A route to limit cycles via unfolding the pitchfork with feedback

Paul B. Reverdy

***Abstract*— Motivated by the problem of developing a system that can switch between low-level control vector fields, we study the feedback interconnection of a decision mechanism based on the pitchfork bifurcation with a plant whose dynamics are given by the linear combination of two control vector fields; the state of the pitchfork system then determines the linear combination coefficients for the plant. We show how the plant state can be fed back to the pitchfork by using the Lyapunov functions of the control vector fields as an unfolding parameter in the pitchfork, and derive conditions under which the closed-loop system exhibits a Hopf bifurcation leading to stable limit cycles.**

An effective method for developing automated systems that are capable of complex behavior is to construct a number of base behaviors and then to construct a mechanism that can select the appropriate base behavior to perform at any given time. For example, a hybrid dynamical system may have a number of different continuous controllers and the active controller is selected by a finite-state automaton which monitors the continuous state and switches among controllers according to some predetermined logic. The construction and analysis of such hybrid systems has received considerable attention in the literature in recent years, e.g., [1], [2], [3].

The recent work [4] introduced a new method to select among base controllers that is purely continuous. In contrast to the hybrid systems approach which uses a finite-state automaton to perform the selection, the authors of [4] develop a selection mechanism based on a continuous dynamical system by adapting a bio-inspired model of decision-making from [5]. The resulting dynamical system has only continuous states. Under appropriate conditions [4], this dynamical system exhibits a stable limit cycle under which the system alternatively selects each of two base control vector fields.

This paper is dedicated to performing analysis to elucidate the dynamical mechanisms at play in [4]. It is well understood [5] that the decision-making mechanism from [5] exhibits a pitchfork bifurcation. Intuitively, the system studied in [4] takes this decision-making system with a pitchfork bifurcation and incorporates it in a feedback loop to produce a Hopf bifurcation. In this paper, we provide rigorous justification for this intuition by showing how the feedback can be incorporated into the pitchfork through a mechanism that can be interpreted as an unfolding of the pitchfork and use this interpretation to understand the dynamics of the closed-loop feedback system.

The state of the pitchfork system is used to choose between two vector fields defined on a physical domain, and the state of the physical system follows the dynamics defined by the chosen vector field. The unfolding parameters of the pitchfork are defined as functions of this physical system state, which closes the feedback loop. The authors of [4] studied the dynamics of the closed-loop system for a particular choice of vector fields on the physical domain and showed that the closed-loop dynamics exhibit a Hopf bifurcation leading to a stable limit cycle. The contributions of this paper are twofold: first, to investigate the fundamental dynamical mechanism at play in [4], and second to find a broader class of physical vector fields that admit the Hopf bifurcation and resulting stable limit cycle.

The remainder of the paper is organized as follows. Section I summarizes the relevant theory concerning the unfolding of bifurcations. Section II shows how an unfolded pitchfork system can be put into a feedback interconnection with a simple bistable system to result in a Hopf bifurcation. Section III summarizes the dynamical system from [4]. Section IV shows how the unfolded pitchfork studied in Section II can be substituted for the decision-making mechanism used in [4] and states and proves the main result: sufficient conditions under which this new system exhibits a Hopf bifurcation. Section V presents the results of a numerical implementation of the main result and VI concludes.

I. Unfolding of a pitchfork bifurcation

Consider the normal form of the supercritical pitchfork bifurcation:

$$\dot{x} = f(x, \mu) := x(\mu - x^2), \tag{1}$$

which defines f. This system has the following equilibria:

$$x = \begin{cases} 0, & \forall \mu \\ \pm\sqrt{\mu}, & \mu > 0. \end{cases}$$

The nonzero equilibria are stable when they exist; the equilibrium $x = 0$ is stable for $\mu < 0$ and unstable for $\mu > 0$.

We are interested in investigating ways in which a system exhibiting a pitchfork bifurcation can be incorporated in a feedback loop such that the closed-loop system exhibits a Hopf bifurcation. A natural way to carry out such analysis is to consider the feedback loop as a perturbation of the pitchfork system (1). In turn, a standard way to study perturbations of bifurcations is through the use of the so-called *singularity theory* of bifurcations, which is a general theory of the qualitative properties of bifurcation problems.

In the singularity theory approach, one calls F a *perturbation* of a bifurcation problem $f(x, \mu) = 0$ if F is a function $F(x, \mu, \alpha_1, \ldots, \alpha_k)$ forming a k-parameter family of bifurcation problems such that $F(x, \mu, 0, \ldots, 0) = f(x, \mu)$. For

The author is with the Department of Aerospace and Mechanical Engineering, University of Arizona, Tucson, AZ 85721 USA. preverdy@email.arizona.edu. This work has been supported in part by startup funds provided by the University of Arizona, as well as Air Force Research Laboratory grant FA8650-15-D-1845 subcontract 669737-6.

many bifurcation problems f, a special perturbation F exists such that any perturbation of f whatsoever is (topologically) equivalent to $F(\cdot,\cdot,\alpha)$ for some $\alpha \in \mathbb{R}^k$ near the origin. Such a special perturbation is called a *universal unfolding* of f. It is a well-known result [6, I.1.13] that

$$F(x,\mu,\alpha) = x(\mu - x^2) + \alpha_1 + \alpha_2 x^2 \tag{2}$$

is a universal unfolding of the pitchfork (1).

