\text{mid}})}{|\tan\varphi|} \left(-\lambda_1 |\eta| + \lambda_2 \left(\eta^2 + \left| \frac{\varphi\eta}{\sin\varphi \cos\varphi} \right| \right) \right). \end{aligned} \tag{10}$$

It can be verified that $\left| \frac{\varphi}{\sin\varphi \cos\varphi} \right| < 1.2$ if $\varphi < \frac{1}{2}$, so in order to show the claim we must show that

$$\lambda_1 |\eta| > \lambda_2 (\eta^2 + 1.2 |\eta|),$$

which is implied by

$$1.7\lambda_2 |\eta| > \lambda_2 (\eta^2 + 1.2 |\eta|),$$

which is equivalent to

$$0.5 |\eta| > \eta^2,$$

which is true for $\eta < 0.5$, and the claim is proven.

Since $\frac{d}{dt} |\dot{d}| < 0$, the maximum $|\dot{d}|$ occurs at time t_0 , so

$$|\dot{d}|_{\text{max}} = \frac{\lambda_2 \eta_0 (d_0 + d_{\text{mid}})}{\tan\varphi_0}$$

and

$$t_{\text{safe}} = \frac{D \tan\varphi_0}{\lambda_2 \eta_0 (d_0 + d_{\text{mid}})}.$$

Controller Φ_7 (drive $\varphi \rightarrow 0, d \rightarrow d_f$) Consider the controller

$$u = \begin{bmatrix} \text{sign}(d - d_f) \\ - \left| \frac{\cos \varphi}{\lambda_3(d - d_f)} \right| \left(\lambda_2 \varphi + \lambda_3 \frac{(d - d_f) \tan \varphi}{d + d_{\text{mid}}} \right) \end{bmatrix},$$

where $\lambda_2, \lambda_3 > 0$. If

$$|\eta_0| + \frac{1}{d_{\text{min}}} |(d_0 - d_f) \tan \varphi_0| < \epsilon$$

then

1. there exists $T > t_0$ such that $\|q(T) - q_f\|_\infty < \epsilon$ for all $t > T$.
2. $q(t) \in \Omega_S$ for all $t > 0$.

Substituting u into the unicycle kinematics reveals and

$$\begin{bmatrix} \dot{\eta} \\ \dot{\varphi} \end{bmatrix} = \begin{bmatrix} -\cos(\varphi) \text{sign}(d - d_f) \\ -\lambda_2 \left| \frac{\cos(\varphi)}{\lambda_3(d - d_f)} \right| \varphi \end{bmatrix}.$$

The derivatives of the absolute values are then

$$\begin{aligned} \frac{d}{dt} |d - d_f| &= \dot{d} \text{sign}(d - d_f) = -\cos(\varphi) \\ \frac{d}{dt} |\varphi| &= \dot{\varphi} \text{sign}(\varphi) = -\frac{\lambda_2}{\lambda_3} \left| \frac{\varphi \cos(\varphi)}{d - d_f} \right|, \end{aligned}$$

For $\varphi < \varphi_c < \pi/4$ we have $\cos(\varphi) > 0$, so both $d - d_f$ and φ go monotonically toward zero with t . Because of this, we can write

$$\frac{d}{dt} |d - d_f| < -\cos(\varphi_0)$$

and conclude that the time it takes d to reach d_f is at most

$$T(\varphi_0, d_0) = \left| \frac{d_0 - d_f}{\cos(\varphi_0)} \right| + t_0.$$

Since $d_0 - d_f$ appears in the denominator of $\frac{d}{dt} |\varphi|$, we are also guaranteed that φ reaches zero before $T(\varphi_0, d_0)$.

All that is left then is to show that $|\eta(t)| < \epsilon$ for all $t \leq T(\varphi_0, d_0)$. Substituting u into the unicycle kinematics yields

$$|\dot{\eta}| = \left| \frac{-\sin \varphi}{d + d_{\text{mid}}} \text{sign}(d - d_f) \right| < \left| \frac{\sin(\varphi_0)}{d_{\text{min}}} \right|.$$

This combined with the fact that $\dot{\eta} = 0$ when $\varphi = 0$ yields the required bound on $F(t)$ for all $t > t_0$:

$$\begin{aligned} |\eta(t)| &< |\eta_0| + \left| \frac{\sin(\varphi_0)}{d_{\text{min}}} \right| (T(\varphi_0, d_0) - t_0) \\ &= |\eta_0| + \frac{1}{d_{\text{min}}} |(d_0 - d_f) \tan(\varphi_0)| \\ &< \epsilon. \end{aligned}$$

