PROCEEDINGS A

royalsocietypublishing.org/journal/rspa

HJ-Gauss

Check for updates

Article submitted to journal

Subject Areas:
dynamical systems, game theory, control systems, runtime safety assurance

Keywords:
reachability analysis, differential games, dynamical systems, verification, rigorous approximation

Correspondence:
lekanmolu@scriptedonachip.com

Sampling-based HJ Reachability

Lekan Molu

Seattle, WA 98109

Backward reachable sets and tubes for differential games are characterized by Hamilton–Jacobi equations whose numerical solution becomes costly as state dimension increases. We study a generalized Cole–Hopf transformation of the viscous HJ Cauchy equation associated with reachability while retaining its residual. The transformed problem is a semilinear heat equation whose mild representation couples Gaussian initial-condition and nonlinear source integrals. We derive its spatial gradient, establish conditional finite-time convergence of source iteration on a positive invariant set, and obtain a residual certificate by comparison. Multidimensional verification combines exact solutions, direct integration of the full transformed equation, and independent discretizations of the same viscous equation for control and pursuit–evasion problems. A starling-inspired 4D murmurations model supports evaluation at 100,000 states, with computations that assess discretization and population-query accuracy.

Animated comparisons

Reference ⟶ HJ-Gauss ⟶ Difference

Rockets | **Relative Dubins**

Double integrator | **Murmurations**

Code and reproducibility workflow

© The Authors. Published by the Royal Society under the terms of the Creative Commons Attribution License http://creativecommons.org/licenses/by/4.0/, which permits unrestricted use, provided the original author and source are credited.

THE ROYAL SOCIETY PUBLISHING

1. Introduction

Modern complex systems — with their diverse interconnection of actuators, sensors, control systems, and the software modules that coordinate their various processes and events — require software verification tools capable of geometrically reasoning about their future behaviors. Given the continuous and/or hybrid systems characterization of their dynamics, simulation and testing, verification, observation, and further testing loop cycles are the standard norm for detecting unforeseen behavior ahead of deployment. Owing to their scale, however, analyzing system behavior based on closed-form analytic mathematical tools is often infeasible. This has stimulated efforts that complement traditional mathematical analysis tools from dynamical systems with emerging techniques in reachability [1], viability [2], or invariance [16] theories. These new tools allow exploring all possible scenarios where faults may arise for safety assurance; that is, ascertaining the freedom of a system from harm by generating evidence that a system, or any its components satisfy all specified requirements including functional and allocated baselines (i.e., *verification*) [11]. Examples of complex systems include smart manufacturing, ground [23] and air traffic management systems [35], flight control systems [26], or distributed systems.

Differential optimal control and game theories are useful formal verification tools for certifying the runtime safety assurance [16] of multiple agents interacting over a shared space in a *geometric* sense [18]. Both rely on a resolution of the Hamilton-Jacobi-Bellman (HJB) and/or the Hamilton-Jacobi-Isaacs (HJI) equation [10]. Being partial differential equations (PDEs), general high-dimensional HJB/HJI equations rarely admit closed-form solutions. Thus, approximation techniques have become the default mechanism for formal systems verification in these settings. Reachability theory [27, 35] is one such approximation technique. It entails computing the essentially non-oscillating (ENO) [33] Lax-Friedrichs [9] approximations to the HJ PDE via levelset methods [28, 30, 31]. Propagating the fronts of a conservation law on a mesh, the inviscid Euler equations' solution reduces to the derivative of the HJ problem's solution [13]; such that integrating this solution with total variation diminishing Runge-Kutta (TVD-RK) [32] produces generalized *viscosity solutions* to the original HJ problem.

HJ equation systems are employed within reachability analysis to evaluate a dynamical system's behavioral (or trajectory) evolution throughout a state space — for example, given its initial point sets (states). The problem is typically setup so that the initial set must satisfy *input or state constraints* in a *least restrictive sense* and under a *worst-possible disturbance* [24]. Reachability is thus a useful geometric reasoning tool for determining the states that can be reached from an initial condition in a finite number of steps [27, 30] . *Reachability* can be analyzed in a (i) *forward* sense, where system trajectories are examined to determine if they enter certain states from an *initial set*; (ii) *backward* sense, where system trajectories are examined to determine if they enter certain *target sets*; (iii) *reach set* sense, in which they are examined to see if states reach a *set* at a *particular time*; or (iv) *reach tube* sense, in which they are examined to determine that they reach a set *during a time interval*.

This article studies the Cauchy equations associated with terminal reachability and backward reachable tubes. A terminal reachable set concerns target entry at a fixed horizon; a tube permits entry during an interval. Reach–avoid problems additionally constrain the trajectory before target entry and require an obstacle or stopping formulation. Classical tensor-grid discretizations require storage that grows exponentially with the number of state coordinates at fixed resolution [4]; avoiding that storage does not alone establish numerical accuracy or favorable complexity.

Consider the evolution-form of the Cauchy-type HJ PDE

$$\boldsymbol{v}_t + \boldsymbol{H}(t; \boldsymbol{x}, D\boldsymbol{v}) = 0 \text{ in } \Omega \times (0, T], \quad \boldsymbol{v}(0; \boldsymbol{x}) = \boldsymbol{g}(0; \boldsymbol{x}) \text{ on } \Omega \times \{t = 0\} \tag{HJ-IVP}$$

where $\boldsymbol{x} \in \Omega \subseteq \mathbb{R}^n$, the Hamiltonian $\boldsymbol{H}$ is continuous, and $D\boldsymbol{v}$ is the spatial gradient. The initial condition $\boldsymbol{g}$ is bounded and uniformly continuous. Crossing characteristics can destroy classical differentiability even when the Hamiltonian and initial condition are smooth [7]. Viscosity solutions provide generalized solutions. We study a fixed, positive-viscosity regularization before considering its approximation error.

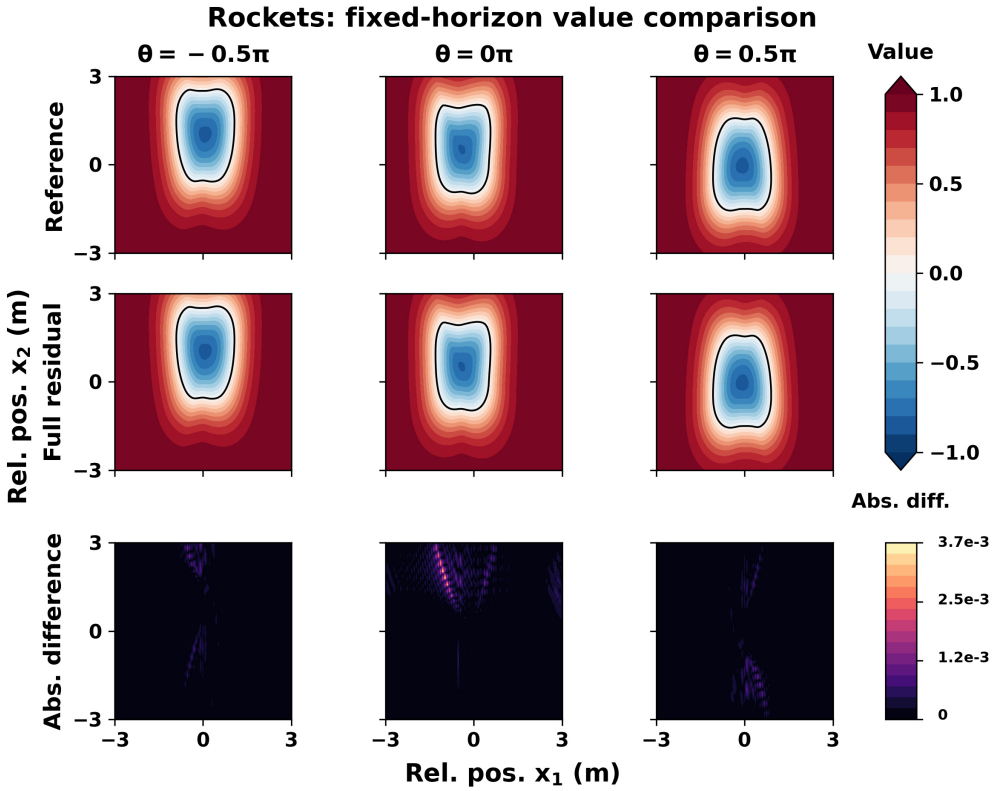


Figure 1: Normalized rockets benchmark at elapsed horizon $T = 0.5$, with $\delta = 0.08$ with fixed-horizon values. Rows show the independent finite-difference reference [28, 30], direct integration of the quasilinearized variational HJ equation, and their absolute difference, respectively. Columns show the slices of the BRTs at $\theta = -\pi/2$, 0, and $\pi/2$. Both solvers use the same initial condition and viscosity parameter. Black contours mark the zero levelset. The numerical validation protocol is given in Section 3(c).

In this paper, we apply a generalized Cole–Hopf transformation on the viscous form of (HJ-IVP) to yield a semilinear heat equation under positive logarithmic inversion. This retains the Hamiltonian and value derivatives. It is further decomposed into Gaussian kernel linear subproblems which are dependent on an unknown space–time value and gradient. Our construction uses the semilinear parabolic framework of [15, 17] — which does not imply a new dimension-independent complexity result. In essence, our results concern (i) the transformation identity and its differentiated mild representation; (ii) a finite-time source iteration that contracts on a positive invariant set under explicit value-and-gradient Lipschitz assumptions; and (iii) a comparison that converts a uniformly bounded residual into a value-error budget and conditional sign certificate.

The rest of this paper is structured as follows: §2 discusses background material and presents our theoretical machinery, §3 provides numerical validation on moderate to high-dimensional problems, and we conclude the paper in §4. Comprehensive proofs to our results are presented in Appendix A, convergence, residual certification and robustness analyses are given in B, while further results are presented in Appendix C.

