Geometric Control, Actuator Feasibility, and Algebraic Equilibria in a Four-Wheel X-Drive Robot by Hassan Shahid

Abstract

Four-wheel X-drive robots can translate laterally and rotate simultaneously, making them useful platforms for autonomous navigation. Their control problem connects three mathematical perspectives: Riemannian geometry for trajectory costs, the Lie group $SE(2)$ for rigid-body pose error, and computational algebraic geometry for equilibrium analysis. This paper develops these connections for an idealized symmetric X-drive with four independently driven omniwheels. Its main analytical findings are an explicit wheel-speed map, a flat left-invariant Riemannian metric induced by squared wheel speeds, and an exact polyhedral constraint on feasible body twists. A logarithmic pose-error feedback law is then normalized against this constraint; a local Lyapunov argument establishes convergence to a stationary target under perfect state knowledge and ideal velocity realization. A parameterized polynomial heading model illustrates how algebraic elimination detects additional equilibrium branches, while also identifying why that model is not evidence of bifurcations in the actual drivetrain. Finally, reproducible numerical integrations of the ideal closed-loop error system illustrate the regulation and speed-bound results. No tests of a physical robot are reported. The framework offers a mathematically explicit starting point for encoder- and inertial-sensor-based navigation and identifies the identification, estimation, and experimental work needed before claims of practical reliability can be made.

1. Introduction

An X-drive robot uses omniwheels whose driven directions are typically arranged at approximately 45° to the chassis axes. Because each wheel supplies a directional component of force while its passive rollers allow transverse motion, the platform can command two translational velocities and a yaw rate without first steering its wheels. These kinematics make X-drives attractive for rapid movement among field locations, but they also create a subtle optimization problem: not every chassis velocity is available at a given motor-speed limit, and the attainable speed depends on direction.

Existing literature already addresses omnidirectional wheel kinematics, constrained trajectory generation, and Lie-group feedback. Lynch and Park develop the relationship between wheel velocities and body twists; Balkcom, Kavathekar, and Mason show that motor-speed constraints can make time-optimal paths different from shortest geometric paths; and Purwin and D'Andrea study trajectory generation and control for four-wheeled omnidirectional systems. Separately, Teng and colleagues investigate logarithmic costs on Lie groups. These works motivate an integrated treatment, but their existence also makes an unsupported claim of novelty inappropriate (Lynch and Park; Balkcom et al.; Purwin and D'Andrea; Teng et al.).

The question here is narrower and testable: **What geometric structure and actuator constraints follow from a symmetric four-wheel X-drive model, and how can they be used**

to construct a feasible stationary-pose controller whose local convergence is provable? A second question asks where polynomial equilibrium tools genuinely apply and where their use would be mathematically misleading.

The paper is an analytical and numerical *model study*, not a hardware performance report. The intended application is a VEX V5 robot with motor encoders and an inertial sensor. VEX documents encoder-based motor rotation feedback and inertial heading measurements, but the actual wheel mounting, radius, effective contact geometry, motor control response, sensor calibration, and field slip must still be measured on the particular robot (VEX Robotics, "V5 Motor Overview"; VEX Robotics, "Using the Inertial Sensor"). A vision system is an optional future extension, not a sensor assumed in the present analysis.

Four conclusions will be established. First, the symmetric wheel map has rank three, permitting arbitrary planar rigid-body velocity in the ideal model. Second, the squared-wheel-speed metric is flat, so introducing Riemannian terminology does not create curved geodesics in this symmetric limit. Third, hard wheel-speed bounds define a nonsmooth convex body in twist space, fundamentally different from the smooth quadratic effort metric. Fourth, a radially normalized logarithmic regulator respects the hard bounds and is asymptotically stable in a local logarithmic chart under explicit idealizations. The algebraic section then separates demonstrable mathematics from conjectural physical mechanisms.

2. Related Work and Scope

The kinematic foundation of omniwheel control is established in Chapter 13 of *Modern Robotics*. The same source distinguishes odometry, wheel-limited chassis motion, and trajectory tracking (Lynch and Park). Direction-dependent speed limits have a history in motion-planning research. Balkcom, Kavathekar, and Mason derive time-optimal curves for a different three-omniwheel geometry, finding that shortest paths and fastest paths need not coincide. Purwin and D'Andrea address a four-wheel omnidirectional vehicle and dynamic trajectory generation. Their results caution against presenting the wheel-limit geometry below as a new discovery.

Geometric controller design treats rigid-body motion on a configuration space rather than treating orientation as an arbitrary Euclidean coordinate. The appropriate planar rigid-body group is $SE(2)$, not $SO(3)$, because the robot moves on a planar field and has one orientation angle. Teng and colleagues demonstrate the broader value and mathematical nuance of designing costs in a Lie algebra. Their approach is more general than the elementary local regulator derived here (Teng et al.).

Computational algebraic geometry, including polynomial elimination and Gröbner bases, is useful when equilibrium equations can be represented by polynomial ideals. Hajnová and Přibylová use such methods to study bifurcation manifolds in nonlinear dynamical models. That result does not transfer automatically to an X-drive: trigonometric angles, logarithms on $SE(2)$, absolute values, and saturation laws are not globally polynomial expressions. Algebraic computation must therefore be tied to a carefully specified polynomial subsystem or an exact

polynomial lifting (Hajnová and Přibylová).

This paper's contribution is the **explicit synthesis and limitation analysis** of these three viewpoints for one idealized symmetric configuration. It makes no priority claim for the component methods, the normalization law, or the proposed theoretical results relative to the full robotics literature.

3. Kinematic Model and Assumptions

3.1 Configuration and frames

Let the chassis pose be $T \in SE(2)$, with world-frame position $p = (x, y)^{\top}$ and heading $\theta \in S^1$:

$$T = [R(\theta)\ p\ 0\ 1\ ], \qquad R(\theta) = [cos\theta\ -\ sin\theta\ sin\theta\ cos\theta\ ].$$

The commanded body twist is $\xi = \left(v_x, v_y, \omega\right)^{\top}$, with translational components expressed in chassis coordinates and ω the yaw rate. Define its Lie-algebra matrix by

$$\hat{\xi} = \left[0\ -\ \omega\ v_x\ \omega\ 0\ v_y\ 0\ 0\ 0\ \right], \qquad \dot{T} = T\hat{\xi}.$$

The analysis assumes: (i) four identical omniwheels in a symmetric X arrangement; (ii) effective wheel radius $r > 0$; (iii) equal effective rotational lever arm $R > 0$; (iv) ideal roller behavior without slip or lift-off; (v) independently commanded wheel speeds with the same magnitude limit $\Omega > 0$; and (vi) negligible dynamics at the outer kinematic control level. The parameter R is the kinematic moment arm in the wheel-speed map, **not** a value to substitute from chassis dimensions without calibration.