3.6 Sequential Composition

The controllers presented in the previous section can be used to realize the VCC inspired behavior by sequentially driving η , φ , and $(d - d_f)$ to zero. Specifically controllers Φ_1 , Φ_2 , Φ_3 , and Φ_4 can be applied to make η arbitrarily small, then Φ_5 and Φ_6 can be applied to make φ arbitrarily small while holding η small, and finally Φ_7 brings $(d - d_f)$ to zero while explicitly keeping φ small and implicitly keeping η small. In this section, we use sequential composition to carefully connect these controllers to create a global feedback controller that provably solves Problem 1.

Proposition 1 Let $q_f \in \Omega_S$ and $\epsilon > 0$. Define $\mathcal{I} = \{1, 2, \dots, 7\}$. Consider the palette of controllers $\mathcal{U} = \{\Phi_i | i \in \mathcal{I}\}$ defined in Section 3.5 together with goal sets $\mathcal{G}(\Phi_i)$ and domains $\mathcal{D}(\Phi_i)$ defined from highest to lowest priority:

$$\begin{aligned}\mathcal{G}(\Phi_7) &= \{q \in \Omega_S | |q - q_f| < \epsilon\} \\ \mathcal{D}(\Phi_7) &= \left\{q \in \Omega_S | |\eta| + \frac{1}{d_{\min}} |(d - d_f)\tan\varphi| < \epsilon\right\}\end{aligned}\quad (11)$$

$$\begin{aligned}\mathcal{G}(\Phi_6) &= \left\{q \in \Omega_S | |\eta| < \frac{\epsilon}{2}, \quad |\varphi| < \frac{d_{\min}\pi\epsilon}{16D}\right\} \\ \mathcal{D}(\Phi_6) &= \left\{q \in \Omega_S | |\eta| < f(\epsilon), \quad |\varphi| < h(\epsilon), \quad |\eta| < 2|\varphi|\right\}\end{aligned}\quad (12)$$

$$\begin{aligned}\mathcal{G}(\Phi_5) &= \left\{q \in \Omega_S | |\eta| < \min\left(f(\epsilon), \frac{2}{3}h(\epsilon)\right), \quad |\varphi - \eta| < \frac{|\eta|}{2}\right\} \\ \mathcal{D}(\Phi_5) &= \left\{q \in \Omega_S | |\eta| < \min\left(f(\epsilon), \frac{2}{3}h(\epsilon)\right)\right\}\end{aligned}\quad (13)$$

$$\begin{aligned}\mathcal{G}(\Phi_4) &= \mathcal{D}(\Phi_5) \\ \mathcal{D}(\Phi_4) &= \left\{q \in \Omega_S | |M(q) - 1| < \min(\epsilon_M, \delta_1), \quad \varphi \neq 0\right\}\end{aligned}\quad (14)$$

$$\begin{aligned}\mathcal{G}(\Phi_3) &= \mathcal{D}(\Phi_4) \cup \mathcal{D}(\Phi_5) \\ \mathcal{D}(\Phi_3) &= \left\{q \in \Omega_S | M(q) < 1, \quad \varphi \neq 0\right\}\end{aligned}\quad (15)$$

$$\begin{aligned}\mathcal{G}(\Phi_2) &= \mathcal{D}(\Phi_4) \cup \mathcal{D}(\Phi_5) \\ \mathcal{D}(\Phi_2) &= \left\{q \in \Omega_S | M(q) > 1, \quad \varphi \neq 0\right\}\end{aligned}\quad (16)$$

$$\begin{aligned}\mathcal{G}(\Phi_1) &= \mathcal{D}(\Phi_2) \cup \mathcal{D}(\Phi_3) \\ \mathcal{D}(\Phi_1) &= \Omega_S\end{aligned}\quad (17)$$

where $f(\epsilon) = \frac{\epsilon}{2}e^{\frac{\lambda_1 D}{2\lambda_2 d_{\max}}}$ and $h(\epsilon) = \frac{d_{\min}\pi\epsilon}{16D}e^{\frac{D}{2d_{\max}}}$. Under this composition, there exists a $T > t_0$ such that the solution $q(t)$ to the unicycle kinematics in Equation 1 for any initial condition $q(t_0) \in \Omega_S$ satisfies $\|q(t) - q_f\|_\infty < \epsilon$ for all $t > T$.