2. Background and Quasilinearization Framework

Backward reachability associates a differential game with a Hamilton–Jacobi–Isaacs equation [27]. Let $\mathcal{U} \subset \mathbb{R}^m$ and $\mathcal{W} \subset \mathbb{R}^p$ be compact control sets. At physical time $s \in [0,T]$, let $\overline{\mathcal{U}}(s)$ and $\overline{\mathcal{W}}(s)$ contain the measurable control signals on $[s,T]$. Let $\mathcal{B}(s)$ contain the nonanticipative maps $\beta : \overline{\mathcal{U}}(s) \to \overline{\mathcal{W}}(s)$. The trajectory ξ starts in state $\boldsymbol{x}$ and satisfies

$$\dot{\xi}(r) = f(r, \xi(r), \boldsymbol{u}(r), \beta[\boldsymbol{u}](r)), \qquad \xi(s) = \boldsymbol{x}, \quad s \le r \le T, \tag{2.1}$$

where f is bounded, continuous and uniformly Lipschitz in $\boldsymbol{x}$. For bounded uniformly continuous running cost l and terminal cost $\boldsymbol{g}$, the lower value of the differential game is

$$\boldsymbol{v}_{\mathrm{phys}}(s,\boldsymbol{x}) = \inf_{\beta \in \mathcal{B}(s)} \sup_{\boldsymbol{u} \in \overline{\mathcal{U}}(s)} \left\{ \int_s^T l(r, \xi(r), \boldsymbol{u}(r), \beta[\boldsymbol{u}](r))\, dr + \boldsymbol{g}(\xi(T)) \right\}. \tag{2.2}$$

The target is represented by the zero sublevel set of $\boldsymbol{g}$. We assign $\boldsymbol{u}$ to the evader, which maximizes the payoff to avoid this set, and $\boldsymbol{w}$ to the pursuer, which minimizes the payoff to force target entry. The Hamiltonian is

$$\boldsymbol{H}_{\mathrm{phys}}(s,\boldsymbol{x},\boldsymbol{p}) = \max_{\boldsymbol{u} \in \mathcal{U}} \min_{\boldsymbol{w} \in \mathcal{W}} \{\boldsymbol{p} \cdot f(s,\boldsymbol{x},\boldsymbol{u},\boldsymbol{w}) + l(s,\boldsymbol{x},\boldsymbol{u},\boldsymbol{w})\}. \tag{2.3}$$

For pure target-reach problems, the running cost is set to zero and the target is represented by $\boldsymbol{g}$. The elapsed time $t = T - s$ converts the terminal condition into an initial condition so that if we write $\boldsymbol{v}(t,\boldsymbol{x}) = \boldsymbol{v}_{\mathrm{phys}}(T-t,\boldsymbol{x})$ and define $\boldsymbol{H}(t,\boldsymbol{x},\boldsymbol{p}) = -\boldsymbol{H}_{\mathrm{phys}}(T-t,\boldsymbol{x},\boldsymbol{p})$, the elapsed equation is $\boldsymbol{v}_t + \boldsymbol{H}(t,\boldsymbol{x},D\boldsymbol{v}) = 0$. Fixing $\delta > 0$ gives the viscous parabolic regularization

$$\boldsymbol{v}^\delta_t + \boldsymbol{H}(t,\boldsymbol{x},D\boldsymbol{v}^\delta) = \frac{\delta}{2}\Delta \boldsymbol{v}^\delta, \qquad \boldsymbol{v}^\delta(0,\boldsymbol{x}) = \boldsymbol{g}(\boldsymbol{x}). \tag{2.4}$$

Equation (2.4) supplies comparison and approximation results under additional hypotheses on the Hamiltonian and initial condition [7, 8]. The transformation below concerns classical positive-time solutions of this regularized equation.

(a) HJ Linearization

Write $\boldsymbol{v} = \boldsymbol{v}^\delta$ within this subsection and choose a prescribed coefficient $\boldsymbol{c} \in C_b^{1,2}([0,T] \times \mathbb{R}^n)$ with $0 < \boldsymbol{c}_- \le \boldsymbol{c} \le \boldsymbol{c}_+$. Define the transformed value and its initial condition by

$$\boldsymbol{\omega} = e^{-\boldsymbol{c}\boldsymbol{v}}, \qquad \boldsymbol{\omega}_0(\boldsymbol{x}) = e^{-\boldsymbol{c}(0,\boldsymbol{x})\boldsymbol{g}(\boldsymbol{x})}. \tag{2.5}$$

A constant positive coefficient avoids coefficient derivatives and the singular ratio $\boldsymbol{H}/|D\boldsymbol{v}|^2$, while a variable coefficient permits local scaling. For a scalar $z > 0$ and vector $q \in \mathbb{R}^n$, introduce

$$V(t,\boldsymbol{x},z) = -\frac{\log z}{\boldsymbol{c}(t,\boldsymbol{x})}, \qquad P(t,\boldsymbol{x},z,q) = -\frac{q}{\boldsymbol{c}z} + \frac{\log z}{\boldsymbol{c}^2} D\boldsymbol{c}, \tag{2.6}$$

$$\boldsymbol{R}_{\mathrm{alg}}[\boldsymbol{c},V,P] = \boldsymbol{c}\left(\boldsymbol{H}(t,\boldsymbol{x},P) - \frac{\delta \boldsymbol{c}}{2}|P|^2\right), \tag{2.7}$$

$$\boldsymbol{R}_{\mathrm{der}}[\boldsymbol{c},V,P] = -\boldsymbol{c}_t V + \frac{\delta}{2}(\Delta \boldsymbol{c})V - \frac{\delta}{2}|D\boldsymbol{c}|^2 V^2 + \delta(1 - \boldsymbol{c}V) D\boldsymbol{c} \cdot P, \tag{2.8}$$

$$F(t,\boldsymbol{x},z,q) = z\{\boldsymbol{R}_{\mathrm{alg}}[\boldsymbol{c},V,P] + \boldsymbol{R}_{\mathrm{der}}[\boldsymbol{c},V,P]\}. \tag{2.9}$$

Then, the transformed equation is

$$\boldsymbol{\omega}_t - \frac{\delta}{2}\Delta \boldsymbol{\omega} = F(t,\boldsymbol{x},\boldsymbol{\omega},D\boldsymbol{\omega}), \qquad \boldsymbol{\omega}(0) = \boldsymbol{\omega}_0. \tag{2.10}$$

For positive classical solutions, logarithmic inversion recovers the original viscous equation. It reduces to a heat equation when the sum in (2.9) vanishes along the solution.

Algorithm 1 Source iteration on a finite interval

1: Choose $\delta > 0$, a prescribed smooth coefficient $\boldsymbol{c} \geq \boldsymbol{c}_- > 0$, and an interval $[0, \tau]$.
2: Set $\boldsymbol{\omega}_0 = \exp(-\boldsymbol{c}(0,\cdot)\boldsymbol{g})$ and $z_0(t) = S_t\boldsymbol{\omega}_0$.
3: Specify a positive invariant admissible set containing z_0 and a contraction bound $\theta < 1$ for the source map.
4: **for** $k = 0, 1, \ldots$ **do**
5: Evaluate $f_k = F(\cdot, z_k, Dz_k)$ on the required $\Omega \times (0, T]$.
6: Compute the value and gradient updates (2.12).
7: Bound sampling, quadrature, and reconstruction errors; verify admissibility and positivity.
8: Recover the value and gradient using (2.13).
9: Stop when the controlled error budget meets tolerance; otherwise continue or shorten the interval.
10: **end for**

Let S_t denote the whole-space heat semigroup with covariance $\delta t I_n$. The transformed equation admits the *mild representation*

$$\boldsymbol{\omega}(t) = S_t\boldsymbol{\omega}_0 + \int_0^t S_{t-s}F(s, \cdot, \boldsymbol{\omega}(s), D\boldsymbol{\omega}(s))\, ds, \tag{2.11}$$

coupling intermediate-time values and gradients. Each mild update substitutes the preceding iterate into both source arguments; an initial-condition Gaussian average alone omits the source integral. Finite-time convergence requires the source and gradient bounds in Theorem 2.2.

For $f_k(s, \boldsymbol{x}) = F(s, \boldsymbol{x}, z_k, Dz_k)$, the value and gradient update as an ordered pair:

$$(z_{k+1}, Dz_{k+1})(t) = (S_t, DS_t)\boldsymbol{\omega}_0 + \int_0^t (S_{t-s}, DS_{t-s})f_k(s)\, ds. \tag{2.12}$$

$$\boldsymbol{v}_{k+1}(t, \boldsymbol{x}) = -\frac{\log z_{k+1}(t, \boldsymbol{x})}{\boldsymbol{c}(t, \boldsymbol{x})}, \qquad D\boldsymbol{v}_{k+1} = -\frac{Dz_{k+1}}{\boldsymbol{c}z_{k+1}} - \frac{\boldsymbol{v}_{k+1}}{\boldsymbol{c}} D\boldsymbol{c}. \tag{2.13}$$

Appendix A derives the Gaussian forms of the source integral and coefficients. Bounded transformed initial conditions suffice for a bounded-source mild representation. Uniform logarithmic stability also requires a positive lower bound; an unbounded signed-distance initial condition does not supply that bound automatically. The coefficient is prescribed throughout this iteration. The one-dimensional verification uses exponential-Euler time stepping, rather than the whole-interval iteration in Algorithm 1.

(b) Relations to Backward Reachable Sets/Tubes

The target and backward reachable tube satisfy

$$\mathcal{L}_0 = \{\boldsymbol{x} : \boldsymbol{g}(\boldsymbol{x}) \leq 0\}, \quad \mathcal{L}([s, T], \mathcal{L}_0) = \{\boldsymbol{x} : \boldsymbol{v}_{\mathrm{tube,phys}}(s, \boldsymbol{x}) \leq 0\}. \tag{2.14}$$

For zero running cost, the tube value replaces the terminal payoff by $\min_{r\in[s,T]} \boldsymbol{g}(\xi(r))$. Under the game's reachability characterization, its zero sublevel set represents target entry during the interval [26, 27]. In particular, $\boldsymbol{x} \in \mathcal{L}([s, T], \mathcal{L}_0) \iff \boldsymbol{v}_{\mathrm{tube,phys}}(s, \boldsymbol{x}) \leq 0$. When the minimizing strategy and worst-case payoff are attained, the tube has the strategy characterization

$$\mathcal{L}([s, T], \mathcal{L}_0) = \{\boldsymbol{x} : \exists \beta \in \mathcal{B}(s)\, \forall \boldsymbol{u} \in \overline{\mathcal{U}}(s)\, \exists r \in [s, T],\ \xi_{s,\boldsymbol{x}}^{\boldsymbol{u},\beta[\boldsymbol{u}]}(r) \in \mathcal{L}_0\}. \tag{2.15}$$

A reachable set at a fixed terminal horizon uses the terminal payoff instead. For autonomous dynamics, a time-independent target, and zero running cost, the standard frozen-dynamics

characterization gives the physical-time tube equation [27]

$$\partial_s \boldsymbol{v}_{\mathrm{tube,phys}} + \min\{0, \boldsymbol{H}_{\mathrm{phys}}(\boldsymbol{x}, D\boldsymbol{v}_{\mathrm{tube,phys}})\} = 0, \qquad \boldsymbol{v}_{\mathrm{tube,phys}}(T) = \boldsymbol{g}. \tag{2.16}$$

A positive-viscosity Cauchy regularization of this autonomous equation is

$$\partial_t \boldsymbol{v}^{\delta}_{\mathrm{tube}} + \max\{0, \boldsymbol{H}(\boldsymbol{x}, D\boldsymbol{v}^{\delta}_{\mathrm{tube}})\} = \frac{\delta}{2}\Delta \boldsymbol{v}^{\delta}_{\mathrm{tube}}, \qquad \boldsymbol{v}^{\delta}_{\mathrm{tube}}(0) = \boldsymbol{g}. \tag{2.17}$$

The regularized Cauchy problem uses $\max\{0, \boldsymbol{H}\}$. Applying the transformation to the unmodified terminal-game Hamiltonian solves a different problem. The frozen-dynamics representation only applies to autonomous dynamics. For time-dependent dynamics, allowing a trajectory to pause can alter target reachability. Moreover, the positive-viscosity regularization retains diffusion when the Hamiltonian contribution vanishes. It therefore regularizes the autonomous tube equation rather than modeling a diffusion stopped at target entry. Reach–avoid problems require trajectories to satisfy a state constraint until target entry. If $h(\boldsymbol{x}) \le 0$ describes that constraint, then both h and the target function enter the stopping payoff and its terminal condition. Their value satisfies an obstacle problem. The Cauchy-source formulas apply in its continuation region but do not enforce obstacle complementarity or the stopping rule. No reach–avoid guarantee follows from the Cauchy analysis alone.