3.2 Wheel-to-chassis map

With one consistent wheel ordering and motor sign convention, the ideal relationship is

$$\dot{\phi} = \frac{1}{r}J\xi, \qquad J = \left[\frac{1}{\sqrt{2}}\ \frac{1}{\sqrt{2}}\ R\ -\frac{1}{\sqrt{2}}\ \frac{1}{\sqrt{2}}\ -\ R\ -\frac{1}{\sqrt{2}}\ \frac{1}{\sqrt{2}}\ R\ \frac{1}{\sqrt{2}}\ \frac{1}{\sqrt{2}}\ -\ R\right].$$

Here $\dot{\phi} \in R^4$ contains wheel angular rates. A physical wiring layout can permute rows or reverse row signs without changing the principal results; an asymmetric chassis or different wheel-angle convention requires a new measured Jacobian. Direct multiplication gives

$$J^{\top}J = diag\left(2, 2, 4R^2\right), \qquad rankJ = 3.$$

Proposition 1. Under the assumptions above, the ideal platform has three independent instantaneous chassis velocities. The least-squares body-twist estimate from four wheel rates is

$$\hat{\xi}_{odo} = r\left(J^{\top}J\right)^{-1}J^{\top}\dot{\phi}.$$

Proof. Since $r, R > 0$, the displayed diagonal matrix $J^{\top}J$ is positive definite, so J has full column rank and admits the stated left pseudoinverse. The estimate equals the true twist if the wheel constraints are exact and noiseless. ▫

This algebraic rank statement does not prove that the real robot can track arbitrary accelerations or resist impact disturbances. Those questions depend on mass, torque, traction, motor latency, and field conditions.

3.3 Sensitivity to wheel geometry and encoder errors

The matrix above is useful precisely because it can be falsified by measurement. A wheel installed at the wrong angle, a different effective radius, or a contact point displaced from the assumed position changes its Jacobian row. Four encoder signals then provide four constraints on only three chassis-twist variables. The remaining redundancy can help detect inconsistency, although it cannot tell whether an error arose from slip, geometry, or encoder faults without additional information.

To compare translations and rotation in the same physical units, define the scaled twist $z = \left(v_x, v_y, R\omega\right)^{\top}$ and the dimensionless nominal matrix

$$A = Jdiag\left(1, 1, R^{-1}\right), \qquad A^{\top}A = diag(2, 2, 4).$$

With this scaling, all components of z have units of linear velocity. If the measured wheel tangential speeds are $y = r\dot{\phi}$, the ideal observation equation is $y = Az$. Now represent calibration deviations by a matrix ΔA and measurement or contact disturbance by n, giving $y = (A + \Delta A)z + n$. When the nominal pseudoinverse is used to interpret these measurements,

$$\hat{z} - z = A^{\dagger}\Delta A\, z + A^{\dagger}n.$$

The minimum singular value of A is $\sqrt{2}$. Therefore the induced Euclidean norm of $A^{\dagger}$ equals $1/\sqrt{2}$, and the elementary error bound is

$$\| \hat{z} - z\|_2 \leq \frac{\|\Delta A\|_2 \|z\|_2 + \|n\|_2}{\sqrt{2}}.$$

Proposition 1A. If $\| \Delta A\|_2 < \sqrt{2}$, the perturbed map $A + \Delta A$ retains full column rank.

Proof. For any unit vector u, the reverse triangle inequality implies $\| (A + \Delta A)u\|_2 \geq \| Au\|_2 - \| \Delta Au\|_2 \geq \sqrt{2} - \| \Delta A\|_2 > 0$. Thus the perturbed map has no nonzero null vector. ▫

The condition is sufficient, not necessary. A real calibration error need not be small in this matrix norm, and even a full-rank encoder map can produce a biased translation estimate during slip. Nonetheless, the inequalities suggest a practical identification procedure: command short

motions with independently measured world displacement, fit the Jacobian by least squares, and compare the empirical residuals with sensor noise. The measured matrix, rather than the symmetric idealization, should ultimately determine control allocation.

4. Quadratic Wheel Effort and Riemannian Geometry

A natural smooth surrogate for total wheel activity is the sum of squared angular speeds,

$$Q(\xi) = \| \dot{\phi} \|_2^2 = \xi^\top G \xi, \qquad G = \frac{1}{r^2} diag\left(2, 2, 4R^2\right).$$

Because G is symmetric and positive definite, it defines an inner product on each tangent space using body velocity. Such a smoothly varying family of inner products is a Riemannian metric. Written in world coordinates, its differential line element is

$$ds^2 = \frac{2}{r^2}\left(dx^2 + dy^2\right) + \frac{4R^2}{r^2} d\theta^2.$$

Proposition 2. The wheel-speed quadratic metric of the ideal symmetric X-drive is flat. On a local angular chart, its unforced geodesics have $\ddot{x} = \ddot{y} = \ddot{\theta} = 0$.

Proof. The translational body velocity is $v = R(\theta)^\top \dot{p}$. Orthogonality of $R(\theta)$ implies $v_x^2 + v_y^2 = \dot{x}^2 + \dot{y}^2$. Substituting this identity into Q produces the displayed metric with constant coefficients in (x, y, θ). All coordinate derivatives of these metric coefficients vanish, so the Christoffel symbols in this chart vanish. The geodesic equations consequently reduce to zero coordinate acceleration. ▫

The result is less exotic than the terminology might suggest. Under symmetry, effort-minimizing fixed-duration, unconstrained motion consists of a straight world-frame translation and a constant angular rate, using a locally shortest angular displacement. The circle topology of S^1 still matters: at antipodal headings, more than one shortest angular lift exists. Obstacles, changing traction, configuration-dependent weights, and hard actuator limits can substantially alter optimal trajectories. A smooth metric also does not by itself encode the nonsmooth motor-speed ceiling.

