

Symplectic Matrices

From phase space to optimal control

A rotation preserves lengths. A symplectic map preserves the pairing between a coordinate and its conjugate momentum. This is why the same matrix identity appears in planetary motion, optical systems, and the Riccati equation behind LQR.

1. What can change while the essential geometry stays?

Imagine a cloud of particles described by position q and momentum p . Stretch every position by 2 and compress every momentum by 2:

$$\begin{pmatrix} q' \\ p' \end{pmatrix} = \underbrace{\begin{pmatrix} 2 & 0 \\ 0 & \frac{1}{2} \end{pmatrix}}_S \begin{pmatrix} q \\ p \end{pmatrix}. \quad (1)$$

The circle becomes an ellipse. Lengths change, and the oscillator energy $H = \frac{q^2 + p^2}{2}$ generally changes. Yet the oriented area between any two displacements is unchanged. The shear below has the same property.




Figure 1: Three symplectic maps in one degree of freedom. Dashed: original unit circle. Filled: its image. Every filled region has area π , but not the same shape.

Now write $z = (q, p)$, with $q, p \in \mathbb{R}^n$, **in that order**, and define

$$J = \begin{pmatrix} 0 & I_n \\ -I_n & 0 \end{pmatrix}, \quad \omega(u, v) = u^T J v = u_q^T v_p - u_p^T v_q. \quad (2)$$

The alternating bilinear form ω measures canonical pairing; in one degree of freedom it is the signed area of the parallelogram spanned by u, v . A real $2n \times 2n$ matrix S is **symplectic** when

$$S^T J S = J. \quad (3)$$

Indeed, $\omega(Su, Sv) = u^T S^T J S v = \omega(u, v)$ for every pair. In higher dimensions this preserves the **sum** of canonical pairings, not each coordinate-plane projection separately.

Reading route. Pages 2–3 establish the algebra and the connection to flow. Pages 4–5 use it in physics. Pages 6–9 explain why optimal control needs it. Page 10 gives reproducible checks and questions. Only linear algebra and basic ODEs are assumed. All matrices are real unless quantum operators are mentioned.

2. The algebra: stronger than volume preservation

The set of matrices satisfying Equation 3 is written $\text{Sp}(2n, \mathbb{R})$. Since $J^T = -J$ and $J^2 = -I$, the defining identity gives

$$S^{-1} = -JS^T J. \quad (4)$$

This proves invertibility without assuming it. Products and inverses are symplectic: $(S_1 S_2)^T J (S_1 S_2) = S_2^T J S_2 = J$.

2.1. What does the condition say about the blocks?

For $S = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ with $n \times n$ blocks, multiplication gives

$$a^T c = c^T a, \quad b^T d = d^T b, \quad a^T d - c^T b = I. \quad (5)$$

Useful building blocks are

$$\begin{pmatrix} I & B \\ 0 & I \end{pmatrix} \quad (B = B^T), \quad \begin{pmatrix} I & 0 \\ C & I \end{pmatrix} \quad (C = C^T), \quad \begin{pmatrix} A & 0 \\ 0 & A^{-T} \end{pmatrix} \quad (A \text{ invertible}). \quad (6)$$

Here $A^{-T} = (A^T)^{-1}$. A general position coordinate change must transform momentum by its **inverse transpose**, not by the same matrix. Otherwise the canonical pairing changes. These building blocks are not claimed to be a single global factorization chart for every symplectic matrix.

2.2. Why the determinant is +1, not merely ± 1

Taking determinants in Equation 3 gives only $\det(S)^2 = 1$. To fix the sign, write the two-form $\omega = \sum_i dq_i \wedge dp_i$. Its n -fold wedge product is a nonzero constant multiple of oriented phase-space volume. Because S preserves ω , it preserves this top-degree form too; hence $\det(S) = 1$. This is an algebraic statement about the map, not an assumption of stability.