To prove this proposition, it suffices to show

1. The collection of domains cover the workspace, i.e. $\bigcup_{i \in \mathcal{I}} \mathcal{D}(\Phi_i) = \Omega_S$.
2. The goal set of the highest level controller is a subset of the overall goal set.
3. The domain of the highest level controller must be positive invariant.

and for each $i \in \mathcal{I}$

4. Φ_i reaches its goal set, i.e. for any initial condition $q_0 \in \mathcal{D}(\Phi_i)$, there exists a finite $T > t_0$ such that the solution $q(t)$ to the unicycle kinematics under the controller Φ_i satisfies $q(T) \in \mathcal{G}(\Phi_i)$.

and for each $i \in \mathcal{I}, i < 7$

5. Φ_i prepares a set of higher priority controllers, i.e. $\mathcal{G}(\Phi_i) \subset \bigcup_{j > i} \mathcal{D}(\Phi_j)$.
6. $\mathcal{D}(\Phi_i)$ is conditionally invariant under Φ_i , i.e. if $q_0 \in \mathcal{D}(\Phi_i)$ and $T > t_0$ is the smallest time such that $q(T) \in \mathcal{G}(\Phi_i)$, then $q(t) \in \mathcal{D}(\Phi_i)$ for all $t \in [0, T]$.

Condition 1 is trivially true because $\mathcal{D}(\Phi_1) = \Omega_S$. Condition 2 is true because the goal set is identical to $\mathcal{D}(\Phi_7)$. For Controllers Φ_1, Φ_2, Φ_3 , and Φ_4 , conditions 4, 5, and 6 follow directly from the properties the controllers and the definitions of the domains. We proceed to check conditions 4, 5, and 6 for each remaining controller.

Controller Φ_5 : Conditions 4 and 6 follow from the properties of Controller Φ_5 . To show condition 5, note that $|\varphi - \eta| < \frac{|\eta|}{2} \Leftrightarrow |\eta| < 2|\varphi| < 3|\eta|$. Also, $|\eta| < 2h(\epsilon) \Rightarrow |\varphi| < h(\epsilon)$. As a result, $\mathcal{G}(\Phi_5) \subset \mathcal{D}(\Phi_6)$.

Controller Φ_6 : Condition 6 follows from the properties of Controller Φ_6 . To show condition 5 we use the fact that $|\tan\varphi| < \frac{4|\varphi|}{\pi}$ for $|\varphi| < \varphi_c < \frac{\pi}{4}$ and the fact that $|d - d_f| < 2D$ to compute

$$|\eta| + \frac{1}{d_{\min}} |(d - d_f)\tan\varphi| < |\eta| + \frac{8D}{\pi d_{\min}} |\varphi| < \epsilon,$$

so $\mathcal{G}(\Phi_6) \subset \mathcal{D}(\Phi_7)$.

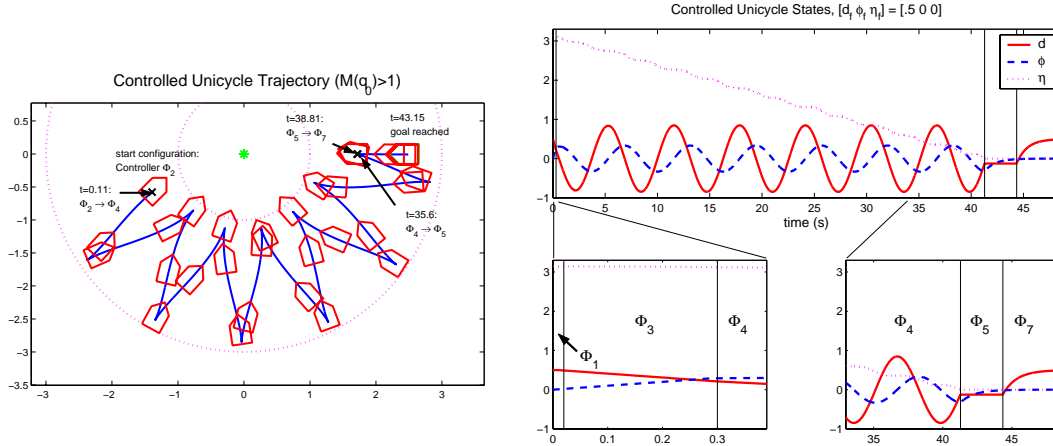


Figure 2: This figure depicts the simulated unicycle trajectory for an initial condition q_0 such that $M(q_0) > 1$. The sequential composition starts in Controller Φ_2 , then proceeds through Φ_4 , Φ_5 , and Φ_7 before achieving the goal condition of $\|q(t) - q_f\| < \epsilon$.