Most importantly, Q is a **wheel-speed effort surrogate**, not measured electrical energy. Actual electrical work depends on motor torque, current, efficiency, rolling resistance, acceleration, and load. Describing these geodesics as experimentally established minimum-energy paths would therefore be unjustified.

4.1 A fixed-duration lower bound

The metric permits a simple closed-form calculation that is useful as a consistency check for any subsequent numerical planner. Consider motion from $\left(p_0, \theta_0\right)$ to $\left(p_f, \theta_f\right)$ during a prescribed

positive duration τ. Choose a continuous local angular lift, with $\Delta\theta = \theta_f - \theta_0$, and write $\Delta p = p_f - p_0$. The integral of the quadratic cost equals

$$E = \int_0^\tau Q(\xi(t))\,dt = \frac{2}{r^2}\int_0^\tau \|\dot{p}(t)\|_2^2\,dt + \frac{4R^2}{r^2}\int_0^\tau \dot{\theta}(t)^2\,dt.$$

By the Cauchy-Schwarz inequality, each integral of squared velocity is bounded below by the square of the corresponding net displacement divided by τ. Therefore

$$E \geq \frac{2\|\Delta p\|_2^2 + 4R^2(\Delta\theta)^2}{r^2\tau}.$$

Proposition 2A. Ignoring obstacles and hard wheel-speed bounds, the lower bound is attained by constant world-frame translational velocity $\dot{p} = \Delta p/\tau$ and constant angular rate $\dot{\theta} = \Delta\theta/\tau$.

Proof. Equality in the two Cauchy-Schwarz inequalities requires each velocity to be constant almost everywhere. The proposed velocities satisfy the endpoint constraints and attain equality in both terms. ▫

The formula reveals a common conceptual trap: extending the travel time reduces this squared-speed integral, even though a real motor may consume nonzero baseline power while it waits. Moreover, constant *world-frame* velocity with simultaneous rotation generally demands wheel rates that vary with heading, so this unconstrained extremal curve may violate the exact speed limit. The result is consequently a lower bound and diagnostic reference, not automatically an executable motion plan.

5. Exact Actuator-Speed Geometry

Let U be the admissible body-twist set when all wheel rates obey $|\dot{\phi}_i| \leq \Omega$. Define

$$h(\xi) = \frac{|v_x| + |v_y|}{\sqrt{2}} + R|\omega|.$$

Proposition 3. For the symmetric wheel map, the four wheel bounds are equivalent to one exact convex inequality:

$$\max_{1\leq i\leq 4} |(J\xi)_i| = h(\xi), \qquad U = \{\xi : h(\xi) \leq r\Omega\}.$$

Proof. Set $a = (v_x + v_y)/\sqrt{2}$ and $b = (-v_x + v_y)/\sqrt{2}$. The entries of $J\xi$ are $a + R\omega$, $b - R\omega$, $b + R\omega$, and $a - R\omega$. Hence

$$\max_i |(J\xi)_i| = max(|a|, |b|) + R|\omega|.$$

The real-number identity $max(|x + y|, |x - y|) = |x| + |y|$ yields the stated expression. Dividing by r gives the equivalence to the individual wheel-speed constraints. ▫

The set U is a centrally symmetric convex polytope. For zero yaw rate its section is the diamond

$$|v_x| + |v_y| \leq \sqrt{2}\, r\Omega.$$

Consequently, ideal pure axial motion can attain speed $\sqrt{2}r\Omega$, whereas motion at a 45° diagonal has maximum Euclidean speed $r\Omega$. Pure rotation has maximum rate $r\Omega/R$. The speed in these statements is **chassis speed**, not individual wheel surface speed. Directional anisotropy follows from the hard componentwise motor bounds despite the isotropic *quadratic* translation metric.

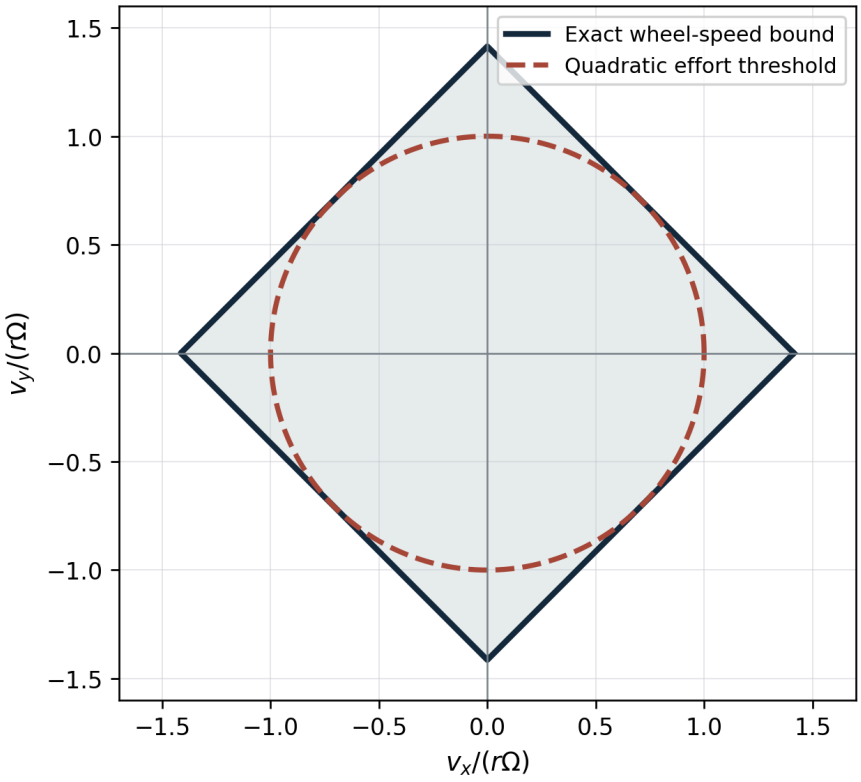


Figure 1. Exact fixed-heading actuator-feasible translation velocities (solid diamond) and a quadratic wheel-effort threshold (dashed circle). The circle is conservative, not an equivalent motor limit.

The corresponding time-minimization problem should impose $\xi(t) \in U$ throughout the trajectory. A minimum-squared-wheel-speed problem instead minimizes $\int \xi^\top G \xi \, dt$; the objectives are distinct. In particular, shortest paths in a flat metric need not be time-optimal under a body-frame polytope that rotates with the heading (Balkcom et al.; Purwin and D'Andrea).