II. Pitchfork plus feedback yields Hopf

We will consider a particular perturbation with $\alpha_1 = y, \alpha_2 = -y$ and consider the feedback interconnection of the perturbed pitchfork with a bistable system that encodes dynamics of y. In particular, we let $\dot{y} = (-x + \mu y)$ and obtain the following system.

$$\dot{x} = F(x,\mu,y) = x(\mu - x^2) + y - yx^2 \tag{3}$$
$$\dot{y} = G(y,\mu,x) = (-x + \mu y). \tag{4}$$

The bistable dynamics (4) are designed to have $y = x/\mu$ as a fixed point. Theorem 2 below shows that the feedback system (3)–(4) exhibits a Hopf bifurcation as μ is raised through $\mu = 0$. The proof relies on the following theorem from [7] which summarizes the conditions under which a system undergoes Hopf bifurcation. The coefficient $\ell_1 \in \mathbb{R}$, given by [7, (3.4.29)], is known as the Lyapunov coefficient.

Theorem 1 (Hopf bifurcation, [7, Theorem 3.4.2]): Suppose that the system $\dot{z} = f(z,\mu), z \in \mathbb{R}^n, \mu \in \mathbb{R}$, has an equilibrium (z_0,μ_0) and the following properties are satisfied:

1) The Jacobian $D_z f|_{(z_0,\mu_0)}$ has a simple pair of pure imaginary eigenvalues $\lambda(\mu_0)$ and $\bar{\lambda}(\mu_0)$ and no other eigenvalues with zero real parts,
2) $\mathrm{d}(\mathrm{Re}\ \lambda(\mu))/\mathrm{d}\mu|_{\mu=\mu_0} = d \neq 0$.

Property 1) implies that there is a smooth curve of equilibria $(z(\mu),\mu)$ with $z(\mu_0) = z_0$. The eigenvalues $\lambda(\mu), \bar{\lambda}(\mu)$ of $D_z f|_{(z(\mu_0),\mu_0)}$ which are imaginary at $\mu = \mu_0$ vary smoothly with μ. If, moreover, Property 2) is satisfied, then there is a unique three-dimensional center manifold passing through (z_0,μ_0) in $\mathbb{R}^n \times \mathbb{R}$ and a smooth system of coordinates (preserving the planes μ=const.) for which the Taylor expansion of degree 3 on the center manifold is given by [7, (3.4.8)]. If $\ell_1|_{(z_0,\mu_0)} \neq 0$, there is a surface of periodic solutions in the center manifold which has quadratic tangency with the eigenspace of $\lambda(\mu_0), \bar{\lambda}(\mu_0)$ agreeing to second order with the paraboloid $\mu = -(\ell_1|_{(z_0,\mu_0)}/d)(x^2 + y^2)$. If $\ell_1|_{(z_0,\mu_0)} < 0$, then these periodic solutions are stable limit cycles, while if $\ell_1|_{(z_0,\mu_0)} > 0$, the periodic solutions are repelling.

The following theorem makes precise the result concerning the Hopf bifurcation of the feedback system (3)–(4).

Theorem 2: The system (3)–(4) represents the feedback interconnection of a perturbed pitchfork bifurcation with a bistable system. The following statements hold.

i) The feedback system (3)–(4) exhibits a Hopf bifurcation at $(x_0,y_0,\mu_0) = (0,0,0)$.

ii) $F(x,\mu,0)$ given by (3) has a symmetric pitchfork singularity at $(x_0,\mu_0) = (0,0)$.

iii) For $y \neq 0$, the bifurcation problem $F(x,\mu,y)$ is a one parameter unfolding of the symmetric pitchfork.

Proof: i) The origin is an equilibrium for the feedback dynamics (3)–(4). The Jacobian J_0 of the feedback dynamics evaluated at the origin is

$$J_0 = \begin{bmatrix} \mu & 1 \\ -1 & \mu \end{bmatrix},$$

which has eigenvalues $\lambda = \mu \pm i$, so the requisite conditions of Theorem 1 are satisfied, and the feedback dynamics exhibit a Hopf bifurcation.

ii) $F(x,\lambda,0)$ is precisely the normal form for a pitchfork with singularity at $(x_0,\mu_0) = (0,0)$.

iii) By construction, $F(x,\mu,y)$ is a special case of the universal unfolding (2) of the symmetric pitchfork with $\alpha_1 = y, \alpha_2 = -y$. ■

III. Motivation dynamics for composing control vector fields

Theorem 2 shows that an unfolded pitchfork can be put in a feedback interconnection with a bistable system to yield a system that exhibits a Hopf bifurcation. We are ultimately interested in replacing the decision-making mechanism used in [4] with the unfolded pitchfork mechanism (3). In this section, we review the problem formulation from [4].