For a 2×2 matrix, $S^T J S = \det(S) J$, so determinant one is also **sufficient**. For $2n \geq 4$ it is not. In ordering (q_1, q_2, p_1, p_2) , take

$$D = \text{diag}\left(2, 1, 1, \frac{1}{2}\right). \quad (7)$$

Its determinant is one, but

$$\omega(D e_1, D e_3) = 2 \neq 1 = \omega(e_1, e_3). \quad (8)$$

One canonical plane expands while another contracts. Total volume survives; the canonical pairings do not.

2.3. Reciprocal eigenvalues, not automatic stability

Equation 4 makes S^{-1} similar to S^T . Thus the spectrum is closed under $\lambda \mapsto \frac{1}{\lambda}$, including algebraic multiplicities. Real entries also give complex-conjugate pairs. A generic off-axis, off-unit-circle eigenvalue therefore belongs to a quartet

$$\lambda, \quad \bar{\lambda}, \quad \frac{1}{\lambda}, \quad \frac{1}{\bar{\lambda}}. \quad (9)$$

Coincidences reduce the quartet. The squeeze on page 1 has eigenvalues $2, \frac{1}{2}$: symplectic maps can be unstable. A whole symplectic map cannot be strictly contracting in all directions; some expansion must compensate contraction.

Do not substitute a weaker test. Orthogonality preserves Euclidean lengths; determinant one preserves volume; symplecticity preserves canonical pairings. They are different conditions. Rotations of a single (q, p) plane happen to be both orthogonal and symplectic.

3. Hamiltonian matrices generate symplectic motion

For a smooth physical Hamiltonian $H(q, p)$, Hamilton's equations are

$$\dot{q} = \nabla_p H, \quad \dot{p} = -\nabla_q H, \quad \dot{z} = J \nabla H(z). \quad (10)$$

For a quadratic Hamiltonian $H(z) = \frac{1}{2} z^T K z$ with $K = K^T$,

$$\dot{z} = Fz, \quad F = JK, \quad F^T J + JF = 0. \quad (11)$$

A matrix satisfying the last identity is called **Hamiltonian**. Conversely, $K = -JF$ is symmetric whenever that identity holds. In block form,

$$F = \begin{pmatrix} A & B \\ C & -A^T \end{pmatrix}, \quad B = B^T, \quad C = C^T. \quad (12)$$

This is a condition on an infinitesimal generator, not on a finite step.

3.1. A proof that explains the distinction

Let $S(t) = e^{tF}$, so $\dot{S} = FS$ and $S(0) = I$. Then

$$\frac{d}{dt}(S^T JS) = S^{T(F^T J + JF)} S = 0. \quad (13)$$

The initial value is J , so $S(t)^T JS(t) = J$. The same proof works for a time-dependent Hamiltonian generator $F(t)$, using its fundamental solution rather than a single matrix exponential.

Hamiltonian matrix $\rightarrow$ differential equation $\rightarrow$ symplectic transition. $F^T J + JF = 0$ describes the generator; $S^T JS = J$ describes the flow. Confusing these identities is like confusing angular velocity with a rotation.

3.2. The harmonic oscillator: exact motion

For mass $m > 0$ and stiffness $k > 0$,

$$H = \frac{p^2}{2m} + k \frac{q^2}{2}, \quad F = \begin{pmatrix} 0 & \frac{1}{m} \\ -k & 0 \end{pmatrix}. \quad (14)$$

With $\nu = \sqrt{\frac{k}{m}}$, exponentiation gives

$$S(t) = \begin{pmatrix} \cos(\nu t) & \frac{\sin(\nu t)}{m\nu} \\ -m\nu \sin(\nu t) & \cos(\nu t) \end{pmatrix}. \quad (15)$$

This map is symplectic and preserves this particular H . For $m = k = 1$ it is a clockwise Euclidean rotation in (q, p) ; for other parameters its trajectories are ellipses. Energy preservation is an additional property of this flow, not a consequence of symplecticity of an arbitrary matrix.

3.3. Nonlinear physics: the Jacobian is symplectic

A nonlinear flow φ_t is not a matrix. Its tangent map $S(t) = D\varphi_{t(z_0)}$ satisfies

$$\dot{S} = J \nabla^2 H(\varphi_{t(z_0)}) S, \quad S(0) = I. \quad (16)$$

Since the Hessian is symmetric, the same argument proves $S(t)^T JS(t) = J$. This local statement is what “the nonlinear flow is symplectic” means. For an autonomous H , $\dot{H} = \nabla H^T J \nabla H = 0$ as well. With explicit time dependence, symplecticity survives even though $\dot{H} = \partial_t H$ need not vanish.