(c) Weighted Importance Sampling of Reachable Sets

Importance sampling changes the Gaussian proposal while preserving the integral through a density ratio. For a displacement $\vartheta \in \mathbb{R}^n$, variance $\sigma^2 > 0$, and sample $\boldsymbol{y} \sim \mathcal{N}(\boldsymbol{x} + \vartheta, \sigma^2 I_n)$, the "importance weight" is

$$w_\vartheta(\boldsymbol{y}) = \exp\left(\frac{|\vartheta|^2 - 2\vartheta \cdot (\boldsymbol{y} - \boldsymbol{x})}{2\sigma^2}\right). \tag{2.18}$$

Use $\sigma^2 = \delta t$ for the initial-condition integral and $\sigma^2 = \delta(t - s)$ for the source integral. Each changed proposal requires its own importance weights and estimator moment bounds. Effective sample size describes sampled-weight balance, not an estimation-error bound. The evaluation count M counts queried states and sample count N counts independent random draws at each query and update. Source evaluation requires intermediate-time values and gradients at sampled states. Their reconstruction and storage must enter the cost analysis, so a single endpoint sample array does not establish the cost of the full nonlinear solve.

(d) Sampling Complexity and Convergence Guarantees

Conditional on the preceding iterate, suppose $a_0 \le \boldsymbol{\omega}_0 \le b_0$ and $|f_k| \le B_F$. Uniform source-time sampling gives the unbiased transformed-value sample

$$A = \boldsymbol{\omega}_0(\boldsymbol{x} + \sqrt{\delta t} Z) + t f_k(tU, \boldsymbol{x} + \sqrt{\delta t(1-U)} Z), \qquad U \sim \mathcal{U}(0,1), \quad Z \sim \mathcal{N}(0, I_n), \tag{2.19}$$

with independent U and Z. Each sample lies in $[a_0 - tB_F, b_0 + tB_F]$, of width $D = b_0 - a_0 + 2tB_F$, and has mean equal to the prescribed transformed update. Logarithmic value recovery and gradient ratios can introduce bias.

Theorem 2.1 (Finite-query concentration for a fixed source)**.** *Fix positive integers M, K, N, representing the number of queries, the maximum number of updates, and the samples in each batch, and let $\eta \in (0,1)$. Assume the preceding source bounds hold conditionally at every update with common sample-width bound D_*. If each batch contains N conditionally independent samples with the stated common conditional law, then all transformed-value sampling errors are at most*

$$h_N = D_* \sqrt{\frac{\log(2MK/\eta)}{2N}} \tag{2.20}$$

with probability at least $1 - \eta$.

The proof is given in Appendix B.

Corollary 2.1 (Sample count and logarithmic recovery)**.** *For transformed-value tolerance $\varepsilon_\omega > 0$, it suffices to choose*

$$N \geq \frac{D_*^2}{2\varepsilon_\omega^2} \log \frac{2MK}{\eta}. \tag{2.21}$$

If the target mild update is at least $a > 0$ and $h_N < a$, then its sampled logarithmic inverse differs by at most $h_N/[\mathbf{c}_-(a - h_N)]$ on the same event.

The proof is given in Appendix B.

These statements do not give a continuum or gradient error bound. The endpoint-weighted source-time proposal used by the prototype has interval width $b_0 - a_0 + 4tB_F$ for its value sample; Appendix A derives the corresponding finite-second-moment gradient sample. Unbounded adaptive iteration counts require a summable allocation of failure probabilities. Write $\mathrm{BUC}^1(\mathbb{R}^n)$ for continuously differentiable functions with bounded, uniformly continuous values and gradients.

Theorem 2.2 (Conditional finite-time convergence of the full source map)**.** *Let $X_\tau = C([0,\tau]; \mathrm{BUC}^1)$ with the sum of the uniform value and gradient norms. Suppose a nonempty closed set $\mathcal{A} \subset X_\tau$ is invariant under the mild map, its functions satisfy $a \leq z \leq b$ and $|Dz| \leq Q$ with $a > 0$, and the bounded source is uniformly Lipschitz in (z, Dz) with constant L. If*

$$\theta = L\left(\tau + 2\sqrt{n\tau/\delta}\right) < 1, \tag{2.22}$$

then the full source map has a unique fixed point in $\mathcal{A}$ and its iterates starting in $\mathcal{A}$ converge geometrically in X_τ. If that fixed point is classical, its inverse solves the viscous Hamilton–Jacobi equation.

The proof is given with Theorem B.1 in Appendix B, together with an invariant-ball criterion and a posteriori stopping bound. Invariance, regularity, and controlled gradient reconstruction are necessary hypotheses.

(e) Notes on Monte Carlo and Viscosity Approximation Errors

For a bounded smooth approximation U with bounded residual

$$r_U = U_t + \boldsymbol{H}(t, \boldsymbol{x}, DU) - \frac{\delta}{2}\Delta U, \tag{2.23}$$

comparison gives a value-error bound by the initial-condition error plus $\int_0^t \|r_U(s)\|_\infty ds$. If an independent vanishing-viscosity theorem also gives a bound $C_{\mathrm{vis}}\sqrt{\delta}$, the combined error bound is

$$E(t) = \|U(0) - \boldsymbol{g}\|_\infty + \int_0^t \|r_U(s)\|_\infty ds + C_{\mathrm{vis}}\sqrt{\delta}. \tag{2.24}$$

The transformed residual satisfies $d = -\boldsymbol{c}\exp(-\boldsymbol{c}U)r_U$ (see Appendix B).

Corollary 2.2 (Conditional sign certification)**.** *Suppose $|U - \boldsymbol{v}^0| \leq E$ on a specified set, with $\boldsymbol{v}^0 \leq 0$ denoting the target-reaching class. Then $U > E$ implies $\boldsymbol{v}^0 > 0$, and $U < -E$ implies $\boldsymbol{v}^0 < 0$. The remaining states form the abstention band $|U| \leq E$.*

The proof is given in Appendix B.

Reducing δ can reduce viscosity bias while increasing source sensitivity, score variance, and logarithmic instability. The full-source Lipschitz and positivity constants also affect the admissible time interval. The global complexity theory for multilevel Picard methods has additional source and derivative-growth assumptions [15, 17]; Appendix B explains why they do not automatically hold here.

Table 1: Maximum sampled discrepancies for the full residual computations. Refinement and independent-reference columns use the common observation queries; the residual column gives the largest sampled original-equation residual over native states at three interior times. All payoffs are dimensionless. The grid column lists the fine spatial and heading counts. Pending or failed checks are disclosed in the reproducibility record.

System	Grid	Refinement	Reference	Residual	Status
Rockets	$160^2 \times 80$	0.007481	0.003223	0.1433	Pass
Relative Dubins	$224^2 \times 112$	0.005717	0.002511	0.1379	Pass
Single Dubins	$128^2 \times 64$	0.009297	0.001415	0.02992	Pass
Double integrator	128^2	0.006741	5.60×10^{-4}	0.01187	Pass
Bird model	$192^2 \times 96$	0.008817	6.44×10^{-4}	0.1295	Pass

3. Numerical results

The experiments test full-residual discretization accuracy and population-scale field evaluation. Paired solvers share the viscosity, Hamiltonian, bounded initial condition, and horizon, separating numerical discrepancy from viscosity bias. The one-dimensional analytic verification is reported in Appendix C(e).

(a) Games of two vehicles on a plane

The pursuit–evasion examples use relative Dubins dynamics [25, 26] and a normalized rockets benchmark motivated by [12, 29]. Appendix C specifies the physical control optimization and its elapsed-time sign. Figure 1 compares the rockets value fields at three headings. The relative-Dubins example in Figure 3 provides a second game validation. The rows in Table 1 denote terminal-horizon Cauchy problems, where a zero contour represents the numerical terminal reachable set for the viscosity solution. The normalized rockets system is defined directly by its controlled vector field; the physical-unit illustrations and launch problem are given in [29].

(b) Population-scale evaluation of starling-inspired dynamics

The attacked-agent model has two relative positions, altitude, and heading (Appendix C(d)). Equal vertical control bounds cancel the vertical Hamiltonian terms; an altitude-independent payoff then reduces the solution to the three-dimensional relative-Dubins equation. Independent enumeration of the four-dimensional controls checks this reduction.

We evaluate the resulting field at 100,000 freshly generated four-component states, using a fixed saved random seed. Positions are independently uniform on $[-3, 3]^2$, altitude on $[-1, 1]$, and heading on $[-\pi, \pi)$. Every queried value is retained, together with its independent-reference value. The population contains 100,000 finite paired evaluations. Of these, 99,788 have a positive approximate value; 2 change sign between the two numerical fields. These counts describe the saved query population. They are not a simulation of coupled flock trajectories or a population-wide safety certificate. The maximum paired value discrepancy is 6.52×10^{-4}. The 2 sign disagreements occur near the zero contour: their reference magnitudes are at most 5.03×10^{-5}. Thus the two discretizations agree on most classifications in this population, while the near-interface disagreements identify states whose sign needs further error control. Appendix C(d) relates these comparisons to the displayed geometry.

(c) Refinement and sensitivity

For the rockets, relative-Dubins, single-Dubins, and double-integrator cases, the transformed solver integrates the complete source in (2.9) by second-order exponential Runge–Kutta time

stepping on a periodic Fourier grid. A fixed cutoff retains source modes with component index at most one third of the corresponding grid count. This is a collocation approximation to the full source; the equation residual is evaluated against the untruncated source. A separate fourth-order finite-difference discretization with classical fourth-order Runge–Kutta time stepping evolves the original viscous equation. The bird model instead uses fourth-order centered finite differences and classical fourth-order Runge–Kutta integration for the transformed equation, with no source-mode cutoff. Its independent original-variable reference uses sixth-order centered differences and fourth-order Runge–Kutta integration. Both transformed discretizations retain every analytic source term. Positivity is checked at every transformed stage, without clipping. The deterministic comparisons test these discretizations; they do not test the convergence of stochastic source iteration.

The equations are computed in normalized coordinates. Figures assign the reference scales $L_* = 1\,\mathrm{m}$ and $T_* = 1\,\mathrm{s}$, so positions are displayed in meters and elapsed time in seconds. Double-integrator velocity uses $L_*/T_* = 1\,\mathrm{m\,s^{-1}}$; headings are in radians and payoffs remain normalized. These choices specify display units without changing the normalized dynamics or numerical values. For a circular target of radius r, the initial condition in the observation region is

$$\boldsymbol{g}(\boldsymbol{x}) = \tanh\left(\sqrt{\boldsymbol{x}_1^2 + \boldsymbol{x}_2^2 + \epsilon_g^2} - \sqrt{r^2 + \epsilon_g^2}\right), \qquad \epsilon_g = 0.08. \tag{3.1}$$

The single-Dubins pose target adds $4\sin^2(\theta/2)$ under the first square root. The bird example takes the minimum of seven translated radial payoffs, with target centers on a radius-1.15 ring and disk radius 0.6. This specifies the actual viscous initial condition; preserving a signed-distance function's initial zero set does not preserve its positive-time viscous solution.

The baseline spatial box is $[-6, 6]^2$, with periodic heading $[-\pi, \pi)$ for three-dimensional cases. A smooth completion makes the initial condition equal to 1 near each spatial seam while leaving it unchanged on $[-4.5, 4.5]^2$. The observation box is $[-2, 2]^2$ for the single-Dubins, relative-Dubins, and double-integrator cases and $[-3, 3]^2$ for rockets and birds. The prescribed coefficient is

$$\boldsymbol{c}(t, \boldsymbol{x}) = b + 0.04\cos t \prod_{j=1}^{d} \cos(2\pi \boldsymbol{x}_j / L_j), \qquad b = 0.2, \tag{3.2}$$

where L_j is the corresponding periodic length. The computation includes its time derivative, gradient, and Laplacian. The sensitivity calculation changes b to 0.35 while keeping the Hamiltonian and initial condition fixed.