A. Base vector fields

We suppose that our system evolves on a state space $\mathcal{D} \subseteq \mathbb{R}^n$ and that the state $x \in \mathcal{D}$ can be controlled by selecting one of the two control vector fields

$$F_i : \mathcal{D} \to T\mathcal{D}, i \in \{1,2\}. \tag{5}$$

We assume that each vector field F_i has an associated Lyapunov function $f_i : \mathcal{D} \to \mathbb{R}$ such that

$$\dot{x} = F_i(x) \Rightarrow \dot{f}_i = (\nabla^T f_i) F_i \leq 0 \ \forall x \in \mathcal{D}, \tag{6}$$

with equality only for $x \in A_i$, where $A_i \subset \mathcal{D}$ is the attracting set for vector field F_i. The authors of [4] considered specific choices of vector fields F_i where the attracting sets were the singletons $A_i = \{x_i^*\} \in \mathcal{D}$, but this assumption can be relaxed to allow for more general attracting sets.

B. Motivation state

The system studied in [4] has an additional state $m = (m_1, m_2, m_U) \in \Delta^2$ referred to as a *motivation* state, where $\Delta^2 = \{x \in \mathbb{R}^3 : x_i \geq 0, \sum_i x_i = 1\}$ is the 2-simplex. The motivation state encodes the system's motivation to perform its three possible actions: m_1 corresponds to activating vector field F_1, m_2 to activating vector field F_2, and $m_U = 1 - m_1 - m_2$ to activating neither. The normalization condition $\sum_i x_i = 1$ that defines the simplex encodes the notion that motivation is limited, so the vector fields cannot be activated to an arbitrary degree.

In a standard hybrid system approach, the currently-selected action would be encoded in a discrete switching variable. The motivation state can be thought of as a convex relaxation of this discrete variable: setting $m = (1,0,0)$

encodes the motivation to activate only vector field F_1, while setting $m = (0, 1, 0)$ encodes the motivation to activate only vector field F_2. The convex relaxation allows the motivation state to encode actions that would be impossible to encode in a discrete switching variable, e.g., setting $m = (1/2, 1/2, 0)$ encodes the motivation to split the system's effort equally between both vector fields.

C. Composing vector fields

Given the motivation state m, the motivation dynamical system composes the base vector fields F_1 and F_2 by linear combination using m_1 and m_2 as weights such that the state dynamics are given by

$$\dot{x} = m_1 F_1(x) + m_2 F_2(x). \tag{7}$$

It is often useful to rewrite (7) by changing from the standard coordinates on Δ^2 to mean-difference coordinates defined by

$$\Delta m = m_1 - m_2,\ \bar{m} = (m_1 + m_2)/2.$$

In the following section, Δm plays a role analogous to that of x in (3). Note that the transformation from the standard coordinates (m_1, m_2, m_U) to $(\Delta m, \bar{m})$ is invertible at all points in Δ^2.

In these coordinates, the dynamics (7) take the form

$$\begin{aligned}\dot{x} &= \frac{2\bar{m} + \Delta m}{2} F_1(x) + \frac{2\bar{m} - \Delta m}{2} F_2(x) \\ &= \frac{\Delta m}{2} \Delta F(x) + 2\bar{m}\bar{F}(x),\end{aligned} \tag{8}$$

where $\Delta F(x) = F_1(x) - F_2(x)$ and $\bar{F}(x) = (F_1(x) + F_2(x))/2$, respectively.

D. Motivation dynamics

The motivation dynamical system studied in [4] allows its action selection to evolve by endowing the motivation state with its own dynamics $\dot{m} = f_m(m; v, \sigma)$, where $v = (v_1, v_2) \in \mathbb{R}^2_+$ are positive values that encode the urgency of activating each of the two vector fields and $\sigma > 0$ is a parameter. Analogous to the definitions of Δm and $\bar{m}$, we define mean-difference coordinates for the value space as

$$\Delta v = v_1 - v_2,\ \bar{v} = (v_1 + v_2)/2.$$

In the mean-difference coordinates, the components of f_m are given by [4, Eqns. (14), (15)]:

$$\begin{aligned}\dot{\Delta m} &= f_{\Delta m}(\Delta m, \bar{m}; \bar{v}, \sigma; \Delta v) \\ &= -\left(\frac{2\bar{m} + \Delta m}{2\bar{v} + \Delta v} - \frac{2\bar{m} - \Delta m}{2\bar{v} - \Delta v}\right) \\ &\quad + \bar{v}\Delta m(1 - 2\bar{m}) + \Delta v(1 - 2\bar{m})(1 + \bar{m}),\end{aligned} \tag{9}$$