Where it stops. Add linear drag to a unit-mass oscillator: $F = \begin{pmatrix} 0 & 1 \\ -k & -\gamma \end{pmatrix}$ with $\gamma > 0$. Then $F^T J + JF = -\gamma J$, and $S(t)^T JS(t) = e^{-\gamma t} J$. Dissipation contracts phase area. The damped physical system is not symplectic in these canonical coordinates.

4. Why numerical physics cares about this identity

For the normalized oscillator $\dot{q} = p$, $\dot{p} = -q$, ordinary explicit Euler updates both coordinates from the old state:

$$E_h = \begin{pmatrix} 1 & h \\ -h & 1 \end{pmatrix}, \quad \det(E_h) = 1 + h^2. \quad (17)$$

Every step expands area, and, more strongly for this example,

$$H_{j+1} = (1 + h^2)H_j. \quad (18)$$

The simulated orbit spirals outward even though the real oscillator does not.

Instead **kick** momentum first, then **drift** position using the new momentum:

$$p_{j+1} = p_j - hq_j, \quad q_{j+1} = q_j + hp_{j+1}. \quad (19)$$

This is symplectic Euler, with matrix

$$S_h = \begin{pmatrix} 1 - h^2 & h \\ -h & 1 \end{pmatrix} = \begin{pmatrix} 1 & h \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 1 & 0 \\ -h & 1 \end{pmatrix}, \quad S_h^T J S_h = J. \quad (20)$$

Each factor is a symplectic shear. The method preserves exactly the quadratic quantity

$$\tilde{H}_h = \frac{1}{2}(q^2 + p^2 - hqp), \quad (21)$$

as direct substitution shows. This is **not** the original energy.




Figure 2: Starting at $(q,p) = (1,0)$, with $h = 0.2$. Left: the first 100 explicit Euler steps versus the exact circle and 300 symplectic steps. Right: energy over all 300 steps. The logarithmic axis makes Euler's exponential drift visible.

When $|h| < 2$, Equation 21 is positive definite, so its constant level sets bound the numerical trajectory. With $h = 0.2$, $H_0 = \frac{1}{2}$, it implies

$$\frac{1}{1.1} \leq \frac{H_j}{H_0} \leq \frac{1}{0.9}. \quad (22)$$

At step 300, explicit Euler gives $\frac{H_{300}}{H_0} = 1.04^{300} \approx 128825$. The symplectic method oscillates within the stated bounds. But for $|h| > 2$ the eigenvalues of S_h leave the unit circle: **symplectic does not mean stable for any step size**. At $|h| = 2$ a defective boundary case permits growth.

For general separable mechanics $H = T(p) + V(q)$, use the second-order kick-drift-kick (Störmer-Verlet) scheme:

$$p_{j+\frac{1}{2}} = p_j - \frac{h}{2} \nabla V(q_j), \quad q_{j+1} = q_j + h \nabla T(p_{j+\frac{1}{2}}), \quad (23)$$

$$p_{j+1} = p_{j+\frac{1}{2}} - \frac{h}{2} \nabla V(q_{j+1}). \quad (24)$$

It composes exact subflows, so it is symplectic. Under appropriate smoothness, step-size, and bounded-orbit assumptions, backward error analysis explains excellent long-time energy behavior via a nearby modified Hamiltonian—not exact conservation of the original energy or an exact trajectory [1], [2].

5. Three physical settings, one canonical structure

5.1. Classical mechanics and molecular dynamics

For coupled masses, let $M = M^T \succ 0$ and $K = K^T \succeq 0$:

$$H(q, p) = \frac{1}{2} p^T M^{-1} p + \frac{1}{2} q^T K q, \quad F = \begin{pmatrix} 0 & M^{-1} \\ -K & 0 \end{pmatrix}. \quad (25)$$

The off-diagonal blocks are symmetric, hence F is Hamiltonian. In nonlinear molecular or celestial mechanics replace the quadratic potential by $V(q)$. Symplectic integration preserves the canonical structure over many oscillations or orbits. It does not remove force-model error, close-encounter difficulties, phase error, or the need to choose an adequate step size.