5.1 Exact travel times for restricted maneuvers

For some elementary motions, the polyhedral actuator constraints give exact minimum-time

formulas without numerical optimization. Suppose the heading is fixed at θ_0 throughout a translation. Let $\Delta p_b = R(\theta_0)^\top (p_f - p_0)$ be the desired displacement resolved in that unchanged body frame. Because $\omega = 0$, each feasible trajectory satisfies $|v_x| + |v_y| \le \sqrt{2} r\Omega$. Integration and the triangle inequality give

$$\| \Delta p_b \|_1 \le \int_0^\tau \left(|v_x| + |v_y| \right) dt \le \sqrt{2} r\Omega \tau.$$

Proposition 3A. Among ideal fixed-heading motions constrained only by wheel angular speeds, the minimum time to translate by Δp_b is

$$\tau_{min,translation} = \frac{|\Delta p_{b,x}| + |\Delta p_{b,y}|}{\sqrt{2} r\Omega}.$$

Proof. The preceding inequality establishes a lower bound. A constant commanded body velocity $v_b = \Delta p_b / \tau_{min,translation}$ achieves it and obeys the wheel bound with equality whenever $\Delta p_b \neq 0$. ▫

Likewise, for a pure rotation with $v_x = v_y = 0$, the shortest angular lift $\Delta\theta$ gives

$$\tau_{min,rotation} = \frac{R|\Delta\theta|}{r\Omega}.$$

These formulas are intentionally restricted. If the robot may change heading while translating, its feasible translation diamond rotates in the world frame. A sequence combining rotations and translations might beat a constant-heading maneuver for some endpoint conditions, and independent acceleration limits would prevent instantaneous changes in velocity. Consequently, applying the fixed-heading formula to every general pose-to-pose planning problem would overstate its scope. It is useful as an exact benchmark for isolated straight, lateral, diagonal, and turn-in-place trials.

6. A Locally Stable Feasible Geometric Regulator

6.1 Lie-group pose error

Let $T_d \in SE(2)$ be a stationary desired pose. Define the right-hand target displacement in body coordinates as

$$E = T^{-1} T_d, \qquad e = (\log E)^\vee \in R^3.$$

The principal logarithm is smooth on a domain containing the identity but excluding the rotation branch cut at angle $\pm\ \pi$. Translation coordinates in e are Lie-logarithmic coordinates; they should not generally be confused with a direct subtraction of world positions. This representation

handles orientation as a group element rather than as an unbounded raw angle (Teng et al.).

Without wheel constraints, an intuitive stationary-target command is $\xi_c = ke$ for a gain $k > 0$. Because h is positively homogeneous, a scalar normalization can enforce the hard motor limits:

$$\lambda(e) = \{k,\ e = 0,\ min\left(k, \frac{r\Omega}{h(e)}\right),\ e \neq 0, \qquad \xi_c = \lambda(e)e.$$

If the orientation and body-translation units are considered together, k has units of inverse time. This law contains no integral term and does not require a numerically solved nonlinear optimization program. It is a kinematic pose regulator, not a complete motor-voltage controller.

Proposition 4. Under exact state knowledge, instantaneous realization $\xi = \xi_c$, and a stationary target inside the logarithm chart, the command is wheel-speed feasible for every error state, and the error converges asymptotically to zero. Near the target the convergence is exponential.

Proof. For $e \neq 0$, positive homogeneity gives

$$h\left(\xi_c\right) = \lambda(e)h(e) \leq r\Omega.$$

For a constant T_d, differentiation of $E = T^{-1}T_d$ gives $\dot{E} = -\hat{\xi}E$. Here $\hat{\xi} = \lambda(e)logE$, which commutes with E because the logarithm is an analytic function of E on the chosen chart. Therefore $\dot{E} = -\lambda(e)ElogE$, and the local logarithmic coordinates satisfy

$$\dot{e} = -\lambda(e)e.$$

With $V(e) = \frac{1}{2}e^{\top}e$, differentiation yields

$$\dot{V} = -\lambda(e) \parallel e\parallel_2^2 < 0 \quad (e \neq 0).$$

The coefficient $\lambda(e)$ is positive. The solution remains on the ray through the initial error, and its norm decreases, so the error stays inside the initial logarithm chart whenever that chart contains the ray to the origin. On a bounded initial-error set, $h(e)$ is bounded and thus $\lambda(e)$ is bounded away from zero. This ensures convergence. In the neighborhood $h(e) < r\Omega/k$, the gain is $\lambda = k$, giving $e(t) = e\left(t_0\right)e^{-k\left(t-t_0\right)}$. ▫

The argument is intentionally local in orientation. A smooth, unique global logarithmic regulator on the entire circle cannot be obtained by simply ignoring its branch cut. The theorem also assumes an ideal inner velocity loop, which actual VEX motors do not satisfy exactly.

6.2 Moving-target extension and limitations

For a differentiable reference pose $T_d(t)$, define ξ_d by $\dot{T}_d = T_d\hat{\xi}_d$. The exact error dynamics satisfy

$$\dot{E} = -\hat{\xi}E + E\hat{\xi}_d.$$

The unconstrained feedforward-plus-feedback choice

$$\xi_{nom} = Ad_E\xi_d + ke$$

reduces the error equation to the same radial logarithmic dynamics when perfect twist tracking is possible. Directly clipping ξ_{nom} at the wheel limit, however, changes those dynamics and can invalidate the tracking proof. A practical planner should reserve actuator authority for correction and impose the full wheel constraints during reference generation. A constrained tracking law would require a separate feasibility and stability analysis; it is not claimed here.

This distinction matters for autonomous navigation. A path is a geometric curve, a *trajectory* attaches timing to that curve, and a tracking controller corrects deviations from the timed reference. Combining the three requires attention to wheel-speed, acceleration, and estimator constraints. The stationary regulator is a baseline that can be tested and compared against a conventional controller, not a guarantee that the fastest route will be tracked.