Successive field refinements can change both spatial resolution and time step; fixed-grid time halving is checked separately. Refinement and solver discrepancies are measured on the same 65×65 observation points and, in three dimensions, 32 equally spaced headings. These are maximum sampled discrepancies of interpolated fields. Equation residuals use native grid states and spectral spatial derivatives and centered differences at one quarter, one half, and three quarters of the elapsed horizon. They are finite diagnostics rather than uniform space–time bounds.

The acceptance thresholds were fixed before the refinement cycle: 0.01 for successive field refinements, 0.02 for the independent reference, 0.002 for time halving and coefficient changes, and 0.005 for each domain-extension test. Sampled original-equation residuals must have maximum at most 0.15 and root-mean-square at most 0.02. The box experiment extends the half-width to 7.5 at fixed spatial spacing and heading resolution; it also changes the time step and coefficient wavelength, so it measures their combined sensitivity. A separate boundary-sensitivity test evolves the original equation on the enlarged box while fixing the time step, spatial spacing, heading count, Hamiltonian, and initial condition on the common domain. This separates boundary sensitivity from changes in the transformation coefficient and time step. Diffusion has infinite propagation speed, so neither finite extension supplies a whole-space boundary-error certificate.

Table 2 separates time-step, coefficient, and domain sensitivity. All comparisons meet their thresholds; these sampled checks do not establish the source iteration's invariant-set hypotheses.

Table 2: Maximum sampled changes under time halving, coefficient variation, isolated boundary extension, and combined box sensitivity. Boundary extension fixes the reference time step, spacing, heading count, Hamiltonian, and initial condition on the common box. The combined test also changes the time step and coefficient wavelength. Payoffs are normalized.

System	Time	Coefficient	Boundary	Combined box
Rockets	7.19×10^{-6}	1.84×10^{-4}	1.68×10^{-13}	0.00136
Relative Dubins	3.97×10^{-5}	2.79×10^{-5}	8.88×10^{-16}	5.26×10^{-5}
Single Dubins	1.47×10^{-6}	1.02×10^{-5}	0	8.83×10^{-6}
Double integrator	6.85×10^{-7}	1.73×10^{-6}	0	4.40×10^{-5}
Bird model	6.23×10^{-7}	7.46×10^{-6}	1.49×10^{-6}	5.46×10^{-4}

(d) Computational scope

The experiments distinguish population size from joint state dimension. A model of 15 individually steered planar vehicles has the dynamics

$$\dot{x}_i = a_i \cos\theta_i, \quad \dot{y}_i = a_i \sin\theta_i, \quad \dot{\theta}_i = u_i, \qquad i = 1, \dots, 15, \tag{3.3}$$

and a joint state in $\mathbb{R}^{45}$. A tensor grid with 101 points per coordinate would contain 101^{45} states. The present full residual computations resolve two- and three-dimensional fields, including an exact reduction of the four-component bird model. Evaluating one such field at many states does not establish accuracy or computational scaling for a joint 45-dimensional equation. Source reconstruction for general high-dimensional problems remains a separate computational requirement.

4. Conclusion

Retaining the transformation source preserves the viscous Hamilton–Jacobi equation under positive classical logarithmic inversion. The resulting mild representation includes both the initial-condition convolution and a nonlinear source integral. Its gradient also includes the coefficient derivative. Homogeneous heat flow is a special case, rather than the transformed equation for every game Hamiltonian. Source iteration converges on a finite interval under the stated positivity, invariance, and contraction assumptions. The sampling bounds control value-update errors at selected queries. A uniform equation-residual bound, combined with initial-condition and viscosity-error bounds where required, controls the value error and certifies its sign outside the uncertainty band. Extending sampled bounds to the entire domain requires additional error control.

References

1 M. Althoff, G. Frehse, and A. Girard. Set Propagation Techniques for Reachability Analysis. *Annual Review of Control, Robotics, and Autonomous Systems*, 4:369–395, 2021. ISSN 25735144.

2 J.-P. Aubin, A. M. Bayen, and P. Saint-Pierre. *Viability Theory*. Number August 2015. 2011. ISBN 9783642166846.

3 M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. interaction Ruling Animal Collective Behavior Depends On Topological Rather Than Metric Distance: Evidence From A Field Study. *Proceedings of the National Academy of Sciences*, 105(4):1232–1237, 2008. ISSN 0027-8424. doi: 10.1073/pnas.0711437105. URL https://www.pnas.org/content/105/4/1232.

4 R. Bellman. *Dynamic programming*. Princeton University Press, 1957. ISBN 0-486-42809-5.

5 S. P. Bhat and D. S. Bernstein. Continuous finite-time stabilization of the translational and rotational double integrators. *IEEE Transactions on automatic control*, 43(5):678–682, 1998.

6 A. Cavagna, A. Cimarelli, I. Giardina, G. Parisi, R. Santagati, F. Stefanini, and M. Viale. Scale-ffee correlations in starling flocks. *Proceedings of the National Academy of Sciences*, 107(26):11865–11870, 2010.

7 M. G. Crandall and P.-L. Lions. Viscosity solutions of hamilton-jacobi equations. *Transactions of the American mathematical society*, 277(1):1–42, 1983.

8 M. G. Crandall and P.-L. Lions. Two Approximations of Solutions of Hamilton-Jacobi Equations. *Mathematics of Computation*, 43(167):1–19, 1984.

9 M. G. Crandall and A. Majda. Monotone Difference Approximations For Scalar Conservation Laws. *Mathematics of Computation*, 34(149):1–21, 1980.

10 M. G. Crandall, L. C. Evans, and P. L. Lions. Some Properties of Viscosity Solutions of Hamilton-Jacobi Equations. *Transactions of the American Mathematical Society*, 282(2):487, 1984. ISSN 00029947.

11 Defense Acquisition University. Verification, 2023. URL https://www.dau.edu/tools/se-brainbook/Pages/Technical%20Processes/verification.aspx. Accessed April 5, 2023.

12 S. E. Dreyfus. Control Problems With Linear Dynamics, Quadratic Criterion, and Linear Terminal Constraints. Technical report, Rand Corp, Santa Monica Calif, 1966.

13 L. Evans and P. E. Souganidis. Differential Games And Representation Formulas For Solutions Of Hamilton-Jacobi-Isaacs Equations. *Indiana Univ. Math. J*, 33(5):773–797, 1984. ISSN 0022-2518.

14 A. Hatcher. *Algebraic Topology*, chapter 2. Cambridge University Press, 2002. Euler characteristic: $\chi(X) = \sum_{n=0}^{\infty}(-1)^n \text{rank}(H_n(X))$.

15 D. Henry. *Geometric Theory of Semilinear Parabolic Equations*, volume 840 of *Lecture Notes in Mathematics*. Springer, 1981. doi: 10.1007/BFb0089647.

16 K. L. Hobbs, M. L. Mote, M. C. Abate, S. D. Coogan, and E. M. Feron. Runtime assurance for safety-critical systems. *IEEE Control Systems Magazine*, 43:28–65, 2023.

17 M. Hutzenthaler and T. Kruse. Multilevel Picard Approximations of High-Dimensional Semilinear Parabolic Differential Equations with Gradient-Dependent Nonlinearities. *SIAM Journal on Numerical Analysis*, 58(2):929–961, 2020. doi: 10.1137/17M1157015.

18 R. Isaacs. Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization., 1966.

19 R. Isaacs. *Differential Games: A Mathematical Theory with Applications to Warfare and Pursuit, Control and Optimization.* Kreiger, Huntigton, NY, 1999.

20 A. Jadbabaie, J. Lin, and A. S. Morse. Coordination of groups of mobile autonomous agents using nearest neighbor rules. *IEEE Transactions on automatic control*, 48(6):988–1001, 2003.

21 P. Karnakov, S. Litvinov, and P. Koumoutsakos. Computing Foaming Flows Across Scales: From Breaking Waves to Microfluidics. *arXiv preprint arXiv:2103.01513*, 2021.

22 Y. Kim, M.-C. Lai, and C. S. Peskin. Numerical Simulations of Two-Dimensional Foam by the Immersed Boundary Method. *Journal of Computational Physics*, 229(13):5194–5207, 2010.

23 C. Livadas and N. A. Lynch. Formal verification of safety-critical hybrid systems. In *International Workshop on Hybrid Systems: Computation and Control*, pages 253–272. Springer, 1998.

24 J. Lygeros. On reachability and minimum cost optimal control. *Automatica*, 40(6):917–927, 2004.

25 A. Merz. The game of two identical cars. *Journal of Optimization Theory and Applications*, 9(5): 324–343, 1972.

26 I. Mitchell. A Robust Controlled Backward Reach Tube with (Almost) Analytic Solution for Two Dubins Cars. *EPiC Series in Computing*, 74:242–258, 2020.

27 I. M. Mitchell, A. M. Bayen, and C. J. Tomlin. A Time-Dependent Hamilton-Jacobi Formulation of Reachable Sets for Continuous Dynamic Games. *IEEE Transactions on Automatic Control*, 50 (7):947–957, 2005. ISSN 00189286.

28 L. Molu. The Python LevelSet Toolbox (LevelSetPy). In *2024 IEEE 63rd Conference on Decision and Control (CDC)*, pages 8938–8945, 2024. doi: 10.1109/CDC56724.2024.10886640.

29 L. Molu. The python levelset toolbox (levelsetpy). In *IEEE 63rd Conference on Decision and Control (CDC)*, pages 8938–8945, 2024. doi: 10.1109/CDC56724.2024.10886640.

30 L. Molu. LevelSetPy: A GPU-Accelerated Package for Hyperbolic Hamilton-Jacobi Partial Differential Equations' Solubility. *The ACM Transactions on Mathematical Software*, 2025.

31 S. Osher and R. Fedkiw. Level Set Methods and Dynamic Implicit Surfaces. *Applied Mechanics Reviews*, 57(3):B15–B15, 2004. ISSN 0003-6900.
32 S. Osher and J. A. Sethian. Fronts Propagating with Curvature-Dependent Speed: Algorithms based on Hamilton-Jacobi Formulations. *Journal of Computational Physics*, 79(1):12–49, 1988. ISSN 10902716.
33 S. Osher and C.-W. Shu. High-Order Essentially Nonoscillatory Schemes for Hamilton-Jacobi Equations. *SIAM Journal of Numerical Analysis*, 28(4):907–922, 1991. ISSN 15333868.
34 R. I. Saye and J. A. Sethian. The Voronoi Implicit Interface Method for Computing Multiphase Physics. *Proceedings of the National Academy of Sciences of the United States of America*, 108(49): 19498–19503, 2011. ISSN 10916490.
35 C. J. Tomlin, J. Lygeros, and S. S. Sastry. A game theoretic approach to controller design for hybrid systems. *Proceedings of the IEEE*, 88(7):949–970, 2000.
36 W. M. Wonham. Linear multivariable control: A geometric approach. *Applications of Mathematics*, 10, 1985.
37 A. Zaitzeff, S. Esedoglu, and K. Garikipati. On the Voronoi Implicit Interface Method. *SIAM Journal on Scientific Computing*, 41(4):A2407–A2429, 2019.

Appendices

A. Transformation and sampling formulas

The full residual changes the representation formula. Throughout this appendix, the domain is $\mathbb{R}^n$, the diffusion coefficient is $\delta/2$, and $\boldsymbol{c}$ satisfies the bounds in §2(a).