To show condition 4, first compute

$$t_{\text{safe}}(q_0) = \left| \frac{D \tan \varphi_0}{\lambda_2 \eta_0 (d_0 + d_{\text{mid}})} \right| > \left| \frac{D \varphi_0}{\lambda_2 \eta_0 d_{\text{min}}} \right| > \left| \frac{D(\eta_0/2)}{\lambda_2 \eta_0 d_{\text{min}}} \right| = \frac{D}{2\lambda_2 d_{\text{min}}}.$$

So $|\eta_0| < \frac{\epsilon}{2} e^{\frac{\lambda_1 D}{2\lambda_2 d_{\text{min}}}} \Rightarrow |\eta(t_{\text{safe}})| < |\eta_0| e^{-\frac{\lambda_1 D}{2\lambda_2 d_{\text{min}}}} < \frac{\epsilon}{2}$ and $|\varphi_0| < \frac{d_{\text{min}} \pi \epsilon}{16D} e^{\frac{D}{2d_{\text{min}}}} \Rightarrow |\varphi(t_{\text{safe}})| < |\eta_0| e^{\frac{D}{2d_{\text{min}}}} < \frac{d_{\text{min}} \pi \epsilon}{16D}$.

Controller Φ_7 : Conditions 5 and 6 do not apply since Φ_7 is the highest priority controller. Condition 4 follows directly from the properties of Controller Φ_7 . Condition 3 is all that is left to show. Define

$$Z(q) \triangleq |\eta| + \frac{1}{d_{\text{min}}} |(d - d_f) \tan \varphi|.$$

and compute

$$\begin{aligned} \dot{Z} &= \frac{d}{dt} |\eta| + \frac{1}{d_{\text{min}}} \left(\frac{d}{dt} |\tan \varphi| |d - d_f| + |\tan \varphi| \frac{d}{dt} |d - d_f| \right) \\ &= -\frac{\sin \varphi}{d + d_{\text{mid}}} \text{sign}((d - d_f) \eta) + \frac{1}{d_{\text{min}}} \left(\frac{d}{dt} |\tan \varphi| |d - d_f| - |\sin \varphi| \right) < 0. \end{aligned}$$

As result, if $Z(q_0) < \epsilon$, then $Z(q(t)) < \epsilon$ for all time $t > t_0$ and the set $\mathcal{D}(\Phi_7)$ is invariant.

4 Results

The sequential composition stated in Proposition 1 was verified experimentation as well as numerical simulation using MATLAB.

4.1 Simulation Results

The results of a sample simulation run are depicted in Figure 2. Note that the solution proceeds through a sequence of controllers with monotonically increasing priority until successfully reaching a state in the goal set. The example also exhibits opportunistic switching by skipping directly from Φ_5 to Φ_7 when the state enters the $\mathcal{D}(\Phi_7)$.




Figure 3: The CMU Deminer robot used in the experiment (left) and cube stack used to measure robot pose (right).

4.2 Experimental Results

The controller presented here was tested experimentally using the CMU Deminer (Figure 3), a differentially steered wheeled mobile robot that is modelled quite well by the unicycle kinematics of Equation 1. The robot was equipped with a forward facing color camera. An engineered landmark was constructed using two of the cubes described in Section 3.1. The cubes were stacked vertically as shown in Figure 3. In order to determine the pose relative to a cube, three faces of the cube must be visible. Due to perspective issues, the workspace around a single cube will contain many “blind spots” where only two faces are visible. To deal with this, the top cube was rotated relative to the bottom cube (about the vertical axis) so that blind spots of the two cubes do not overlap. The camera–cube combination was able to provide the full $SE(2)$ pose of the robot in polar coordinates at a rate of about 8 Hz. The resulting experimental set-up was very similar to the schematic shown in Figure 1.

In order to better handle sensor noise, the sequential composition stated in Proposition 1 was modified slightly by adding hysteresis. At each time step, the domain of the controller that was active in the previous step is enlarged slightly. This enlargement is accomplished practically by adding positive constants to larger side of the test inequality, for example $|\eta| < f(\epsilon)$ would be replaced by $|\eta| < f(\epsilon) + C$ for some $C > 0$. Admittedly, this violates our claim that the controller has no added state since the sequential composition must remember which controller was invoked last. However, the amount of state added is small and the performance gains in the presence of sensor noise are worth the compromise.