$$\begin{aligned}\dot{\bar{m}} =& f_{\bar{m}}(\Delta m, \bar{m}; \bar{v}, \sigma; \Delta v) \\ =& \frac{1}{2}\Bigg(-\frac{2\bar{m} + \Delta m}{2\bar{v} + \Delta v} - \frac{2\bar{m} - \Delta m}{2\bar{v} - \Delta v} \\ &+ \frac{2\bar{v} + \Delta v}{2}(1 - 2\bar{m})(1 + \frac{2\bar{m} + \Delta m}{2}) \\ &+ \frac{2\bar{v} - \Delta v}{2}(1 - 2\bar{m})(1 + \frac{2\bar{m} - \Delta m}{2}) \\ &- \frac{\sigma}{2}(2\bar{m} + \Delta m)(2\bar{m} - \Delta m)\Bigg).\end{aligned} \tag{10}$$

These dynamics were first studied by [5] to model the decision-making process of honeybee swarms. It is well understood [5], [8] that these dynamics exhibit a pitchfork bifurcation as σ is raised through a critical value. However, the complex algebraic structure of the functions (9) and (10) makes it difficult to analyze the bifurcation characteristics of the resulting closed-loop system. The goal of the present paper is to use the unfolding theory to understand the fundamental dynamical mechanism at play.

E. Hopf bifurcation

The motivation dynamical system studied in [4] closes the feedback loop from plant state x to motivation state m by letting the values v_i of each vector field F_i depend on the Lyapunov function f_i. In particular, Theorem 2 of [4] studies the case where $v_i = v^* f_i$, where $v^* > 0$ is a feedback gain. The theorem shows that, for a particular choice of vector fields F_i and Lyapunov functions f_i, as well as appropriate values of the parameter σ, the closed-loop dynamics undergo a Hopf bifurcation resulting in stable periodic solutions as v^* is raised through a critical threshold value v_0^*.

The result [4, Theorem 2] is dependent on the particular choice of $F_i(x) = x - x_i^*$ and $f_i(x) = \frac{1}{2}\|x - x_i^*\|_2^2$, where $x_i^* \in \mathcal{D}$ are goal states. In the next section we show how an analogous result can be obtained for a much broader class of vector fields F_i by replacing the motivation dynamics (9)–(10) with an unfolded pitchfork system of the form (3).

IV. Main Result

Suppose as before that our system has two base behaviors F_i defined by (5) with associated Lyapunov functions (6). We replace the motivation dynamics (9)–(10) with an unfolded pitchfork system of the form (3) and derive conditions under which this new system exhibits a Hopf bifurcation analogous to the one found in [4, Theorem 2].

A. A new motivation dynamics

As in the motivation system from [4], let $m = (\Delta m, \bar{m}) \in \Delta^2$. Let the control system act by applying the linear combination of the base vector fields F_i as in (8). We develop a feedback control law for the state m such that the closed-loop system exhibits a Hopf bifurcation. In the system studied in [4, Theorem 2], both Δm and $\bar{m}$ had dynamics given by (9) and (10), respectively. Here we fix $\bar{m} = 1/2$ and endow Δm with dynamics inspired by the unfolded pitchfork dynamics (3). With the assumption $\bar{m} = 1/2$, the dynamics (7) become

$$\dot{x} = \frac{\Delta m}{2}\Delta F(x) + \bar{F}(x). \tag{11}$$

Note that, setting $\bar{m} = 1/2$, one recovers $\dot{x} = F_1(x)$ when $\Delta m = +1$ and $\dot{x} = F_2(x)$ when $\Delta m = -1$, respectively.

Define Δf and $\bar{f}$ analogously to the definitions of ΔF and $\bar{F}$, namely, $\Delta f(x) = f_1(x) - f_2(x)$ and $\bar{f}(x) = (f_1(x) + f_2(x))/2$, respectively. In the motivation dynamics formulation in [4], Δf and $\bar{f}$ play the role of unfolding parameters in the dynamics (9)–(10) of the motivation state m. However, the algebraic structure of those dynamics makes it difficult to understand the role of the unfolding parameters.

Here, in analogy with (3) we set

$$\begin{aligned} \dot{\Delta m} &= F(\Delta m, \mu, \Delta f) \\ &= \Delta m(\mu - \Delta m^2) + \Delta f(1 - \Delta m^2), \end{aligned} \tag{12}$$

where $\mu \in \mathbb{R}$ is a bifurcation parameter and $\Delta f = \Delta f(x)$.

The resulting closed-loop system has state space $\mathcal{D} \times \mathbb{R}$ and dynamics

$$\dot{z} = (\dot{x}, \dot{\Delta m}) = f_z(z; \mu), \tag{13}$$

where the two components of f_z are given by (11) and (12), respectively.

B. Deadlock condition

When certain conditions hold, the closed-loop system defined by (13) exhibits a Hopf bifurcation as μ is raised above a critical value. The result is made precise in Theorem 3 below. The bifurcation occurs at a so-called *deadlock* equilibrium defined as follows.