Lemma A.1 (Transformation identity)**.** *Suppose $\boldsymbol{v}, \boldsymbol{c} \in C^{1,2}$ and set $\boldsymbol{\omega} = e^{-\boldsymbol{c}\boldsymbol{v}}$. Define the original-equation residual $r_v = \boldsymbol{v}_t + \boldsymbol{H}(t, \boldsymbol{x}, D\boldsymbol{v}) - \delta\Delta\boldsymbol{v}/2$. Then*

$$\boldsymbol{\omega}_t - \frac{\delta}{2}\Delta\boldsymbol{\omega} = \boldsymbol{\omega}\{\boldsymbol{R}_{\mathrm{alg}}[\boldsymbol{c}, \boldsymbol{v}, D\boldsymbol{v}] + \boldsymbol{R}_{\mathrm{der}}[\boldsymbol{c}, \boldsymbol{v}, D\boldsymbol{v}] - \boldsymbol{c}r_v\}. \tag{A.1}$$

Consequently, positive classical solutions of the full transformed equation and classical solutions of the original equation correspond under logarithmic inversion.

Proof. Set $\phi = \boldsymbol{c}\boldsymbol{v}$. Differentiation gives $\boldsymbol{\omega}_t = -\boldsymbol{\omega}\phi_t$, $D\boldsymbol{\omega} = -\boldsymbol{\omega}D\phi$, and $\Delta\boldsymbol{\omega} = \boldsymbol{\omega}(|D\phi|^2 - \Delta\phi)$. Expand $D\phi = \boldsymbol{c}D\boldsymbol{v} + \boldsymbol{v}D\boldsymbol{c}$ and $\Delta\phi = \boldsymbol{c}\Delta\boldsymbol{v} + 2D\boldsymbol{c}\cdot D\boldsymbol{v} + \boldsymbol{v}\Delta\boldsymbol{c}$. Substitution cancels the two terms containing $\boldsymbol{c}\Delta\boldsymbol{v}$. The remaining terms give (A.1). Since $\boldsymbol{c}\boldsymbol{\omega} > 0$, the original residual vanishes if and only if the transformed residual vanishes. □

Corollary A.1 (Reduction to homogeneous heat flow)**.** *For a solution of the original equation, homogeneous heat flow holds if and only if $\boldsymbol{R}_{\mathrm{alg}} + \boldsymbol{R}_{\mathrm{der}} = 0$ along that solution. A sufficient condition is a constant coefficient $\boldsymbol{c}_0 > 0$ and $\boldsymbol{H}(t, \boldsymbol{x}, D\boldsymbol{v}) = \delta\boldsymbol{c}_0|D\boldsymbol{v}|^2/2$ along the solution. The Hamiltonian $\boldsymbol{H}(p) = |p|^2/2$ satisfies this condition with $\boldsymbol{c}_0 = 1/\delta$.*

Proof. Apply Lemma A.1 with $r_v = 0$. A constant $\boldsymbol{c}$ eliminates the derivative residual, and the quadratic relation eliminates the algebraic residual. □

A constant $\boldsymbol{c}$ is not necessary for a particular solution. For example, $\boldsymbol{v} = 0$ and $\boldsymbol{H}(t, \boldsymbol{x}, 0) = 0$ give $\boldsymbol{\omega} = 1$ and zero total residual for every admissible variable $\boldsymbol{c}$. Requiring both residuals to vanish separately also excludes possible cancellation between them.

Corollary A.2 (Source representation and value recovery)**.** *Suppose $\boldsymbol{\omega}_0$ is bounded and uniformly continuous, and the source $f(s, \boldsymbol{x}) = F(s, \boldsymbol{x}, \boldsymbol{\omega}(s, \boldsymbol{x}), D\boldsymbol{\omega}(s, \boldsymbol{x}))$ is bounded and continuous. Every*

bounded classical solution satisfies

$$\boldsymbol{\omega}(t,\boldsymbol{x}) = \mathbb{E}[\boldsymbol{\omega}_0(\boldsymbol{x}+\sqrt{\delta t}\,Z)] + \int_0^t \mathbb{E}[f(s,\boldsymbol{x}+\sqrt{\delta(t-s)}\,Z)]\,ds, \tag{A.2a}$$

$$\boldsymbol{v}(t,\boldsymbol{x}) = -\frac{1}{\boldsymbol{c}(t,\boldsymbol{x})}\log\boldsymbol{\omega}(t,\boldsymbol{x}), \qquad Z\sim\mathcal{N}(0,I_n). \tag{A.2b}$$

Proof. Variation of constants for the heat equation gives the mild representation. The whole-space heat kernel has unit mass and equals the indicated Gaussian density. Invert the transformation at the evaluation point. □

The source contains the unknown solution. If $|\boldsymbol{v}|\le G$, then $e^{-\boldsymbol{c}_+ G}\le\boldsymbol{\omega}\le e^{\boldsymbol{c}_+ G}$.

Proposition A.1 (Gradient with the full source)**.** *Under the preceding bounded-source hypotheses, for $t>0$ the gradient satisfies*

$$D\boldsymbol{\omega}(t,\boldsymbol{x}) = \frac{1}{\sqrt{\delta t}}\mathbb{E}[Z\boldsymbol{\omega}_0(\boldsymbol{x}+\sqrt{\delta t}\,Z)] + \int_0^t \frac{\mathbb{E}[Zf(s,\boldsymbol{x}+\sqrt{\delta(t-s)}\,Z)]}{\sqrt{\delta(t-s)}}\,ds, \tag{A.3a}$$

$$D\boldsymbol{v}(t,\boldsymbol{x}) = -\frac{D\boldsymbol{\omega}(t,\boldsymbol{x})}{\boldsymbol{c}(t,\boldsymbol{x})\boldsymbol{\omega}(t,\boldsymbol{x})} - \frac{\boldsymbol{v}(t,\boldsymbol{x})}{\boldsymbol{c}(t,\boldsymbol{x})}D\boldsymbol{c}(t,\boldsymbol{x}). \tag{A.3b}$$

Proof. Differentiate the Gaussian kernel with respect to its center. Its derivative has integrable L^1 norm bounded by $\sqrt{n/(\delta(t-s))}$. This bound justifies differentiation of the source integral. The product rule applied to $\boldsymbol{v}=-\boldsymbol{c}^{-1}\log\boldsymbol{\omega}$ gives the second formula. □

(a) Sampling a fixed source

Monte Carlo sampling estimates a specified source integral. It does not evaluate the unknown source without an approximation of the space–time solution. Conditional on a fixed preceding iterate, write f for its deterministic source. For independent $U_i\sim\mathcal{U}(0,1)$ and $Z_i\sim\mathcal{N}(0,I_n)$, define

$$A_i = \boldsymbol{\omega}_0(\boldsymbol{x}+\sqrt{\delta t}Z_i) + tf(tU_i,\boldsymbol{x}+\sqrt{\delta t(1-U_i)}Z_i), \qquad \widehat{\boldsymbol{\omega}}_N = \frac{1}{N}\sum_{i=1}^N A_i. \tag{A.4}$$

The mean is unbiased for one mild update. The two terms may share the Gaussian draw, since independence is needed between the sample pairs, rather than within each pair. Neither $-\boldsymbol{c}^{-1}\log\widehat{\boldsymbol{\omega}}_N$ nor a ratio gradient estimator is generally unbiased. The logarithm requires a positive estimate. Uniform source-time sampling gives a score gradient whose second moment need not be finite near $s=t$. Instead, draw the time from

$$\pi_t(s) = \frac{1}{2\sqrt{t}\sqrt{t-s}}, \quad 0<s<t, \qquad S = t(1-U^2). \tag{A.5}$$

Using independent copies of S, Z, a gradient sample is

$$B = \frac{Z_0\boldsymbol{\omega}_0(\boldsymbol{x}+\sqrt{\delta t}Z_0)}{\sqrt{\delta t}} + 2\sqrt{\frac{t}{\delta}}\,Z_1 f(S,\boldsymbol{x}+\sqrt{\delta(t-S)}Z_1). \tag{A.6}$$

Its mean equals (A.3*a*). If $|\boldsymbol{\omega}_0|\le W$ and $|f|\le B_f$, then its second moment obeys

$$\mathbb{E}|B|^2 \le \frac{2nW^2}{\delta t} + \frac{8ntB_f^2}{\delta}. \tag{A.7}$$

The bound follows from $|a+b|^2\le 2|a|^2+2|b|^2$ and $\mathbb{E}|Z|^2=n$. No bounded-gradient concentration claim follows from a bounded value sample alone.

Endpoint importance sampling remains valid only after including the density ratio for each sampled integral. For example, for a Gaussian displacement h and variance σ^2, the original-to-shifted density ratio is

$$\frac{\varphi_{\boldsymbol{x}}(y)}{\varphi_{\boldsymbol{x}+h}(y)} = \exp\left(\frac{|h|^2 - 2h\cdot(y-\boldsymbol{x})}{2\sigma^2}\right). \tag{A.8}$$

Apply this ratio at each source time with $\sigma^2 = \delta(t-s)$. A tilt of terminal samples alone does not correct the source integral. Importance sampling may change the variance and requires its own moment bounds.

(b) Linear potential solves

A prescribed bounded continuous potential $\boldsymbol{R}$ and bounded initial condition give another representation. For the bounded mild solution, with $X_r = \boldsymbol{x} + \sqrt{\delta}W_r$ for a standard n-dimensional Brownian motion W, the solution of $\boldsymbol{\omega}_t = \delta\Delta\boldsymbol{\omega}/2 + \boldsymbol{R}\boldsymbol{\omega}$ satisfies

$$\boldsymbol{\omega}(t,\boldsymbol{x}) = \mathbb{E}_{\boldsymbol{x}}\left[\boldsymbol{\omega}_0(X_t)\exp\left\{\int_0^t \boldsymbol{R}(t-r,X_r)\,dr\right\}\right]. \tag{A.9}$$

Freezing $\boldsymbol{R}[\boldsymbol{c},\boldsymbol{v}^{(k)},D\boldsymbol{v}^{(k)}]$ produces a linear potential equation, whose positive solution defines a possible iteration. For the full equation, however, $\boldsymbol{R}$ depends on the unknown value and gradient. The formula then remains implicit. Its potential iteration is distinct from the additive source iteration analyzed below, and requires a separate convergence proof. A stable log-sum-exp evaluation applies to positive path weights after time discretization; it does not remove the path quadrature error.

B. Convergence, residual certification, and robustness

This appendix gives conditional convergence and a residual certificate for the full transformed equation.