6.3 A bounded-disturbance analysis

Actual motor dynamics, unmodeled contact effects, and imperfect pose estimates disturb the ideal radial error law. A first mathematical stress test is an *additive error-coordinate disturbance* $d(t)$, yielding

$$\dot{e} = -\lambda(e)e + d(t), \qquad \|d(t)\|_2 \leq \delta.$$

This is a simplified uncertainty model in local logarithmic coordinates. An arbitrary wheel-slip event need not enter in exactly this additive form. Still, it makes the limits of the original Lyapunov claim explicit. On a bounded region contained inside the logarithm chart, let H be a known upper bound on $h(e)$ and define $\underline{\lambda} = min(k, r\Omega/H) > 0$. The case $H = 0$ corresponds only to $e = 0$, so the positive bound can be taken from any nontrivial neighborhood. For $V = \frac{1}{2}\|e\|_2^2$,

$$\dot{V} = -\lambda(e)\|e\|_2^2 + e^{\top}d \leq -\underline{\lambda}\|e\|_2^2 + \delta\|e\|_2.$$

This derivative is strictly negative whenever $\|e\|_2 > \delta/\underline{\lambda}$. Thus, **provided the trajectory remains in the assumed bounded logarithmic region**, the inequality gives the familiar local practical-stability interpretation: a persistent bounded disturbance generally produces a residual neighborhood rather than exact convergence. The bound is conservative and depends on the chosen region. It is not a claim of global input-to-state stability or a guarantee that the robot can safely handle an arbitrarily large collision.

The inequality also provides an operational interpretation. Increasing k may reduce the

small-error response time, but the available correction is capped by $r\Omega$. Once saturation is active, raising k further does not increase the magnitude of the commanded twist. Any proposal to solve large localization errors merely by increasing proportional gain overlooks this finite actuator authority.

6.4 What discrete implementation preserves

The VEX V5 Brain executes sampled commands, whereas Proposition 4 assumes a continuous-time ideal twist. The numerical integration later in this paper applies the explicit Euler update

$$e_{n+1} = \left(1 - \Delta t\, \lambda\left(e_n\right)\right) e_n.$$

Proposition 4A. For this numerical update, if $0 < k\Delta t < 1$ and the initial error is inside a logarithm chart containing the radial segment to the origin, every nonzero error decreases in norm and the iterates converge to zero.

Proof. Since $0 < \lambda\left(e_n\right) \leq k$, the scalar multiplier lies strictly between 0 and 1. It follows that each step shortens the error without reversing its direction. The sequence stays on the initial ray and remains in the chart. The sequence of norms is decreasing, and the gain is bounded below on its initial compact radial segment; therefore the norms approach zero. ▫

This proposition applies to the *discretization of the derived error equation*. It is not by itself a sampled-data stability theorem for a physical robot whose commanded wheel velocities are held constant between CPU updates. Such a proof must account for the exact $SE(2)$ motion over each sampling interval, delays, and the behavior of the motor loops. In practice, timestamps and update jitter belong in the test log because irregular sampling changes the effective control law.

7. A Unified Navigation Optimization Problem

Let T_0 and T_f denote initial and target poses, and let O represent field obstacles. A consistent continuous-time formulation is

$$\min_{T(\cdot),\xi(\cdot),t_f} \int_0^{t_f} \left(\alpha\, \xi^{\top} G \xi + \beta\right) dt \;\; subject\ to \quad \dot{T} = T\hat{\xi},\ \ h(\xi) \leq r\Omega,\ \ T(0) = T_0, \quad T\left(t_f\right) = T_f,\ \ p(t) \notin O.$$

Here $\alpha \geq 0$ balances the squared-wheel-speed surrogate against a positive time weight $\beta > 0$. The expression is a *proposed design problem*, not an optimization solved in the current study. In a deployed controller, wheel acceleration and current constraints should be added, and obstacles should be inflated to account for the chassis footprint and localization uncertainty.

Its components have different roles. The Riemannian quadratic term scores smooth wheel-speed effort; the group-valued state equation preserves correct rigid-body kinematics; and the convex but nonsmooth h constraint enforces actuator speed ceilings. The algebraic analysis in the next

section concerns equilibria and parameter branches of explicitly polynomial dynamics. No claim is made that a Gröbner basis directly solves this full optimal-control problem.

7.1 Constructing a computationally modest planner

The full continuous optimal-control problem may exceed the computational resources available for a competition robot if it is solved from scratch at every control update. A practical first implementation can separate an offline geometric stage from a fast online feedback stage. The offline planner selects a collision-free sequence of desired positions and headings, estimates how far each segment is from obstacles, and assigns conservative timestamps. The online regulator uses the current pose estimate to correct toward the next designated pose. At each step the wheel map converts the selected twist into motor-speed commands and explicitly enforces the actuator ceiling.

The discrete planner can parameterize each interval by a piecewise-constant body twist ξ_j over duration $\Delta t_j > 0$. Its pose recursion is then group-consistent:

$$T_{j+1} = T_j exp\left(\Delta t_j \hat{\xi}_j\right).$$

The optimization variables are the twists and possibly their durations, subject to $h\left(\xi_j\right) \leq r\Omega$. For any fixed pose, the wheel-speed restriction is convex in the twist, but the complete trajectory problem is generally nonconvex because successive group exponentials couple translation and heading and because obstacle avoidance creates nonconvex constraints. Therefore describing the whole planning problem as a convex quadratic program would be incorrect.

The term β in the objective discourages arbitrarily long movement times, whereas the term with α rewards lower squared wheel speeds. An additional discrete smoothness penalty, such as $\gamma \sum_j \| \xi_{j+1} - \xi_j \|_2^2$, can reduce abrupt changes, but it does not replace a physical acceleration or current constraint. The weights should be selected against measurements such as motor temperature, battery voltage, drive error, and completion time rather than on geometric elegance alone.

7.2 Generalization to unequal motors

The exact diamond in Section 5 relies on symmetry and identical angular speed ceilings. For an identified wheel-rate map $\dot{\phi} = B\xi$ with motor-specific bounds $\left|\dot{\phi}_i\right| \leq \Omega_i$, the actual feasible set becomes

$$U_B = \cap_{i=1}^{4} \{\xi: \left|B_i \xi\right| \leq \Omega_i\}.$$

It is still an intersection of strips and hence a convex polyhedron, but it need not have a diamond-shaped translation section or a closed-form expression involving $\left|v_x\right| + \left|v_y\right|$. A

universally feasible scaling of a candidate twist ξ_* is obtained by setting

$$\gamma = min\left(1, \min_{i:\,|B_i\xi_*|>0} \frac{\Omega_i}{|B_i\xi_*|}\right), \qquad \xi_{safe} = \gamma\xi_*.$$

An empty inner index set is interpreted as $\gamma = 1$. The formula is a direct actuator-allocation safeguard, although it sacrifices some directional freedom compared with a constrained optimization that adjusts the translation and yaw components independently. Testing both choices may reveal whether a more sophisticated allocator is worthwhile for the measured chassis.