Definition 1: Consider $z = (x, \Delta m) \in \mathcal{D} \times \mathbb{R}$ with dynamics given by (13). An equilibrium $z_d = (x_d, 0)$ is called a *deadlock* equilibrium.

The following proposition gives sufficient conditions for a deadlock equilibrium to exist.

Proposition 1: Let F_i and f_i be the base vector fields and associated Lyapunov functions as defined in (5) and (6). Suppose that there exists a point $x_d \in \mathcal{D}$ such that $F_1(x_d) = -F_2(x_d) =: F_d$ and $f_1(x_d) = f_2(x_d)$. Then the point $z_d = (x_d, \Delta m_d) = (x_d, 0)$ is an equilibrium point of the dynamics (13).

Proof: Consider first the dynamics of Δm given by (12). From the definition of Δf, we have $\Delta f(x_d) = f_1(x_d) - f_2(x_d) = 0$. Note that, when $\Delta f = 0$, $\Delta m = 0$ is an equilibrium of (12) for all values of μ.

Now consider the dynamics of x given by (11). From the definitions of ΔF and $\bar{F}$, we have $\Delta F(x_d) = F_1(x_d) - F_2(x_d) = 2F_d$ and $\bar{F}(x_d) = (F_1(x_d) + F_2(x_d))/2 = 0$, respectively. Therefore, $z_d = (x_d, \Delta m_d) = (x_d, 0)$ is an equilibrium of (11). ■

C. Hopf bifurcation result

The bifurcation properties of the closed-loop dynamics (13) are governed by the Jacobian of the closed-loop vector field evaluated at the deadlock equilibrium z_d. This Jacobian, in turn, depends on the Jacobian $\nabla^T \bar{F}$ of the vector field $\bar{F}$. We define

$$\bar{H}(x) = \nabla^T \bar{F}(x), \ \bar{H}_d = \bar{H}(x_d),$$

where x_d is associated with a deadlock equilibrium z_d. We assume that $\bar{H}_d$ has n eigenvectors $\{v_i\}, i \in \{1, \ldots, n\}$ with associated eigenvalues $\{\beta_i\}, i \in \{1, \ldots, n\}$.

We can now state and prove the main result.

Theorem 3: Suppose that there exists an $x_d \in \mathcal{D}$ such that Proposition 1 holds and $z_d = (x_d, 0)$ is a deadlock equilibrium for (13). Let $\bar{H}_d$ be the Jacobian of the vector field $\bar{F}$ evaluated at the deadlock state $x = x_d$ with eigenvectors $\{v_i\}$ and eigenvalues $\{\beta_i\}$ as defined above. Suppose further that $v_1 = F_d$ with $\beta_1 \in \mathbb{R}$, and that the $n-1$ other eigenvalues $\beta_i, i \in \{2, \ldots, n\}$ are not purely imaginary, i.e., with non-zero real parts. Finally, assume that $\nabla^T \Delta f(x_d)$ is orthogonal to each $v_i, i \in \{2, \ldots, n\}$.

Then, the system (13) undergoes a Hopf bifurcation as μ is raised through $\mu = \mu_0 := -\beta_1$ as long as

$$\beta_1^2 < -\phi := -(\nabla^T \Delta f(x_d)) F_d.$$

Proof: Let J_d be the Jacobian of the dynamics (13) evaluated at the deadlock state z_d. Straightforward computation shows that J_d is equal to

$$J_d = \begin{bmatrix} \bar{H}_d & F_d \\ \nabla^T \Delta f(x_d) & \mu \end{bmatrix}.$$

By assumption, we have $\bar{H}_d F_d = \beta_1 F_d$. Then, the vector $[aF_d^T \ 1]^T$ is an eigenvector of J_d for some $a \in \mathbb{C}$, since

$$\begin{aligned} J_d \begin{bmatrix} aF_d \\ 1 \end{bmatrix} &= \begin{bmatrix} \bar{H}_d & F_d \\ \nabla^T \Delta f(x_d) & \mu \end{bmatrix} \begin{bmatrix} aF_d \\ 1 \end{bmatrix} \\ &= \begin{bmatrix} a\bar{H}_d F_d + F_d \\ a(\nabla^T \Delta f(x_d)) F_d + \mu \end{bmatrix} \\ &= \begin{bmatrix} (a\beta_1 + 1) F_d \\ a(\nabla^T \Delta f(x_d)) F_d + \mu \end{bmatrix} \\ &= \lambda \begin{bmatrix} aF_d \\ 1 \end{bmatrix}, \end{aligned} \tag{14}$$

which defines the eigenvalue λ. We show that there is a complex conjugate pair of eigenvalues that solve this eigenvalue problem and that these eigenvalues satisfy the conditions for the Hopf bifurcation Theorem 1.