The experiments were run with a goal state $q_f = [0\ 0\ 0]^T$. The tolerance used was $\epsilon = 0.06$, which turns out to be around the smallest value that exhibited reliable switching behavior given the amount of noise present in the measurements. The bounds on the workspace were chosen to be $d_{\min} = 2$ meters, $d_{\max} = 3$ meters, resulting in $D = 0.5$ meters. The field of view constraint was chosen to be $\varphi_c = 0.25$ radians. The starting position for the robot was chosen to have η_0 at about 2 radians while d_0 and φ_0 were chosen somewhat randomly. From the starting position, the sequential composition was started and allowed to run until the goal was reached or until the camera lost sight of the landmark. The experiment was repeated 10 times, and the goal was successfully reached in 9 of the experiments. The single observed failure occurred when the vision software failed to track the cube for an extended period of time. Data collected from one of the successful runs is shown in Figure 4. Note that the controllers invoked do not proceed monotonically due to sensor noise. Still, the goal is reached using a reasonably small number of switches.

5 Conclusion

In this paper we have extended the idea of sequential composition so that it can be applied to the control of underactuated systems. Starting with the notion of variable constraint control, we identified a sequence of desirable behaviors, found closed loop controllers to implement them, and finally tied them together to create a globally convergent feedback controller for vision based unicycle navigation.

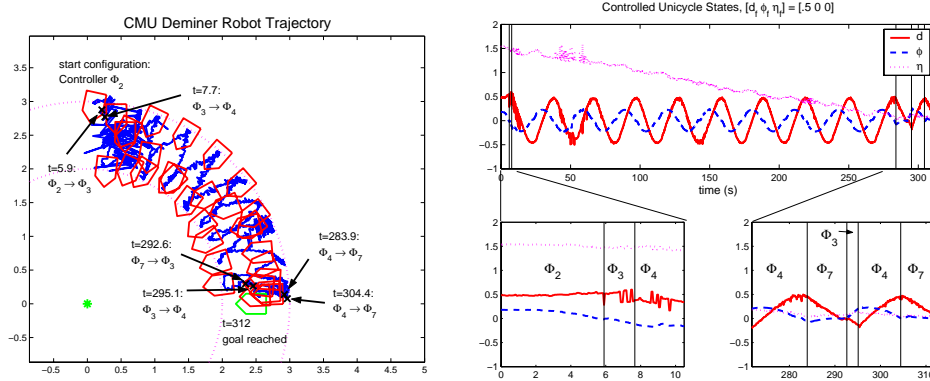


Figure 4: This figure depicts the trajectory of the CMU Deminer robot under sequential composition control. The figure on the left shows the trajectory of the robot in the plane, while the figure on the right shows the evolution of d , φ , and η . The sequential composition starts in Controller Φ_2 , then proceeds through Φ_3 , Φ_4 , Φ_7 , back to Φ_3 , then to Φ_4 and Φ_7 before achieving the goal condition of $\|q(t) - q_f\| < \epsilon = 0.06$.

Future work will proceed in two directions. First, we would like to advance the theoretical ideas presented in this paper so that they can be applied to more general underactuated systems. Our main efforts here will focus on a search for an automated method of generating the subtasks (sequence of manifolds) that can be used to identify the necessary palette of controllers. We will also seek more systematic methods of deriving the individual controllers in the palette. Some aspects of the controllers in this paper exploited special properties of the unicycle, for example its ability to turn in place. However, some of the properties we used are more general than they appear. For example, for a three dimensional system, it will always be possible to hold H constant and vary F , or hold F constant and vary H , so long as $H \neq 0$. We will seek to identify and exploit these types of properties to produce a more general result.

Acknowledgements

We'd like to thank Sarjoun Skaff, Howie Choset, and Al Costa for their help assistance in using the CMU Deminer robot. Thanks also go Noah Cowan providing feedback on an early draft of this paper. This work was supported financially by DARPA/ONR Grant N00014-98-1-0747.