(a) Convergence of source iteration

Gradient dependence requires a norm that controls both values and first derivatives. Let $\mathrm{BUC}^1(\mathbb{R}^n)$ contain continuously differentiable functions with bounded, uniformly continuous values and gradients. On $X_\tau = C([0,\tau];\mathrm{BUC}^1)$, use

$$\|z\|_{X_\tau} = \sup_{s\le\tau}\|z(s)\|_\infty + \sup_{s\le\tau}\|Dz(s)\|_\infty. \tag{B.1}$$

For $\boldsymbol{g}\in\mathrm{BUC}^1$ and bounded coefficient derivatives, the transformed initial condition belongs to the same space. Define the mild update by

$$(\mathcal{M}z)(t) = S_t\boldsymbol{\omega}_0 + \int_0^t S_{t-s}F(s,\cdot,z(s),Dz(s))\,ds. \tag{B.2}$$

Theorem B.1 (Conditional finite-time convergence)**.** *Let $\mathcal{A}\subset X_\tau$ be nonempty, closed, and invariant under $\mathcal{M}$. Suppose all its members satisfy $a\le z\le b$ and $|Dz|\le Q$, with $a>0$. Assume that the source is bounded and uniformly continuous on the relevant coefficient and state ranges, and satisfies*

$$|F(t,\boldsymbol{x},z,q) - F(t,\boldsymbol{x},\widetilde{z},\widetilde{q})| \le L\big(|z-\widetilde{z}| + |q-\widetilde{q}|\big). \tag{B.3}$$

If $\theta = L(\tau + 2\sqrt{n\tau/\delta}) < 1$, then $\mathcal{M}$ has a unique fixed point in $\mathcal{A}$. Its iterates satisfy

$$\|z^{(K)} - \boldsymbol{\omega}\|_{X_\tau} \le \theta^K\|z^{(0)} - \boldsymbol{\omega}\|_{X_\tau}, \qquad \|z^{(K)} - \boldsymbol{\omega}\|_{X_\tau} \le \frac{\theta}{1-\theta}\|z^{(K)} - z^{(K-1)}\|_{X_\tau} \tag{B.4}$$

for $K \geq 1$. Here finite-time convergence means convergence as $K \to \infty$ on the finite elapsed-time interval $[0, \tau]$; it does not assert termination after finitely many updates. If the fixed point is classical, its inverse solves the original viscous equation.

Proof of Theorems 2.2 and B.1. The space $\mathrm{BUC}^1(\mathbb{R}^n)$ is Banach under the sum of its uniform value and gradient norms. Hence X_τ is Banach, and its closed nonempty subset $\mathcal{A}$ is complete. Invariance makes $\mathcal{M}$ a self-map on this set.

For $z, \widetilde{z} \in \mathcal{A}$, write

$$d(s, \boldsymbol{x}) = F(s, \boldsymbol{x}, z(s, \boldsymbol{x}), Dz(s, \boldsymbol{x})) - F(s, \boldsymbol{x}, \widetilde{z}(s, \boldsymbol{x}), D\widetilde{z}(s, \boldsymbol{x})).$$

The source Lipschitz bound gives $\|d(s)\|_\infty \leq L\|z - \widetilde{z}\|_{X_\tau}$ for every $s \in [0, \tau]$. The common initial-condition terms cancel in $\mathcal{M}z - \mathcal{M}\widetilde{z}$. Since the heat semigroup contracts the uniform norm,

$$\|\mathcal{M}z(t) - \mathcal{M}\widetilde{z}(t)\|_\infty \leq \int_0^t \|d(s)\|_\infty \, ds \leq Lt\|z - \widetilde{z}\|_{X_\tau}. \tag{B.5}$$

Gaussian differentiation, followed by $\mathbb{E}|Z| \leq \sqrt{n}$ for $Z \sim \mathcal{N}(0, I_n)$, yields

$$\|DS_r h\|_\infty \leq \frac{\mathbb{E}|Z|}{\sqrt{\delta r}}\|h\|_\infty \leq \sqrt{\frac{n}{\delta r}}\|h\|_\infty, \qquad r > 0. \tag{B.6}$$

The singularity at $r = 0$ is integrable. Differentiating the source integral therefore gives

$$\|D\mathcal{M}z(t) - D\mathcal{M}\widetilde{z}(t)\|_\infty \leq L\int_0^t \sqrt{\frac{n}{\delta(t-s)}}\, ds \, \|z - \widetilde{z}\|_{X_\tau} = 2L\sqrt{\frac{nt}{\delta}}\|z - \widetilde{z}\|_{X_\tau}. \tag{B.7}$$

Taking the two time suprema and adding proves

$$\|\mathcal{M}z - \mathcal{M}\widetilde{z}\|_{X_\tau} \leq L\left(\tau + 2\sqrt{n\tau/\delta}\right)\|z - \widetilde{z}\|_{X_\tau} = \theta\|z - \widetilde{z}\|_{X_\tau}.$$

The assumption $\theta < 1$ and Banach's fixed-point theorem give a unique fixed point $\boldsymbol{\omega} \in \mathcal{A}$. For each starting iterate $z^{(0)} \in \mathcal{A}$, invariance keeps every subsequent iterate in $\mathcal{A}$. Repeated contraction against the fixed point gives the first estimate in (B.4).

For the posterior estimate, contraction between successive iterates gives

$$\|z^{(K+j+1)} - z^{(K+j)}\|_{X_\tau} \leq \theta^{j+1}\|z^{(K)} - z^{(K-1)}\|_{X_\tau}, \qquad j \geq 0.$$

Sum these differences from $j = 0$ to infinity and use convergence to the fixed point. The geometric series has sum $\theta/(1-\theta)$, proving the second estimate. When $\theta = 0$, the map is constant on $\mathcal{A}$ and one update reaches its fixed point.

Finally, $\boldsymbol{\omega} \geq a > 0$ permits logarithmic inversion. If the fixed point is classical, its mild equation implies the full transformed differential equation. Lemma A.1 then shows that its inverse solves the original viscous equation. □

For bounded coefficient derivatives, $\boldsymbol{c} \geq \boldsymbol{c}_- > 0$, and $z \geq a > 0$, the inverse maps have bounded derivatives on bounded (z, q) ranges. A Hamiltonian uniformly bounded and locally Lipschitz in p on those ranges, uniformly over $(t, \boldsymbol{x})$, therefore supplies a finite source constant L. No convexity assumption is needed for this contraction argument.

The global complexity theorem in [17, Corollary 4.8] does not follow from this local argument. Their setting imposes a globally Lipschitz source in the value and gradient variables, together with strong smoothness and derivative-growth bounds for the complexity result. The transformed source need not meet those assumptions. Even with constant $\boldsymbol{c}$ and the linear Hamiltonian $\boldsymbol{H}(p) = ap$ in one dimension, it becomes

$$F(z, q) = -aq - \frac{\delta}{2z}q^2. \tag{B.8}$$

This source is singular as z approaches zero and is quadratic in q. Using a globally Lipschitz extension to analyze the original equation would require a proof that the solution remains inside the bounded positive region on which the sources agree. Transferring their complexity result

would also require verifying their dimension-dependent constants and quadrature regularity. The present contraction theorem establishes neither requirement.

Invariance is a substantive hypothesis. One sufficient construction starts with $b_0(t) = S_t\boldsymbol{\omega}_0$ and an X_τ ball of radius r about b_0. Choose $r < \inf \boldsymbol{\omega}_0$ to preserve positivity. If $|F| \le B_F$ throughout that ball and

$$B_F\left(\tau + 2\sqrt{n\tau/\delta}\right) \le r, \tag{B.9}$$

then the mild map preserves the ball. Under the preceding uniform-in-state Hamiltonian and coefficient assumptions, bounded initial values and gradients give finite source ranges for sufficiently small τ. The case of a constant $\boldsymbol{c} = 1/\delta$ and $\boldsymbol{H} = |p|^2/2$ has $F = 0$ and $\theta = 0$, so the hypotheses include a nonempty solvable case. Extension to a longer horizon requires renewed bounds and invariance on each interval; a global contraction is not asserted.

If $0 < \theta < 1$ and the initial error is bounded by $E_0 > 0$, a sufficient iteration count is

$$K \ge \max\left\{0, \left\lceil \frac{\log(E_0/\varepsilon)}{\log(1/\theta)} \right\rceil \right\}. \tag{B.10}$$

When $\theta = 0$, one update reaches the fixed point. A small change between noisy iterates alone does not certify convergence.

(b) Sampling and accumulated implementation error

For the fixed-source value samples in (A.4), suppose $a_0 \le \boldsymbol{\omega}_0 \le b_0$ and $|f| \le B_f$. Each sample lies in an interval of width $D = b_0 - a_0 + 2tB_f$. At M fixed queries, Hoeffding's inequality and a union bound give

$$\max_{1\le m\le M} |\widehat{\boldsymbol{\omega}}_N(t_m, \boldsymbol{x}_m) - (\mathcal{M}z)(t_m, \boldsymbol{x}_m)| \le D_* \sqrt{\frac{\log(2M/\eta)}{2N}} \tag{B.11}$$

with probability at least $1 - \eta$, using a common width bound D_*. Independence between different queries is unnecessary. The preceding iterate must be fixed conditional on the past, and the new samples must have the stated conditional law. A union bound over K adaptive iterations replaces M by MK.

(i) Proof of finite-query concentration

Proof of Theorem 2.1. Let $\mathcal{F}_{k-1}$ contain the sampling history before update k, for $k = 1, \dots, K$. Conditional on this history, the preceding iterate z_{k-1} and its source f_{k-1} are fixed. Write $(t_m, \boldsymbol{x}_m)$ for query m, where $m = 1, \dots, M$, and let $A_{k,m,i}$ be its ith sample, for $i = 1, \dots, N$.

Independence of the uniform time and Gaussian variables in the sampling law, together with boundedness of the source, permits integration in either order. The conditional sample mean is therefore

$$\mu_{k,m} := \mathbb{E}[A_{k,m,i} \mid \mathcal{F}_{k-1}] = S_{t_m}\boldsymbol{\omega}_0(\boldsymbol{x}_m) + \int_0^{t_m} S_{t_m - s} f_{k-1}(s)(\boldsymbol{x}_m)\, ds = (\mathcal{M}z_{k-1})(t_m, \boldsymbol{x}_m). \tag{B.12}$$

This is the prescribed mild update, rather than the unknown converged solution. Define its Monte Carlo estimate by $\widehat{\mu}_{k,m} = N^{-1}\sum_{i=1}^{N} A_{k,m,i}$.

Each sample lies between $a_0 - t_m B_F$ and $b_0 + t_m B_F$. Its interval width is at most the common deterministic bound D_*. If $D_* = 0$, all samples equal their conditional means almost surely, and the theorem follows with $h_N = 0$. Suppose henceforth that $D_* > 0$. Conditional independence within the batch and Hoeffding's inequality give, for every $h > 0$,

$$\mathbb{P}\left(|\widehat{\mu}_{k,m} - \mu_{k,m}| > h \mid \mathcal{F}_{k-1}\right) \le 2\exp\left(-\frac{2Nh^2}{D_*^2}\right). \tag{B.13}$$

Taking expectations removes the conditioning and preserves this upper bound. A union bound over the at most MK batches now yields

$$\mathbb{P}\left(\max_{k,m} |\widehat{\mu}_{k,m} - \mu_{k,m}| > h\right) \leq 2MK \exp\left(-\frac{2Nh^2}{D_*^2}\right). \tag{B.14}$$

Choosing $h = h_N$ from (2.20) makes the right-hand side equal to η. If the algorithm stops before K updates, its executed batches form a subset of those covered by the same bound. No independence between queries or updates is used; only the stated conditional independence within each batch is required. □

(ii) Proof of sample count and logarithmic recovery

Proof of Corollary 2.1. The sample-count condition (2.21) implies $h_N \leq \varepsilon_\omega$ by direct substitution into (2.20). Theorem 2.1 therefore bounds all covered transformed-value sampling errors by ε_ω on an event of probability at least $1 - \eta$.