Let $\phi = (\nabla^T \Delta f(x_d)) F_d$. The eigenvalue problem (14) is equivalent to the two scalar equations

$$(a\beta_1 + 1) = a\lambda, \quad a\phi + \mu = \lambda.$$

Solving the second equation yields $a = (\lambda - \mu)/\phi$, which can be inserted into the first equation to yield

$$\frac{\beta_1(\lambda - \mu)}{\phi} + 1 = \frac{\lambda(\lambda - \mu)}{\phi}.$$

Rearranging yields the polynomial

$$\lambda^2 - (\mu + \beta_1)\lambda + (\mu\beta_1 - \phi) = 0. \tag{15}$$

It is clear that (15) has two purely imaginary roots $\lambda(\mu_0)$ and $\bar{\lambda}(\mu_0)$ when $\mu = \mu_0 = -\beta_1$ if $\mu_0\beta_1 - \phi > 0$. Substituting in the value of μ_0, this condition becomes $\beta_1^2 < -\phi$, as desired. Therefore, J_d has a pair of purely imaginary eigenvalues under these conditions.

The non-degeneracy property 2) of Theorem 1 is satisfied, since $\mathrm{d}(\mathrm{Re}\ \lambda(\mu))/\mathrm{d}\mu|_{\mu=\mu_0} = \frac{1}{2} \neq 0$.

To show that the eigenvalues of J_d satisfy property 1) of Theorem 1, it suffices to show that $\lambda(\mu_0)$ and $\bar{\lambda}(\mu_0)$ are the only eigenvalues with zero real parts. We do this by straightforward computation.

Let $v_i \neq F_d$ be an eigenvector of $\bar{H}_d$. Then

$$J_d \begin{bmatrix} v_i \\ 0 \end{bmatrix} = \begin{bmatrix} \bar{H}_d & F_d \\ \nabla^T \Delta f(x_d) & \mu \end{bmatrix} \begin{bmatrix} v_i \\ 0 \end{bmatrix} = \begin{bmatrix} \bar{H}_d v_i \\ (\nabla^T \Delta f(x_d)) v_i \end{bmatrix} = \begin{bmatrix} \beta_i v_i \\ 0 \end{bmatrix},$$

so $\begin{bmatrix} v_i^T & 0 \end{bmatrix}^T$ is an eigenvalue of J_d with eigenvalue β_i. Therefore, the eigenvalues of J_d as evaluated at $\mu = \mu_0$ consist of $\{\lambda(\mu_0), \bar{\lambda}(\mu_0)\} \cup \{\beta_i\}_{i=2}^n$. By assumption, all but two of these eigenvalues have non-zero real parts, so $\lambda(\mu_0)$ and $\bar{\lambda}(\mu_0)$ are the only eigenvalues with zero real parts and property 1) of Theorem 1 is satisfied.

Applying Theorem 1 yields the desired result: the system (13) undergoes a Hopf bifurcation as μ is raised through $\mu_0 = -\beta_1$. ■

D. Discussion

The implication of Theorem 3 is that, under mild conditions on the eigenvalues and eigenvectors of $\bar{H}_d$, the feedback interconnection of the plant dynamics (8) with an unfolded pitchfork controller (12), exhibits a Hopf bifurcation as the parameter μ is raised through a threshold value. In the post-bifurcation regime, the closed-loop system exhibits limit cycles where the plant cyclicly follows one vector field, then the other.

Note that the new dynamics (12), as written, do not ensure that $\Delta m \in [-1, 1]$ is an invariant set, which is required for the overall motivation state m to remain in Δ^2. However, minor modifications of (12) can be made to ensure that $\Delta m \in [-1, 1]$ is an invariant set. By inspection of (12), it is clear that $\Delta m \in [-1, 1]$ is an invariant set if $|\mu| \leq 1$. While the bifurcation value $\mu_0 = -\beta_1$ does not necessarily satisfy $|\mu_0| = |\beta_1| \leq 1$, it can be shifted by scaling both the vector fields F_i and Lyapunov functions f_i by some constant $b > 0$ such that the bifurcation value of the scaled system does satisfy $|\mu_0| \leq 1$.

In the original motivation dynamics (9)–(10), there are two additional unfolding parameters, namely, $\bar{v}$ and σ. The new motivation dynamics (12) unfolds the pitchfork (1) to incorporate one unfolding parameter Δf. However, because the perturbation (2) is a universal unfolding of the pitchfork, there must be another unfolding analogous to (3) of the original motivation dynamics which incorporates all three parameters $\Delta v, \bar{v}$, and σ. Such an unfolding can be found through Lyapunov-Schmidt reduction.[1] The analysis using the Lyapunov-Schmidt reduction is in progress and will be the subject of a separate report.

[1]The Lyapunov-Schmidt reduction is the projection of the vector field near its singular point along the null space of the Jacobian at the singular point. The reduction uses the implicit function theorem to compute a local approximation of the vector field in the subspace orthogonal to the null space. The reduced dynamics can be used to infer bifurcations of the vector field. See [6] for details.