An important coordinate caveat: canonical momentum is $p = M\dot{q}$ in this simple model. A vector of positions and velocities is not automatically in standard symplectic coordinates. Likewise, charged-particle canonical momentum contains the vector potential; mechanical momentum alone changes the form.

5.2. Paraxial ray optics and accelerator transfer matrices

In air (refractive index normalized to one), use transverse displacement q and normalized canonical ray momentum p , the small ray angle. A drift of length ℓ and a thin focusing lens of focal length $f \neq 0$ have normalized matrices

$$D(\ell) = \begin{pmatrix} 1 & \ell \\ 0 & 1 \end{pmatrix}, \quad L(f) = \begin{pmatrix} 1 & 0 \\ -\frac{1}{f} & 1 \end{pmatrix}. \quad (26)$$

Their product is symplectic. The lens changes momentum according to position; the drift changes position according to momentum. These are precisely the shears used in kick-drift integration [3].

For a centered ensemble with covariance Σ , a linear map gives

$$\Sigma' = S\Sigma S^T, \quad \det(\Sigma') = \det(\Sigma). \quad (27)$$

Thus the rms phase-space emittance $\varepsilon = \sqrt{\det(\Sigma)}$ is invariant in one canonical plane. This does not say that beam width is fixed: a narrow beam can have a wider angular spread. Coupled multi-plane systems can exchange projected emittances. At refractive-index interfaces, angle alone is not the canonical momentum; index factors must be included before applying this test.

5.3. Quadratic quantum mechanics: preserving commutators

Let $\hat{z} = (\hat{q}, \hat{p})$ obey the canonical commutation relations

$$[\hat{z}_i, \hat{z}_j] = i\hbar J_{ij}. \quad (28)$$

For a real linear transformation $\hat{z}' = S\hat{z}$,

$$[\hat{z}'_i, \hat{z}'_j] = i\hbar (SJS^T)_{ij}. \quad (29)$$

This equals the original commutator matrix exactly when S is symplectic ($SJS^T = J$ is equivalent to Equation 3). Evolution under a quadratic Hamiltonian therefore produces symplectic transformations of quadratures. The squeeze $q' = aq$, $p' = \frac{p}{a}$ narrows one uncertainty while enlarging the other; it is not a way to violate the uncertainty principle. Quantum states evolve unitarily; their **real quadrature coordinates** transform symplectically. These are different levels of description [4].

The applications differ, but the question is identical: does the evolution preserve the canonical relation between each coordinate and its conjugate? “Energy preserving” and “volume preserving” do not capture the whole contract.

6. Why a stabilizing controller creates a Hamiltonian system

Consider continuous-time LQR with $x \in \mathbb{R}^n$, $u \in \mathbb{R}^m$:

$$\dot{x} = Ax + Bu, \quad \mathcal{J} = \frac{1}{2} \int_0^\infty (x^T Q x + u^T R u) dt, \quad (30)$$

where $Q = Q^T \succeq 0$ and $R = R^T \succ 0$. Assume (A, B) is stabilizable and $(Q^{\frac{1}{2}}, A)$ is detectable. These sufficient hypotheses give a unique positive-semidefinite stabilizing Riccati solution. It need not be positive definite when some stable directions carry no state cost [5].