8. Computational Algebraic Geometry and Bifurcations

8.1 Which parts are polynomial

Setting $\dot{e} = 0$ in the feasible geometric regulator gives $-\lambda(e)e = 0$, whose only equilibrium in the valid logarithm chart is $e = 0$. This conclusion follows immediately because $\lambda(e) > 0$; polynomial elimination is unnecessary. Moreover, the regulator uses a Lie logarithm, absolute values, and a minimum operator, so its global equations are not a single polynomial vector field.

Polynomial techniques can become useful after a mechanically justified slip, friction, dead-zone, or other nonlinear model has been identified. For example, each absolute value may be resolved into sign regions, creating a finite family of polynomial inequalities, but the sign-consistency inequalities must be retained. Squaring away constraints without checking real roots can introduce spurious equilibria. The term *Gröbner basis* should be reserved for an actual polynomial ideal, and real-root conditions should be checked separately (Hajnová and Přibylová).

8.2 An explicit polynomial heading example

To show the algebraic procedure without claiming a physical effect, consider the **illustrative** scalar orientation field

$$\dot{\theta} = -\sin\theta\,(a + b\cos\theta),$$

where $a, b \in R$ are arbitrary model parameters. This field is **not** derived from the X-drive motors or from the controller of Section 6. Lift the trigonometric functions onto the unit circle by setting $c = \cos\theta$ and $s = \sin\theta$. The equilibrium set becomes the real algebraic system

$$s(a + bc) = 0, \qquad c^2 + s^2 - 1 = 0.$$

Elimination of s gives the polynomial factor condition

$$\left(c^2 - 1\right)(a + bc)^2 = 0.$$

The roots $c = 1$ and $c = -1$, both with $s = 0$, are always equilibria. If $b \neq 0$ and $|a/b| < 1$, a further pair appears:

$$c = -a/b, \qquad s = \pm\sqrt{1 - a^2/b^2}.$$

They meet one of the original equilibria at the parameter thresholds $a = \pm b$. The linearization of the heading field at a stationary angle is

$$f'(\theta_*) = -a\cos\theta_* - b\cos(2\theta_*).$$

For $b > 0$ and $-b < a < b$, both $\theta = 0$ and $\theta = \pi$ have negative local derivative and are locally stable, while the two additional roots have derivative $(b^2 - a^2)/b > 0$ and are unstable. If $a > b > 0$, only $\theta = 0$ is stable among the two original equilibria; the extra pair does not exist. These are mathematically verified branch and stability statements for the **toy field only**.

The example explains the role of computational algebraic geometry: polynomialization, elimination, real-root screening, parameter thresholds, and stability testing are separate steps. To argue that the real VEX platform exhibits such branches, one would first need to identify a physical nonlinear heading model, measure its coefficients, and show that the model explains observed behavior. A constructed trigonometric term is not experimental evidence of drivetrain bifurcation.

8.3 Real solutions, multiplicity, and false branches

The elimination polynomial contains the squared factor $(a + bc)^2$. Its exponent records algebraic multiplicity in the resultant; it does **not** mean that the physical circle contains two identical states at the same (c, s). The point with $c = -a/b$ lifts to two different real points only when $|a/b| < 1$ and $s \neq 0$. At $|a/b| = 1$ the two lifts coincide at $s = 0$. For $|a/b| > 1$ the remaining formal values of s are imaginary and cannot represent a real robot heading.

For example, the exact parameter choice $a = 3/5$ and $b = 1$ gives $c = -3/5$ and $s = \pm 4/5$. Together with $(c, s) = (1, 0)$ and $(-1, 0)$, the polynomial equations therefore have four distinct real solutions on the unit circle. At $a = 6/5$ and $b = 1$, the proposed extra root has $c = -6/5$, violates $|c| \leq 1$, and is excluded. These fractions make the branch count independently checkable without a numerical root solver.

For a dynamical system with polynomial vector field $\dot{z} = F(z; \mu)$, a standard algebraic route to candidate nonhyperbolic equilibria introduces $F(z; \mu) = 0$ and $detD_z F(z; \mu) = 0$. Eliminating the state coordinates can reveal parameter values where a zero eigenvalue is possible. This necessary condition is not sufficient to establish a particular bifurcation type: one must also inspect real solutions, degeneracy conditions, and the relevant higher-order derivatives. A Hopf bifurcation involves a pair of imaginary eigenvalues, so its detection requires an appropriate characteristic-polynomial condition rather than just a vanishing determinant (Hajnová and Přibylová).

For saturated systems the vector field is only piecewise differentiable. One must distinguish an equilibrium inside one smooth saturation region from a candidate on the switching boundary; ordinary smooth Jacobian criteria cannot be applied blindly across that boundary. This is another reason to treat Gröbner bases as a targeted analysis method within a larger hybrid or semialgebraic model, not as a substitute for a complete nonlinear controller design.

9. Reproducible Ideal-Model Numerical Study

To check the implementation of the theoretical inequalities, a deterministic forward-Euler integration of $\dot{e} = -\lambda(e)e$ was performed. This is a simulation of the **ideal error dynamics**, not a motor, contact, or sensor simulation. The chosen illustrative parameters were $r = 0.05$ m, $R = 0.20$ m, $\Omega = 20$ rad/s, and $k = 1.2\,s^{-1}$. Integration used $\Delta t = 0.01$ s for 8.00 s, with the target pose set to the identity. No parameters were measured from a robot.

The three initial Lie-logarithmic errors were $e_0^{(A)} = (1.2, -0.8, 0.6)$, $e_0^{(B)} = (-1.4, 1.1, -1.0)$, and $e_0^{(C)} = (0.5, 1.6, 1.2)$, where translations are in meters and rotation in radians. At each step, the code computed λ from the exact wheel-speed gauge, commanded $\xi_c = \lambda e$, and applied $e_{n+1} = \left(1 - \lambda_n \Delta t\right) e_n$. The wheel-speed usage was checked from $max_i \left| \left(J\xi_c\right)_i \right| / (r\Omega)$.