V. Numerical implementation

In this section, we present the result of a numerical implementation of Theorem 3 and compare the limit cycle thus obtained to the limit cycle found using the same base vector fields in [4].

A. Vector fields, Lyapunov functions

For purposes of comparison, we adopt the same vector fields and Lyapunov functions as in [4], namely, the point attractors and associated Lyapunov functions

$$F_i(x) = x - x_i^*, \ f_i(x) = \frac{1}{2}\|x - x_i^*\|_2^2, \tag{16}$$

where $x_i^* \in \mathcal{D} = \mathbb{R}^2$ are goal locations. We take $x_1^* = (1, 0)^T$ and $x_2^* = (-1, 0)^T$. With this choice of vector fields, the physical state x of the system is quickly attracted to the convex hull of the goal locations, namely, the set $\{x = (x_1, x_2) \in \mathbb{R}^2 : |x_1| \leq 1, x_2 = 0\}$. For this reason, we plot only the x_1 component of the trajectories.

B. Application of Theorem 3

It is readily verified that, with choices (16) of F_i and f_i, the origin $z_d = (x_d, \Delta m_d) = (0, 0)$ is a deadlock equilibrium. Furthermore, we have $\bar{H}_d = -I_2$ and $F_d = \begin{bmatrix} 1 & 0 \end{bmatrix}^T$, so F_d is an eigenvector of $\bar{H}_d$ and the eigenvalues of $\bar{H}_d$ are $\beta_1 = \beta_2 = -1$. Finally, $\nabla^T \Delta f(x_d) = \begin{bmatrix} -2 & 0 \end{bmatrix}$, so $\phi = (\nabla^T \Delta f(x_d)) F_d = -2$ and we have $\beta_1^2 = 1 < 2 = -\phi$. Therefore, the conditions of Theorem 3 are satisfied and we expect to see a stable limit cycle for values of μ greater than the bifurcation value $\mu_0 = -\beta_1 = 1$.

Two visualizations of a numerical simulation of the system (13) are shown in Figures 1 and 2. The bifurcation parameter μ was set equal to $1.1 > 1 = \mu_0$, so the system was in the post-bifurcation regime. As can be seen in Figures 1 and 2, the system does in fact exhibit a stable limit cycle. As mentioned above, the state x_2 tends asymptotically to zero, so we plot only x_1 and Δm. The amplitude of the oscillations in Δm is larger than the amplitude of the oscillations in x_1, so the geometry of the limit cycle is that of an oblong shape, as can be seen in Figure 2.

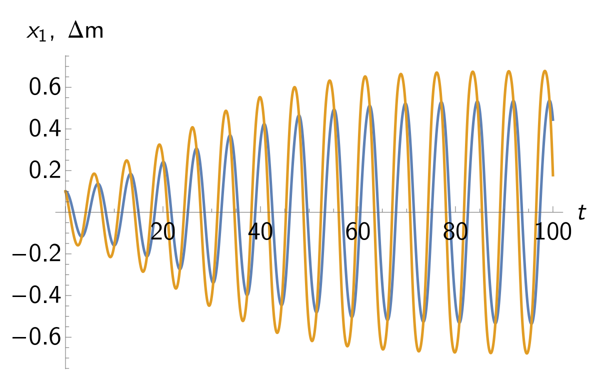

Fig. 1. State trajectory for the solution of (13) with vector fields and Lyapunov functions as specified in (16), for $\mu = 1.1 > \mu_0$. The oscillations of x_1 and the slightly-larger amplitude oscillations of Δm tend towards a periodic orbit.

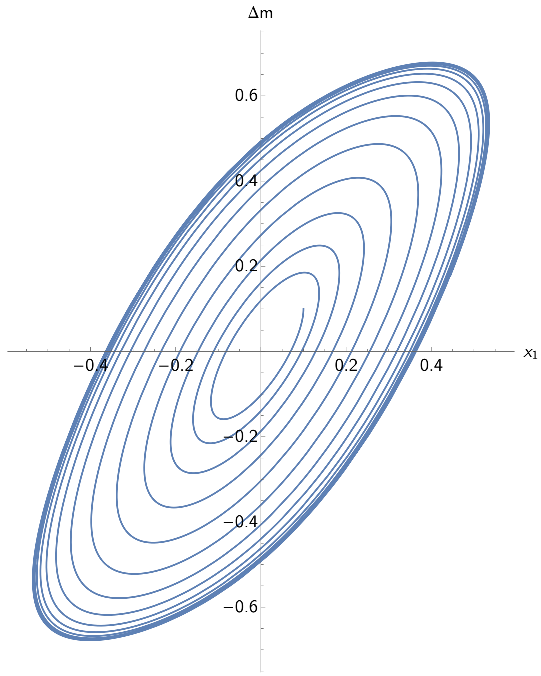

Fig. 2. Trajectory for the solution of (13) with vector fields and Lyapunov functions as specified in (16), for $\mu = 1.1 > \mu_0$. The oscillations of x_1 and the slightly-larger amplitude oscillations of Δm tend towards a periodic orbit, i.e., a limit cycle.