Introduce a costate λ , the marginal future cost of changing x . The optimal-control Hamiltonian (a scalar function, not a physical energy) is

$$\mathcal{H}(x, u, \lambda) = \frac{1}{2} x^T Q x + \frac{1}{2} u^T R u + \lambda^T (Ax + Bu). \quad (31)$$

Stationarity in u gives $Ru + B^T \lambda = 0$. Set $G = BR^{-1}B^T$; after eliminating u , the reduced Hamiltonian is

$$\mathcal{H}_*(x, \lambda) = \frac{1}{2} x^T Q x + \lambda^T A x - \frac{1}{2} \lambda^T G \lambda. \quad (32)$$

Its canonical equations are

$$\dot{x} = Ax - G\lambda, \quad \dot{\lambda} = -Qx - A^T \lambda, \quad (33)$$

or

$$\frac{d}{dt} \begin{pmatrix} x \\ \lambda \end{pmatrix} = \underbrace{\begin{pmatrix} A & -G \\ -Q & -A^T \end{pmatrix}}_{F_L} \begin{pmatrix} x \\ \lambda \end{pmatrix}. \quad (34)$$

Because G, Q are symmetric, $F_L^T J + J F_L = 0$. The full state–costate transition e^{tF_L} is therefore symplectic. The costate is **not** physical momentum, but it plays the same canonical mathematical role.

6.1. The Riccati equation is an invariant-graph condition

Suppose the value function is $V(x) = \frac{1}{2} x^T P x$, with $P = P^T$. Then $\lambda = \nabla V = Px$. For this relation to persist under Equation 34,

$$P\dot{x} = \dot{\lambda}, \quad P(A - GP)x = -Qx - A^T Px. \quad (35)$$

Because this must hold for every x ,

$$\boxed{A^T P + P A - P G P + Q = 0}. \quad (36)$$

This is the continuous-time algebraic Riccati equation (CARE). Equivalently,

$$F_L \begin{pmatrix} I \\ P \end{pmatrix} = \begin{pmatrix} I \\ P \end{pmatrix} (A - GP). \quad (37)$$

The n -dimensional graph $\lambda = Px$ is invariant, and its internal dynamics are $\dot{x} = (A - GP)x$. The optimal feedback is $u = -R^{-1}B^T Px$.

There is no contradiction between symplectic flow and stabilization. The **full** $2n$ -dimensional state–costate flow has contracting and expanding directions. Optimality selects the n -dimensional stable graph. The physical closed-loop state evolves on that graph and can contract.

The Hamiltonian generator has eigenvalue symmetry about zero (and conjugation for real data). With no imaginary-axis eigenvalues, its stable and unstable subspaces both have dimension n . The LQR assumptions above identify the stable one with the graph of the stabilizing P .

7. How stable subspaces turn into a feedback gain

Let columns of $Z = \begin{pmatrix} U \\ V \end{pmatrix}$ span the stable invariant subspace of F_L , with $U, V \in \mathbb{R}^{n \times n}$. If U is nonsingular, then $x = Uc$, $\lambda = Vc$, so

$$P = VU^{-1}. \quad (38)$$

Numerically, solve $PU = V$ rather than explicitly forming an inverse. Use an ordered Schur decomposition (or a suitable pencil/QZ method), not a raw collection of individually computed eigenvectors near multiple eigenvalues [6], [7].

7.1. Why is the result symmetric?

For any stable vectors z_1, z_2 , both $e^{tF_L} z_i$ tend to zero. Their canonical pairing is constant, so

$$z_1^T J z_2 = (e^{tF_L} z_1)^T J (e^{tF_L} z_2) \rightarrow 0. \quad (39)$$

Hence $Z^T J Z = U^T V - V^T U = 0$. Multiplying by U^{-T} and U^{-1} gives $P = P^T$. An n -dimensional subspace with all such pairings zero is called **Lagrangian**. Thus symmetry of the Riccati matrix is a geometric fact.

7.2. A scalar unstable plant: choose the correct root

Take $\dot{x} = x + u$, $Q = R = 1$. Then

$$F_L = \begin{pmatrix} 1 & -1 \\ -1 & -1 \end{pmatrix}, \quad 2P - P^2 + 1 = 0, \quad P_+ = 1 + \sqrt{2}, \quad P_- = 1 - \sqrt{2}. \quad (40)$$

Both roots have zero Riccati residual. But $A - GP_+ = -\sqrt{2}$ is stable, whereas $A - GP_- = +\sqrt{2}$ is unstable. The correct branch matters more than merely solving the polynomial.




Figure 3: Left: the stable and unstable invariant graphs for the scalar plant above; dots on the stable line approach the origin. Right: a different plant, the pure integrator, whose finite-horizon value curvature grows with the remaining horizon.