References

- [1] M. Aicardi, G. Casalino, A. Bicchi, and A. Balestrino. Closed loop steering of unicycle-like vehicles via lyapunov techniques. *IEEE Robotics and Automation Magazine*, 2(1):27–35, 1995.
- [2] R.W. Brockett. Asymptotic stability and feedback stabilization. In R.S. Millman, R.W. Brockett, and H.J. Sussmann, editors, *Differential Geometric Control Theory*. Birkhauser, 1983.
- [3] R. R. Burridge, A. A. Rizzi, and D. E. Koditschek. Sequential composition of dynamically dexterous robot behaviors. *International Journal of Robotics Research*, 18(6):534–555, June 1999.
- [4] N. J. Cowan and D. E. Koditschek. Planar image based visual servoing as a navigation problem. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 611–617, Detroit, MI, May 1999.
- [5] N.J. Cowan, J.D. Weingarten, and D.E. Koditschek. Visual servoing via navigation functions. *IEEE Transactions on Robotics and Automation*, 18(4):521–533, August 2002.

- [6] C.C. de Wit and O.J. Sordalen. Exponential stabilization of mobile robots with nonholonomic constraints. *IEEE Transactions on Automatic Control*, 37(11):1791–1797, 1992.
- [7] K. Hashimoto and T. Noritsugu. Visual servoing of nonholonomic cart. In *Proc. IEEE Int'l. Conf. on Robotics and Automation*, pages 1719–24, Albuquerque, NM, April 1997.
- [8] T. Ikeda, M. Fukaya, and T. Mita. Position and attitude control of an underwater vehicle using variable constraint control. In *Proceedings of the 40th IEEE Conference on Decision and Control*, pages 3758–3763, Orlando, FL, December 2001.
- [9] T. Ikeda, T. Mita, and B.D.O. Anderson. Variable constraint control for a class of nonholonomic systems. available at <http://mita-lab.ctrl.titech.ac.jp/pdf/AJCnew1.pdf>, 2003.
- [10] G. Kantor and A. Rizzi. Feedback control of underactuated systems via sequential composition: Visually guided control of a unicycle. In *11th International Symposium of Robotics Research (ISRR 2003)*, Siena, Italy, October 2003.
- [11] E. Klavins and D. E. Koditschek. A formalism for the composition of concurrent robot behaviors. In *Proc. IEEE Int'l. Conf. on Robotics and Automation*, volume 4, pages 3395–402, San Francisco, CA, April 2000.
- [12] G. Lafferriere and H.J. Sussmann. A differential geometric approach to motion planning. In Z. Li and J.F. Canny, editors, *Nonholonomic Motion Planning*. Kluwer Academic Press, 1993.
- [13] N.E. Leonard and P.S. Krishnaprasad. Motion control of drift-free, left-invariant systems on lie groups. *IEEE Transactions on Automatic Control*, 40:1539–1554, 1995.
- [14] D. Maiwand. System identification, using vision-based localisation, for a hexapod robot. Master's thesis, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, June 2003.
- [15] P. Morin and C. Samson. Practical stabilization of driftless homogeneous systems based on the use of transverse periodic functions. In *Proceedings of the 40th IEEE Conference on Decision and Control*, pages 1761–1766, Orlando, FL, December 2001.
- [16] R.M. Murray and S.S. Sastry. Nonholonomic motion planning: Steering using sinusoids. *IEEE Transactions on Automatic Control*, 38:700–716, 1993.
- [17] P. Murriero, D. Fontanelli, and A. Bicchi. Visual-servoed parking with limited view angle. In B. Siciliano and P. Dario, editors, *Experimental Robotics VIII, STAR 5*, pages 265–274. Springer-Verlag, Berlin Heidelberg, 2003.
- [18] S. Patel, S. Jung, J. Ostrowski, R. Rao, and C. J. Taylor. Sensor based door navigation for a nonholonomic vehicle. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 3081–3086, Washington D.C., May 2002.
- [19] A. A. Rizzi, J. Gowdy, and R. L. Hollis. Distributed coordination in modular precision assembly systems. *International Journal of Robotics Research*, 20(10):819–38, October 2001.
- [20] C. Samson. Control of chained systems: Application to path following and time-varying point stabilization. *IEEE Transactions on Automatic Control*, 40:64–77, 1995.
- [21] O.J. Sordalen and O. Egeland. Exponential stabilization of nonholonomic chained systems. *IEEE Transactions on Automatic Control*, 40:35–49, 1995.
- [22] D.P. Tsakiris, P. Rives, and C. Samson. Extending visual servoing techniques to nonholonomic mobile robots. In D.J. Kreigman, G.D. Hager, and A.S. Morse, editors, *The Confluence of Vision and Control, LNCIS237*, pages 106–117. Springer, London, 1998.