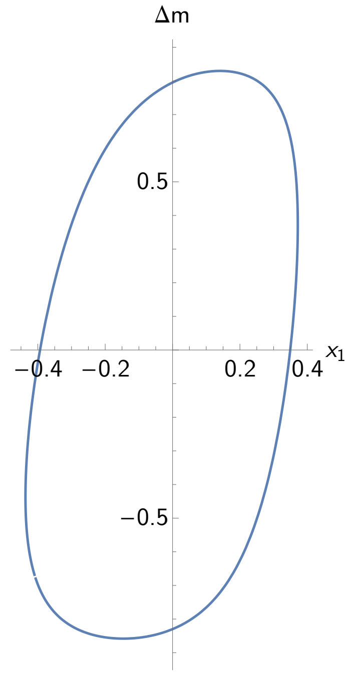

Fig. 3. Numerically-computed limit cycle of the slow system [4, Equation (30)] which is analogous to the system (13) with vector fields and Lyapunov functions as specified in (16). The parameter σ was set equal to 8. As in Figure 2, one sees an oblong-shaped limit cycle with larger excursions in Δm than in x_1.

The behavior of the new system (13) is qualitatively similar to that of the original motivation dynamics system published in [4], as can be seen by comparing Figures 2 and 3. Figure 3 plots the numerically-computed limit cycle for the original motivation dynamics system previously published as Equation (30) of [4]. The parameter σ was set equal to 8. In both figures the limit cycle has an oblong shape with larger oscillations in Δm than in x_1. It may be possible to put the two systems into closer correspondence by careful tuning of the parameters (μ for (13) and σ for [4, (30)]), but the important result is that the limit cycle geometries are similar.

VI. Conclusions

In conclusion, we have studied the problem of composing low-level controllers in order to develop a higher-level behavior, namely, a limit cycle in which the system oscillates between applying one controller and then the next. Inspired by [4], where the authors used a complex dynamical system to perform the composition, we investigated a simpler decision-making system developed from an unfolding of the supercritical pitchfork bifurcation.

We showed how to use feedback to adjust the unfolding parameters of the pitchfork and derived conditions under which the resulting closed-loop system would exhibit a Hopf bifurcation leading to a stable limit cycle. We numerically demonstrated that our new system exhibits limit cycles with similar geometry to that of the original system studied in [4]. The new system helps elucidate the dynamical mechanisms at play in the original work [4] and opens the door to further work using unfolding theory to understand such structures, as has proved useful in other recent work [9]. In particular, we are actively working to use Lyapunov-Schmidt reduction on the original motivation system in order to more completely understand the structure of the unfolding at play. We intend to use this understanding to construct systems with more explicitly-controllable limit cycle geometries. This work will be the subject of a future report.

References

[1] R. Goebel, R. G. Sanfelice, and A. R. Teel, *Hybrid Dynamical Systems: modeling, stability, and robustness*. Princeton University Press, 2012.
[2] J. Liu, N. Ozay, U. Topcu, and R. M. Murray, "Synthesis of reactive switching protocols from temporal logic specifications," *IEEE Transactions on Automatic Control*, vol. 58, no. 7, pp. 1771–1785, 2013.
[3] H. Kress-Gazit, G. E. Fainekos, and G. J. Pappas, "Temporal-logic-based reactive mission and motion planning," *IEEE Transactions on Robotics*, vol. 25, no. 6, pp. 1370–1381, 2009.
[4] P. B. Reverdy and D. E. Koditschek, "A dynamical system for prioritizing and coordinating motivations," *SIAM Journal on Applied Dynamical Systems*, vol. 17, no. 2, pp. 1683–1715, 2018.
[5] T. D. Seeley, P. K. Visscher, T. Schlegel, P. M. Hogan, N. R. Franks, and J. A. Marshall, "Stop signals provide cross inhibition in collective decision-making by honeybee swarms," *Science*, vol. 335, no. 6064, pp. 108–111, 2012.
[6] M. Golubitsky and D. Schaeffer, *Singularities and Groups in Bifurcation Theory*, ser. Applied Mathematical Sciences. Springer, 1985, vol. 51.
[7] J. Guckenheimer and P. J. Holmes, *Nonlinear oscillations, dynamical systems, and bifurcations of vector fields*, ser. Applied Mathematical Sciences. Springer Science & Business Media, 2013, vol. 42.
[8] D. Pais, P. M. Hogan, T. Schlegel, N. R. Franks, N. E. Leonard, and J. A. Marshall, "A mechanism for value-sensitive decision-making," *PloS one*, vol. 8, no. 9, p. e73216, 2013.
[9] R. Gray, A. Franci, V. Srivastava, and N. E. Leonard, "Multi-agent decision-making dynamics inspired by honeybees," *IEEE Transactions on Control of Network Systems*, vol. 5, no. 2, pp. 793–806, 2018.