7.3. A two-state example with a checkable exact answer

For the double integrator,

$$A = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}, \quad B = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad Q = I_2, \quad R = 1, \quad (41)$$

$$P = \begin{pmatrix} \sqrt{3} & 1 \\ 1 & \sqrt{3} \end{pmatrix}, \quad K = B^T P = (1 \quad \sqrt{3}). \quad (42)$$

Substitution verifies Equation 36. Closed-loop poles are $-\frac{\sqrt{3}}{2} \pm \frac{i}{2}$, strictly in the left half-plane. The script checks this answer against Octave's care, the graph identity Equation 37, and the symplectic identity for a nonzero-time transition.

8. Finite-horizon control: Riccati as a moving graph

Now minimize the running LQR cost on $[0, T]$ plus $\frac{1}{2}x(T)^T P_f x(T)$, with $P_f = P_f^T \succeq 0$. The terminal condition is $\lambda(T) = P_f x(T)$, and the time-dependent value is $V(t, x) = \frac{1}{2}x^T P(t)x$. Differentiating $\lambda(t) = P(t)x(t)$ gives

$$\dot{P}x + P(A - GP)x = -Qx - A^T Px, \quad (43)$$

so

$$-\dot{P} = A^T P + PA - PGP + Q, \quad P(T) = P_f. \quad (44)$$

This equation is integrated **backward from T**. Reversing that direction without changing the sign solves a different problem.

8.1. A nonlinear matrix equation hidden inside a linear flow

Let a symplectic transition from time t_0 to t_1 have blocks

$$\begin{pmatrix} x_1 \\ \lambda_1 \end{pmatrix} = \begin{pmatrix} a & b \\ c & d \end{pmatrix} \begin{pmatrix} x_0 \\ \lambda_0 \end{pmatrix}. \quad (45)$$

If $\lambda_0 = P_0 x_0$, then

$$x_1 = (a + bP_0)x_0, \quad \lambda_1 = (c + dP_0)x_0. \quad (46)$$

Whenever $a + bP_0$ is invertible, the propagated graph is

$$P_1 = (c + dP_0)(a + bP_0)^{-1}. \quad (47)$$

This fractional-linear (Möbius) update is the graph-coordinate version of a linear state–costate transition. A symplectic map takes Lagrangian subspaces to Lagrangian subspaces, so P_1 remains symmetric whenever this chart exists. Symplecticity **alone** does not guarantee positive semidefiniteness for arbitrary maps, time directions, or initial graphs.

To enforce the terminal condition in the constant-coefficient case, propagate **from T back to t**:

$$\begin{pmatrix} X(t) \\ Y(t) \end{pmatrix} = e^{(t-T)F_L} \begin{pmatrix} I \\ P_f \end{pmatrix}, \quad P(t) = Y(t)X(t)^{-1}. \quad (48)$$

The sign $t - T \leq 0$ is essential. If the top block loses invertibility, the subspace may remain well-defined even though this graph coordinate becomes singular. For standard finite-horizon LQR with the stated nonnegative costs, the value remains finite on a finite interval; graph singularities are more general Riccati phenomena, not inevitable LQR blow-up.

8.2. A complete scalar finite-horizon calculation

For $\dot{x} = u$, $Q = R = 1$, and $P_f = 0$, set $\tau = T - t \geq 0$. Here $F_L = \begin{pmatrix} 0 & -1 \\ -1 & 0 \end{pmatrix}$ and $F_L^2 = I$, giving

$$e^{-\tau F_L} = \begin{pmatrix} \cosh(\tau) & \sinh(\tau) \\ \sinh(\tau) & \cosh(\tau) \end{pmatrix}. \quad (49)$$

Apply it to the terminal graph $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$:

$$X = \cosh(\tau), \quad Y = \sinh(\tau), \quad P(t) = \tanh(T - t). \quad (50)$$

Indeed $\dot{P} = P^2 - 1$, $P(T) = 0$, and $P(t) \rightarrow 1$ as the remaining horizon grows. The corresponding feedback is $u = -\tanh(T - t)x$. With no terminal penalty, control weakens as the endpoint approaches because there is less future state cost left to avoid. This is the curve on page 7, right.