On that same event, consider a target update $\mu \geq a > 0$ and its estimate $\widehat{\mu}$, with $|\widehat{\mu} - \mu| \leq h_N < a$. Then $\widehat{\mu} \geq a - h_N > 0$, so both logarithms are defined. Every point between μ and $\widehat{\mu}$ is at least $a - h_N$. Integrating the derivative of the logarithm along this interval gives

$$|\log \widehat{\mu} - \log \mu| \leq \frac{|\widehat{\mu} - \mu|}{a - h_N} \leq \frac{h_N}{a - h_N}.$$

At the query point, logarithmic inversion multiplies this difference by $1/\boldsymbol{c}(t, \boldsymbol{x}) \leq 1/\boldsymbol{c}_-$. Hence

$$\left| -\frac{\log \widehat{\mu}}{\boldsymbol{c}(t, \boldsymbol{x})} + \frac{\log \mu}{\boldsymbol{c}(t, \boldsymbol{x})} \right| \leq \frac{h_N}{\boldsymbol{c}_-(a - h_N)}. \tag{B.15}$$

The inequality holds simultaneously at every covered query satisfying the lower-bound hypothesis, without another failure-probability allocation. It bounds recovery of the prescribed update and makes no assertion of unbiased logarithmic recovery or error relative to the converged solution. □

(iii) Accumulated implementation error

The prototype instead uses the source-time density π_t for both value and gradient samples. Its value-source weight is $1/\pi_t(S) \leq 2t$, so the corresponding interval width is $b_0 - a_0 + 4tB_f$. The same concentration argument applies with that larger width.

This is a finite-query value bound. A continuum X_τ guarantee also requires gradient concentration and control of reconstruction between queries. Gaussian score samples are unbounded, so the same Hoeffding argument does not establish that guarantee. The second-moment bound in (A.7) supports variance estimates, but does not by itself control a continuum supremum.

If a numerical implementation separately establishes X_τ errors ϵ_k in each mild update and keeps its iterates in the admissible set, then contraction gives

$$\|\widehat{z}^{(K)} - \boldsymbol{\omega}\|_{X_\tau} \leq \theta^K E_0 + \sum_{j=0}^{K-1} \theta^{K-1-j} \epsilon_j. \tag{B.16}$$

The proof iterates the inequality $e_{k+1} \leq \theta e_k + \epsilon_k$. Thus sampling, source quadrature, and reconstruction errors accumulate through the solve. They do not enter only as a final independent sample average. If both transformed values are at least a, logarithmic inversion yields

$$\|\widehat{\boldsymbol{v}}^{(K)} - \boldsymbol{v}^\delta\|_\infty \leq \frac{\theta^K E_0 + \sum_{j=0}^{K-1} \theta^{K-1-j} \epsilon_j}{\boldsymbol{c}_- a}. \tag{B.17}$$

This conditional bound does not certify the accompanying numerical prototype, which uses time stepping rather than whole-interval Picard iteration.

(c) A posteriori certificate for the original equation

A residual of the full transformed equation measures the residual of the original equation directly. Assume comparison for the original parabolic equation on the whole space. Let U be a bounded classical approximation with bounded residual, and define

$$w = e^{-\boldsymbol{c}U}, \qquad d = w_t - \frac{\delta}{2}\Delta w - F(t, \boldsymbol{x}, w, Dw). \tag{B.18}$$

Theorem B.2 (Residual certificate)**.** *Under these hypotheses, the viscous solution with initial condition* $\boldsymbol{g}$ *obeys*

$$\|U(t) - \boldsymbol{v}^{\delta}(t)\|_{\infty} \le \|U(0) - \boldsymbol{g}\|_{\infty} + \int_0^t \left\| \frac{d(s)}{\boldsymbol{c}(s)w(s)} \right\|_{\infty} ds. \tag{B.19}$$

Proof. Lemma A.1 gives $d = -\boldsymbol{c}wr_U$, with $r_U = U_t + \boldsymbol{H}(t, \boldsymbol{x}, DU) - \delta\Delta U/2$. Let $e_0 = \|U(0) - \boldsymbol{g}\|_{\infty}$ and $E(t) = e_0 + \int_0^t \|r_U(s)\|_{\infty} ds$. Since the Hamiltonian does not depend on the value, $U - E$ is a subsolution and $U + E$ is a supersolution. Their initial conditions bracket $\boldsymbol{g}$. Comparison gives $U - E \le \boldsymbol{v}^{\delta} \le U + E$, proving the bound. □

The source must include both residual components when computing d. If a solver omits $\boldsymbol{R}_{\mathrm{der}}$, then its transformed residual contains the missing term $-w\boldsymbol{R}_{\mathrm{der}}$. The certificate charges that omission. Residuals evaluated only at sampled states are diagnostic; rigorous use requires a bound between those states and through time, including derivative approximation error.

(d) Viscosity, model perturbations, and level sets

Vanishing viscosity remains a separate approximation. Suppose the Hamiltonian and initial condition satisfy a comparison and approximation theorem giving $\|\boldsymbol{v}^{\delta} - \boldsymbol{v}^{0}\|_{\infty} \le C_{\mathrm{vis}}\sqrt{\delta}$ on the chosen horizon. This rate is an additional hypothesis, rather than a consequence of the transformation. Under the same uniform norm and domain, Theorem B.2 then gives

$$\|U(t) - \boldsymbol{v}^{0}(t)\|_{\infty} \le \|U(0) - \boldsymbol{g}\|_{\infty} + \int_0^t \left\| \frac{d(s)}{\boldsymbol{c}(s)w(s)} \right\|_{\infty} ds + C_{\mathrm{vis}}\sqrt{\delta}. \tag{B.20}$$

No universal choice $\delta \sim N^{-1/3}$ follows: source, stability, and sampling constants can depend on δ. Their dependence must be established before optimizing a combined rate.

Proposition B.1 (Hamiltonian and initial condition perturbations)**.** *Suppose two equations with the same diffusion admit comparison, their initial conditions differ by at most* ϵ_g*, and* $|\boldsymbol{H}_1(t, \boldsymbol{x}, p) - \boldsymbol{H}_2(t, \boldsymbol{x}, p)| \le \epsilon_H$ *throughout the required co-state range. For classical solutions, it suffices that this range contains* $D\boldsymbol{v}_2$*. Then*

$$\|\boldsymbol{v}_1(t) - \boldsymbol{v}_2(t)\|_{\infty} \le \epsilon_g + t\epsilon_H. \tag{B.21}$$

Proof. The functions $\boldsymbol{v}_2 - \epsilon_g - t\epsilon_H$ and $\boldsymbol{v}_2 + \epsilon_g + t\epsilon_H$ are a subsolution and a supersolution for the first equation. Comparison proves the claim. For viscosity solutions, impose the Hamiltonian bound on the test-gradient range required by comparison, or uniformly over all p. □

Changing a prescribed positive transformation coefficient does not perturb the Hamiltonian when the full source is retained. Thus a smooth positive regularization that preserves the stated differentiability changes coordinates; it creates no Hamiltonian modeling error by itself. Hard clipping can destroy that differentiability and falls outside these classical results. If the Hamiltonian is deliberately replaced, the perturbation bound applies to that separate change. The matching choice $\boldsymbol{c} = 2\boldsymbol{H}(t, \boldsymbol{x}, D\boldsymbol{v})/(\delta|D\boldsymbol{v}|^2)$ can vanish, change sign, or become singular. At $\boldsymbol{c} = 0$, the exponential transform loses the value information. Recomputing this ratio also requires regularity beyond the prescribed-coefficient hypotheses. The full source permits a positive constant coefficient even in these cases.

If a uniform bound $|U - \boldsymbol{v}^0| \le E$ has been established on a specified domain, then

$$\{U < -E\} \subseteq \{\boldsymbol{v}^0 < 0\}, \qquad \{U > E\} \subseteq \{\boldsymbol{v}^0 > 0\}. \tag{B.22}$$

Proof of Corollary 2.2. Fix a point in the specified set and suppress its arguments. The assumed error bound is equivalent to

$$U - E \le \boldsymbol{v}^0 \le U + E. \tag{B.23}$$

If $U > E$, then the lower bound gives $\boldsymbol{v}^0 \ge U - E > 0$. Thus the point lies outside the target-reaching class $\{\boldsymbol{v}^0 \le 0\}$. If $U < -E$, then the upper bound gives $\boldsymbol{v}^0 \le U + E < 0$, placing the point in the strict negative sublevel set. Since the argument applies at every point covered by the error bound, it proves both inclusions in (B.22).

The states assigned neither strict sign by this rule satisfy $-E \le U \le E$, which is the abstention band $|U| \le E$. At its boundaries, the corresponding estimate in (B.23) is nonstrict, so the corollary makes no strict sign assertion there.

If the assumed error bound holds simultaneously on an event of probability at least $1 - \eta$, both sign implications hold simultaneously on that same event. No additional union bound is needed. The domain of the sign certificate is the domain on which that error bound has been established; a bound at finitely many queries supplies no assertion between those queries. □

Interpreting either sign as safety also requires the corresponding game's value convention and, for reach–avoid problems, a correctly solved obstacle problem.

C. Further Numerical Results

The model descriptions connect the game experiments to their dynamical and geometric interpretation. All recomputed fields use the full residual and the elapsed-time convention of Section 2. Numerical discrepancies and sampled equation residuals are distinguished from uniform value-error certificates.

(a) Relative Dubins pursuit–evasion

Merz's pursuit–evasion game extends the homicidal chauffeur problem to two vehicles with comparable speeds and turn radii [19, 25]. Its reachable-set geometry is a benchmark for numerical reachability [26]. The present experiment computes terminal-horizon values at positive viscosity; the classical barrier analysis supplies context rather than a reference solution for that regularized problem.

Let $\boldsymbol{x} = (x_1, x_2, x_3)$ be the pursuer's position and heading relative to an evader-centered frame. With constant forward speeds $v_e, v_p > 0$ and steering rates $w_e, w_p \in [-w, w]$, the physical-time dynamics are

$$\dot{x}_1 = -v_e + v_p \cos x_3 + w_e x_2, \quad \dot{x}_2 = v_p \sin x_3 - w_e x_1, \quad \dot{x}_3 = w_p - w_e. \tag{C.1}$$

The numerical specialization sets $v_e = v_p = w = 1$. For matched headings, identical steering signals preserve the initial separation. A barrier characterization requires the separate game analysis of [25, 26].

The capture target is the cylinder

$$\mathcal{T} = \{\boldsymbol{x} : x_1^2 + x_2^2 \le r^2\}, \tag{C.2}$$

with capture radius $r > 0$. Its definition does not constrain heading. This target differs from a backward reachable tube, which includes states from which the pursuer can force target entry during an interval. The bounded initial condition in Section 3(c) represents the target through its zero sublevel set.