Case	$\| e_0 \|_2$	$\| e(8\,s) \|_2$	Maximum normalized wheel-speed usage	First unsaturated time
A	1.56205	0.00012638	1.00000	0.71 s
B	2.04206	0.00021740	1.00000	1.14 s
C	2.06155	0.00018675	1.00000	0.90 s

Table 1. Outcomes of ideal-model numerical integration. The mixed-unit Euclidean norm is an analysis coordinate norm, not a physical distance. The normalized wheel-speed usage of 1.0 means the ideal motor-speed ceiling was reached but not exceeded at the sampled instants.

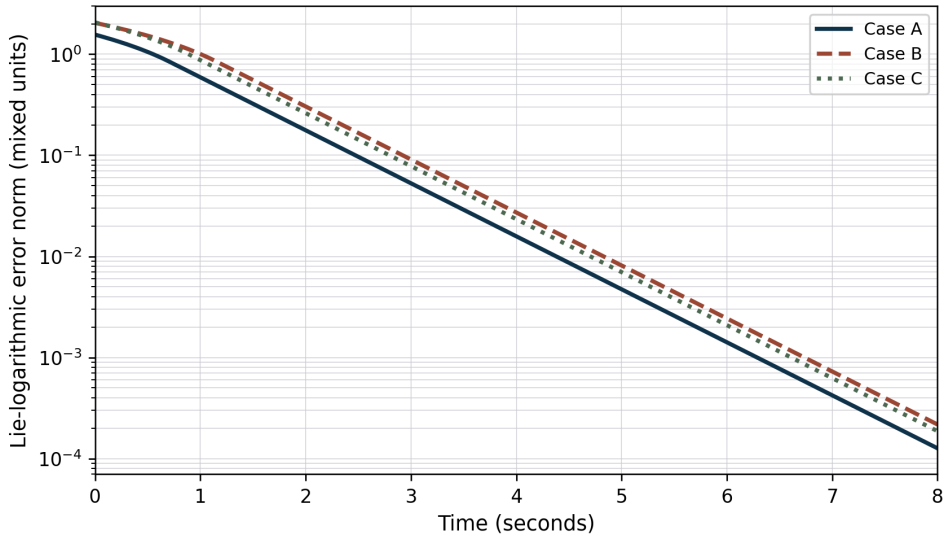


Figure 2. Magnitude of simulated Lie-logarithmic error for three initial conditions. The initial speed-limited interval is followed by exponential decay after motor commands fall below the speed limit.

The numerical traces are consistent with Proposition 4: each idealized error decreases toward zero, and the modeled wheel commands stay within their imposed speed ceiling. They do not establish performance superiority over PID, real-time computational cost, resistance to wheel slip, or robustness to sensor noise. The figures can be reproduced directly from the update equation and the listed parameter values. A meaningful experimental result will require actual pose ground truth or independent visual tracking; encoder integration alone cannot establish absolute position accuracy.

10. Proposed VEX V5 Implementation and Validation

10.1 Estimation and control architecture

A realistic engineering architecture separates the outer pose controller from motor regulation. The input layer records the four wheel encoder increments and inertial yaw; a calibrated forward-kinematic map estimates body displacement; and a world-frame pose estimate is updated using the current orientation. The controller then computes a stationary target error on $SE(2)$, applies the actuator-feasible normalization, and maps the resulting twist back into four wheel-speed requests. Each wheel's actual motor-speed control runs as an inner loop. Motor encoders and inertial sensors have interfaces documented by VEX, but sensor bandwidth, drift, and failure behavior must be established for the installed devices (VEX Robotics, "V5 Motor Overview"; VEX Robotics, "Using the Inertial Sensor").

With encoder increments $\Delta\phi$, a first-order body displacement estimate is

$$\Delta\hat{\zeta} = r\left(J^{\top}J\right)^{-1}J^{\top}\Delta\phi.$$

A practical estimator can use inertial yaw for its orientation update, preserving the encoder-derived translation in the robot frame and rotating it into world coordinates. Sensor fusion should unwrap heading consistently and estimate or test gyroscope bias. If external vision becomes available, it can provide intermittent absolute pose corrections; that option requires known camera calibration, observable landmarks, latency measurement, and a legal hardware configuration for the applicable competition.

10.2 Observability and drift with the existing sensors

Encoders and an inertial sensor answer different measurement questions. An encoder estimates how far a driven wheel has turned; the inertial sensor reports orientation-related quantities. Neither directly measures absolute world position on a featureless field. In particular, the encoder-inertial measurements generated by a robot pose (p, θ) are unchanged if every world position is replaced by $p + c$ for the same constant translation c. That translation symmetry shows why absolute position cannot be inferred from relative-motion measurements alone without a known starting pose or an external position reference.

For a small time interval, a first-order world displacement estimate takes the form

$$\hat{p}_{n+1} = \hat{p}_n + R\left(\hat{\theta}_n\right)\Delta\hat{\zeta}_{xy,n},$$

where $\Delta\hat{\zeta}_{xy,n}$ is the translational portion of the encoder-based increment. If the heading estimate has error ε_θ, the direction used to rotate an otherwise correct body displacement is also wrong. For small ε_θ, the first-order position increment error is approximately $R(\theta)S\,\Delta\zeta_{xy}\,\varepsilon_\theta$, where $S = [0 \;\; -1 \; 1 \; 0]$. This couples orientation bias to translational drift even if the wheel-radius calibration were perfect.

Slip adds another unobserved component: wheel angles can change while the corresponding chassis displacement is smaller than the rolling model predicts. One should not silently treat encoder residuals as independent Gaussian noise when a wheel loses traction. A traction event can produce systematic, correlated errors that persist through several estimator updates. For this reason, an estimator should maintain an uncertainty measure and flag impossible or unusually inconsistent sensor combinations rather than merely integrating every encoder change without scrutiny.

If the competition rules and budget permit a camera observing known field landmarks, a measurement model $z = h(T, \ell) + \nu$ can be added, where ℓ represents fixed landmark geometry and ν is observation error. The usefulness of that correction depends on whether the landmarks are visible and geometrically informative. A single ambiguous feature cannot necessarily identify all three planar pose coordinates, and delayed measurements must be associated with the robot state at the image timestamp. These are observability and calibration questions, not arguments that adding a camera automatically improves every run.