Numerical limitation. Long products of state–costate transitions contain both growing and decaying modes. Forming one huge exponential and dividing its blocks can lose accuracy. Subspace reorthogonalization, graph reinitialization, or specialized Riccati methods are preferable for difficult long horizons.

9. Discrete-time control: symplectic matrices and pencils

For $x_{k+1} = Ax_k + Bu_k$ with stage cost $\frac{1}{2}(x_k^T Q x_k + u_k^T R u_k)$, the optimality conditions are

$$u_k = -R^{-1}B^T\lambda_{k+1}, \quad \lambda_k = Qx_k + A^T\lambda_{k+1}. \quad (51)$$

Again put $G = BR^{-1}B^T$. Then

$$x_{k+1} = Ax_k - G\lambda_{k+1}. \quad (52)$$

Notice the mixed time indices: replacing λ_{k+1} by λ_k does not give the correct discrete optimality equations.

9.1. When A is invertible: an ordinary symplectic transfer

Eliminate $\lambda_{k+1} = A^{-T}(\lambda_k - Qx_k)$ to obtain

$$\begin{pmatrix} x_{k+1} \\ \lambda_{k+1} \end{pmatrix} = S_d \begin{pmatrix} x_k \\ \lambda_k \end{pmatrix}, \quad (53)$$

$$S_d = \begin{pmatrix} A + GA^{-T}Q & -GA^{-T} \\ -A^{-T}Q & A^{-T} \end{pmatrix}. \quad (54)$$

It is symplectic; one proof factors it as

$$S_d = \begin{pmatrix} I & -G \\ 0 & I \end{pmatrix} \begin{pmatrix} A & 0 \\ 0 & A^{-T} \end{pmatrix} \begin{pmatrix} I & 0 \\ -Q & I \end{pmatrix}. \quad (55)$$

Each factor is one of the blocks on page 2. On an invariant graph $\lambda_k = Px_k$,

$$x_{k+1} = (I + GP)^{-1}Ax_k = (A - BK)x_k, \quad (56)$$

where $K = (R + B^T PB)^{-1}B^T PA$. Substitution in the costate equation yields

$$P = Q + A^T P (I + GP)^{-1} A \quad (57)$$

$$P = Q + A^T PA - A^T PB(R + B^T PB)^{-1}B^T PA. \quad (58)$$

This is the DARE. The stabilizing graph selects the eigenvalues of S_d **inside the unit circle**; their reciprocals lie outside. Continuous-time stability uses the left half-plane, not the unit disk.

9.2. Singular dynamics: do not force an inverse that does not exist

Many legitimate systems have singular A , including sampled models with zero modes. Keep the optimality relations as a matrix pair:

$$M \begin{pmatrix} x_k \\ \lambda_k \end{pmatrix} = L \begin{pmatrix} x_{k+1} \\ \lambda_{k+1} \end{pmatrix}, \quad (59)$$

$$M = \begin{pmatrix} A & 0 \\ -Q & I \end{pmatrix}, \quad L = \begin{pmatrix} I & G \\ 0 & A^T \end{pmatrix}. \quad (60)$$

Direct multiplication gives the symplectic-pair identity

$$MJM^T = LJL^T = \begin{pmatrix} 0 & A \\ -A^T & 0 \end{pmatrix}. \quad (61)$$

The generalized eigenproblem $Mv = \mu Lv$ describes its modes. If L is invertible it reduces to $S_d = L^{-1}M$; otherwise use the pencil $M - \mu L$ and a generalized Schur/QZ method. Under the standard discrete stabilizability and detectability assumptions, its stable deflating subspace represents the stabilizing solution. Production solvers may use an **extended** pencil rather than this reduced pair [8].

This connects directly to the site's SDA note and QZ note: doubling composes optimal-control blocks while retaining the pair structure; QZ isolates a stable deflating subspace without forming $L^{-1}M$.

10. What to check in practice

Coordinates first. Identify the ordering and the actual canonical variables. Interleaving $(q_1, p_1, q_2, p_2, \dots)$ permutes J ; positions and velocities may require a different form. A failed test with the wrong J says little.

Test the right object. For a transition test $S^T JS - J$; for a generator test $F^T J + JF$; for a pair test $MJM^T - LLL^T$. Evaluate scaled residuals, not exact floating-point equality. For example,

$$r_S = \frac{\|S^T JS - J\|_F}{\|S\|_2^2 \|J\|_F + \|J\|_F}. \quad (62)$$

A small residual is a useful diagnostic, not a certificate of a small distance to the symplectic group when matrices are badly scaled or ill-conditioned.

For Riccati solutions, also test the branch. Check symmetry, the CARE/DARE residual, and the poles of $A - BK$. A zero equation residual does not distinguish the stabilizing and destabilizing solutions. Poor separation near the imaginary axis or unit circle, and an ill-conditioned upper subspace block, make the problem numerically sensitive [7], [8].

10.1. Reproduce the figures and numerical checks

The companion script uses the existing Octave control package. From the repository root:

```
QT_QPA_PLATFORM=offscreen octave --no-gui --quiet \
  Mathematics/symplectic-matrices/scripts/verify_and_plot.m
typst compile Mathematics/symplectic-matrices/symplectic-matrices.typ
```

It regenerates all three figures and checks: the four-dimensional counterexample; the oscillator's exact modified invariant and both sides of the step-size stability boundary; the exact double-integrator LQR solution; the finite-horizon hyperbolic formula; an asymmetric discrete-time graph; a singular-dynamics pair; and the optical transfer identity. No simulation-derived number is used as an independent expected answer.

10.2. Three questions to test understanding

1. **Does a symplectic squeeze preserve oscillator energy?** No: $(1, 0)$ becomes $(2, 0)$, so $\frac{q^2+p^2}{2}$ increases from $\frac{1}{2}$ to 2 although canonical area stays fixed.
2. **Can every pole of a symplectic matrix be strictly inside the unit circle?** No: reciprocal pairing forbids it. A closed-loop controller contracts on a stable graph, not throughout the entire state–costate space.
3. **Why is the LQR value Hessian symmetric?** The value is a scalar Hessian; equivalently, its graph is Lagrangian. The latter explanation remains visible even when computing the solution from an invariant subspace.

The unifying picture: canonical pairing $\rightarrow$ Hamiltonian generator $\rightarrow$ symplectic flow $\rightarrow$ preserved Lagrangian subspaces $\rightarrow$ symmetric Riccati graphs. Physics uses the full flow; optimal control selects the stable graph inside it.

Sources and further reading

- [1] E. Hairer, “Geometric Numerical Integration: Lecture 2, Symplectic Integrators.” [Online]. Available: https://www.unige.ch/~hairer/poly_geoint/week2.pdf
- [2] E. Hairer, C. Lubich, and G. Wanner, *Geometric Numerical Integration: Structure-Preserving Algorithms for Ordinary Differential Equations*, 2nd ed. Springer, 2006. doi: 10.1007/3-540-30666-8.
- [3] R. Paschotta, “ABCD Matrix.” [Online]. Available: https://www.rp-photonics.com/abcd_matrix.html
- [4] C. Weedbrook *et al.*, “Gaussian Quantum Information.” [Online]. Available: <https://arxiv.org/abs/1110.3234>
- [5] R. Tedrake, “Underactuated Robotics: Linear Quadratic Regulators.” [Online]. Available: <https://underactuated.mit.edu/lqr.html>
- [6] A. J. Laub, “A Schur Method for Solving Algebraic Riccati Equations,” *IEEE Transactions on Automatic Control*, vol. 24, no. 6, pp. 913–921, 1979, doi: 10.1109/TAC.1979.1102178.
- [7] SciPy developers, “solve_continuous_are: Algorithm Notes.” [Online]. Available: https://docs.scipy.org/doc/scipy/reference/generated/scipy.linalg.solve_continuous_are.html
- [8] SciPy developers, “solve_discrete_are: Algorithm Notes.” [Online]. Available: https://docs.scipy.org/doc/scipy/reference/generated/scipy.linalg.solve_discrete_are.html