10.3 A measured performance protocol

An experimental comparison should name its independent reference. The controller's own odometry cannot simultaneously serve as the sole ground truth for its position error. A tape-measured target may suffice for static endpoint accuracy; continuous trajectory accuracy is better measured by overhead video, calibrated field markers, or another independently validated tracking system. The protocol should separate systematic bias from run-to-run dispersion.

For a set of N trials, let $p_f^{(i)}$ denote the independently measured final position and p_d the target. Mean endpoint distance error and root-mean-square endpoint error can be defined by

$$\bar{e}_p = \frac{1}{N}\sum_{i=1}^{N} \| p_f^{(i)} - p_d \|_2, \qquad RMSE_p = \sqrt{\frac{1}{N}\sum_{i=1}^{N} \| p_f^{(i)} - p_d \|_2^2}.$$

Settling time must be defined using a positional and angular tolerance band that the robot remains inside for a stated duration. Constraint violations should be reported both as the number of events and as the largest amount by which a motor-speed limit was exceeded. A comparison with PID is meaningful only if both controllers use the same hardware, pose estimator, saturation limits, trial locations, and allowed tuning time.

10.4 Validation sequence

The following work is necessary before using the theoretical conclusions as claims about a functioning robot:

1. **Identify the geometry.** Measure wheel radius, axis directions, effective yaw lever arm, and motor signs. Use isolated wheel and chassis motions to fit the empirical wheel Jacobian and quantify its residuals.
2. **Validate odometry.** Run straight, lateral, diagonal, turning, and simultaneous translation-plus-rotation trials. Compare encoder-plus-inertial pose estimates with independently measured displacement and heading.
3. **Measure actuator feasibility.** Record wheel speeds, battery voltage, motor current, and commanded twist. Determine whether the assumed uniform rate bound and ideal kinematic map predict actual clipping and motion.
4. **Test stationary regulation.** Sample starting positions and headings, including cases near the logarithm branch threshold but remaining in its valid chart. Measure settling time, overshoot, final position error, and failures.
5. **Compare fair baselines.** Tune a conventional field-centric PID or holonomic controller using the same sensors, actuator limits, paths, and trial conditions. Report distributions across repeated trials rather than selected successful runs.
6. **Evaluate disturbances.** Repeat tests with altered wheel friction, additional payload, deliberate perturbations, IMU bias, and optional vision loss. State explicitly when the original stability proof no longer models the test conditions.

Each trial should retain time-stamped wheel commands, wheel encoder readings, inertial measurements, estimated pose, independently assessed ground truth, and test conditions. Any paper revision reporting experimental results should define sample counts, uncertainty, calibration procedures, and criteria for success before analyzing the data. A gain chosen after viewing the best run is not an unbiased comparative test.

11. Discussion

The three mathematical approaches serve different purposes and should not be collapsed into a single supposed universal controller. Riemannian geometry supplies a smooth quadratic cost. In this symmetric case, its flatness is itself an instructive finding: more sophisticated differential geometry does not automatically produce a different path. The Lie-group model supplies geometrically meaningful pose errors and a concise local proof of convergence. Computational algebraic geometry supplies methods for enumerating roots and studying branches **after** a legitimate polynomial dynamic model has been established.

There are three limits on the present inference. First, motor-speed feasibility is not motor-torque or traction feasibility; acceleration can violate physical constraints while every commanded wheel angular speed remains within its nominal bound. Second, ideal inverse kinematics and motor feedback do not supply accurate global localization when wheels slip. Third, an equilibrium computation for an illustrative orientation field cannot be used to infer bifurcations in a different physical closed loop. These limits directly shape the next experimental program.

The main useful synthesis is the separation of metric-based planning, polyhedral velocity feasibility, and group-consistent pose feedback. The wheel constraints should be enforced as hard inequalities, while the quadratic cost can serve as a secondary preference among feasible commands. Such a separation is especially useful on a small onboard computer because the stationary baseline law uses only elementary operations after pose estimation. The complexity of a moving-obstacle or moving-target planner remains open.

12. Conclusion

For an ideal symmetric four-wheel X-drive, the wheel map has full chassis-twist rank, the squared-wheel-speed Riemannian metric is flat, and the exact wheel-speed region has a diamond-shaped translation section. A normalized logarithmic error controller respects those wheel-speed bounds and converges locally to a stationary pose under ideal sensing and velocity realization. A separate polynomial heading example demonstrates how algebraic elimination can expose equilibrium branches, while clarifying that such branches have not been shown to exist in the robot.

The next step is system identification and controlled experiments on the actual VEX V5 platform. Until then, these ideal-model results do not establish improved competition performance.

Works Cited

Balkcom, Devin J., et al. “Time-Optimal Trajectories for an Omni-Directional Vehicle.” *The International Journal of Robotics Research*, vol. 25, no. 10, 2006, pp. 985-999. Carnegie Mellon University Robotics Institute, https://publications.ri.cmu.edu/time-optimal-trajectories-for-an-omni-directional-vehicle.

Hajnová, Veronika, and Lenka Přibylová. “Bifurcation Manifolds in Predator-Prey Models Computed by Gröbner Basis Method.” *Mathematical Biosciences*, vol. 312, 2019, pp. 1-7. https://doi.org/10.1016/j.mbs.2019.03.008.

Lynch, Kevin M., and Frank C. Park. *Modern Robotics: Mechanics, Planning, and Control*. Cambridge University Press, 2017. Modern Robotics, https://modernrobotics.northwestern.edu/nu-gm-book-resource/13-2-omnidirectional-wheeled-mobile-robots-part-2-of-2/.

Purwin, Oliver, and Raffaello D’Andrea. “Trajectory Generation and Control for Four Wheeled Omnidirectional Vehicles.” *Robotics and Autonomous Systems*, vol. 54, no. 1, 2006, pp. 13-22. https://doi.org/10.1016/j.robot.2005.10.002.

Teng, Sangli, et al. “Lie Algebraic Cost Function Design for Control on Lie Groups.” *arXiv*, 2022, https://arxiv.org/abs/2204.09177.

VEX Robotics. “Using the Inertial Sensor with VEX V5.” *VEX Library*, https://kb.vex.com/hc/en-us/articles/360037382272-Using-the-Inertial-Sensor-with-VEX-V5. Accessed 2 Nov. 2026.

VEX Robotics. “V5 Motor Overview.” *VEX Library*, https://kb.vex.com/hc/en-us/articles/360035591332-V5-Motor-Overview. Accessed 2 Nov. 2026.