A separate single-vehicle example uses $\dot{x}_1 = \cos\theta$, $\dot{x}_2 = \sin\theta$, and $\dot{\theta} = u$, with $|u| \le 1$ and a pose target that includes heading. Its elapsed Hamiltonian is $\boldsymbol{H} = -p_1 \cos\theta - p_2 \sin\theta + |p_3|$. Figure 2 compares the full residual computation with the independent same-viscosity reference. The elapsed-

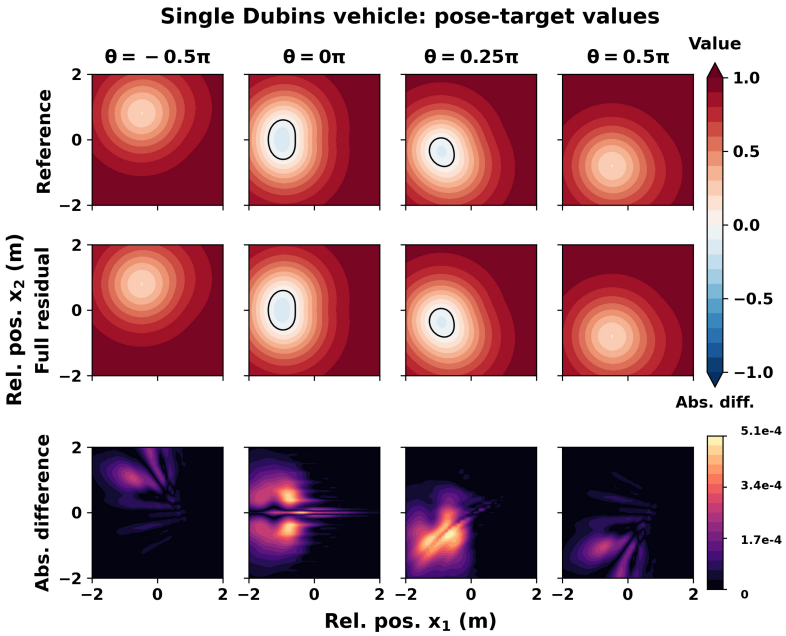


Figure 2: Single-Dubins pose-target values at four headings, $T = 1$, and $\delta = 0.08$. Rows show the independent finite-difference reference, the full residual computation, and their absolute difference. Position axes are in meters under the reference scales in Section 3(c); payoff is normalized. Black contours mark zero where present.

time regularization of the reachable-tube equation uses

$$\ell_t + \max\{0, \boldsymbol{H}(x, D\ell)\} = \frac{\delta}{2}\Delta\ell. \tag{C.3}$$

Optimizing the physical-time dynamics gives the Hamiltonian $\boldsymbol{H}_{\text{phys}} = \max_{w_e} \min_{w_p} p \cdot f$:

$$\boldsymbol{H}_{\text{phys}}(\boldsymbol{x}, p) = p_1(v_p \cos x_3 - v_e) + p_2(v_p \sin x_3) + w|p_1 x_2 - p_2 x_1 - p_3| - w|p_3|. \tag{C.4}$$

With the normalized bounds $v_e = v_p = w = 1$, the physical Hamiltonian becomes

$$\boldsymbol{H}_{\text{phys}}(\boldsymbol{x}, p) = p_1(\cos x_3 - 1) + p_2(\sin x_3) + |p_1 x_2 - p_2 x_1 - p_3| - |p_3|. \tag{C.5}$$

For elapsed backward time, the numerical Cauchy problem uses $\boldsymbol{H} = -\boldsymbol{H}_{\text{phys}}$. For unit bounds this is

$$\boldsymbol{H}(\boldsymbol{x}, p) = p_1(1 - \cos x_3) - p_2 \sin x_3 - |p_1 x_2 - p_2 x_1 - p_3| + |p_3|. \tag{C.6}$$

The recomputed terminal-horizon fields use this Hamiltonian without a stopping operator. A reachable tube requires the separate formulation in Section 2.

The independent-reference comparison is shown in Figure 3; Table 1 reports its refinement checks.

(b) The double integrator

The double integrator illustrates bounded-acceleration reachability [5, 36]. Its classical problem minimizes time to the origin; our experiment uses a fixed horizon and a bounded target-ball payoff. Writing $\ddot{\boldsymbol{x}} = \boldsymbol{u}$ in position–velocity coordinates gives

$$\dot{\boldsymbol{x}}_1(t) = \boldsymbol{x}_2(t), \qquad \dot{\boldsymbol{x}}_2(t) = \boldsymbol{u}(t), \qquad |\boldsymbol{u}(t)| \le 1. \tag{C.7}$$

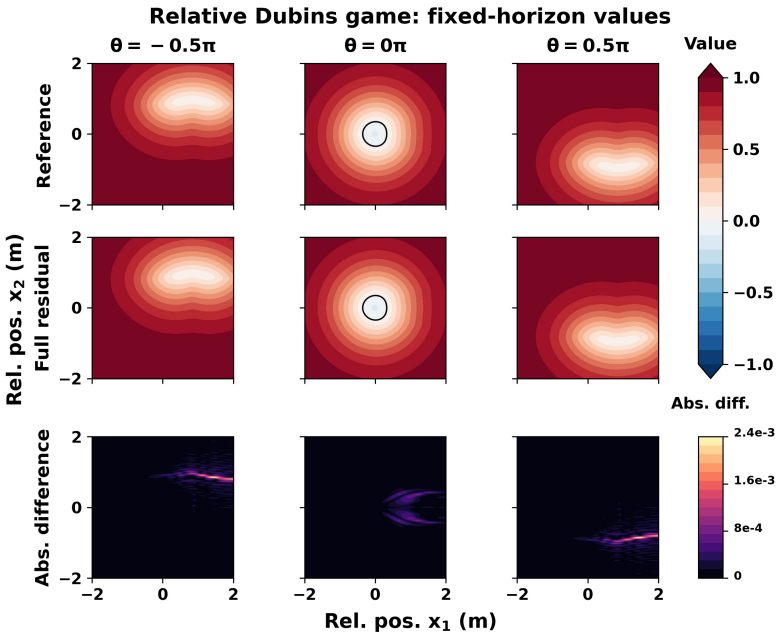


Figure 3: Relative-Dubins terminal-horizon values at $T = 1$ and $\delta = 0.08$. Rows show the independent finite-difference reference, direct integration of the full transformed equation, and their absolute difference. Columns select three headings. Black contours mark zero; position axes are in meters and payoff is normalized. The comparison uses the same Hamiltonian, initial condition, and viscosity in both solvers.

(i) Time-optimal control scheme

The minimizing physical control possesses the Hamiltonian $\boldsymbol{H}_{\text{phys}}(\boldsymbol{x}, p, \boldsymbol{u}) = p_1\boldsymbol{x}_2 + p_2\boldsymbol{u}$ whose minimization gives $\boldsymbol{H}_{\text{phys}} = p_1\boldsymbol{x}_2 - |p_2|$. The elapsed-time Hamiltonian and the minimizing physical control, for $p_2(t) \neq 0$, are

$$\boldsymbol{H}(\boldsymbol{x}, p) = -p_1\boldsymbol{x}_2 + |p_2|, \qquad \boldsymbol{u}^*(t) = -\operatorname{sign}(p_2(t)). \tag{C.8}$$

Any admissible control minimizes the instantaneous Hamiltonian when $p_2(t) = 0$.

The co-state equations give $\dot{p}_1 = 0$ and $\dot{p}_2 = -p_1$, so a nontrivial time-optimal control has at most one switch. Its constant-control sequences are $\{[+1], [-1], [+1, -1], [-1, +1]\}$. The origin itself requires no control interval. On each constant-control segment, $t \in [t_0, t_f]$, write $\boldsymbol{u}(t) = k \in \{-1, +1\}$ so that for initial conditions $\boldsymbol{x}_1(t_0) = \boldsymbol{\xi}_1$ and $\boldsymbol{x}_2(t_0) = \boldsymbol{\xi}_2$, integration gives

$$\boldsymbol{x}_1(t) = \boldsymbol{\xi}_1 + \frac{1}{2}k\left(\boldsymbol{x}_2^2 - \boldsymbol{\xi}_2^2\right), \text{ for } t - t_0 = k\left(\boldsymbol{x}_2(t) - \boldsymbol{\xi}_2\right). \tag{C.9}$$

Figure 4 shows the full residual double-integrator fields at four elapsed horizons. The numerical protocol and sensitivity checks are given in Section 3(c). Equation (C.9) describes each constant-control segment; concatenating segments with opposite controls gives the switching trajectories of the classical time-optimal problem [5, 36]. The numerical refinement and independent-reference comparisons are reported in Table 1.

(c) A normalized rockets benchmark

The rocket-launch examples of Dreyfus [12] and Molu [29] motivate a normalized controlled benchmark. Its numerical definition is the following normalized relative-state system:

$$\dot{x}_1 = a\cos\theta + u_e x_1, \quad \dot{x}_2 = a\sin\theta + a + u_p x_1 - g, \quad \dot{\theta} = u_p - u_e. \tag{C.10}$$

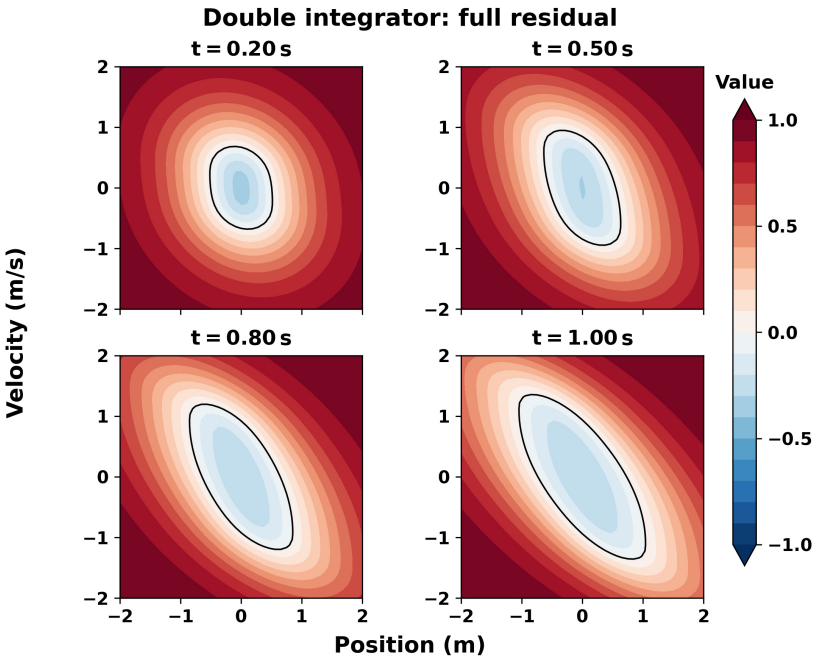


Figure 4: Full residual double-integrator values at elapsed horizons $t = 0.20, 0.50, 0.80$, and 1.00, with $\delta = 0.08$. Black contours mark the zero level set. Position is in meters, velocity in meters per second, and elapsed time in seconds under the reference scales in Section 3(c); payoff is normalized.

Here $\theta \in \mathbb{S}^1$ is an independent relative-heading state, and $u_e, u_p \in [-1, 1]$ are rate controls. The evader maximizes the payoff and the pursuer minimizes it. With $p = D\boldsymbol{v}$, collection of the control terms gives

$$p \cdot f = ap_1 \cos\theta + (a + a\sin\theta - g)p_2 + (p_1 x_1 - p_3)u_e + (p_2 x_1 + p_3)u_p. \tag{C.11}$$

Thus elapsed backward time uses

$$\boldsymbol{H} = -\max_{u_e}\min_{u_p} p \cdot f = -ap_1\cos\theta + (g - a - a\sin\theta)p_2 - |p_1 x_1 - p_3| + |p_2 x_1 + p_3|. \tag{C.12}$$

The recomputation sets $a = 1$, $g = 2$, target radius $r = 1.5$, viscosity $\delta = 0.08$, and horizon $T = 0.5$. Figure 1 compares the full residual solution with the independent reference at the same viscosity. The physical rocket model and its illustrations are given in [29]. The relative-Dubins comparison in Figure 3 supplies the principal pursuit–evasion example in this appendix.

(d) Population-scale evaluation for starling-inspired dynamics

Starling murmurations motivate the use of evolving interfaces to describe collective geometry. Field observations associate flock interactions with a fixed number of nearby birds rather than a fixed metric radius [3]; heading correlations provide further evidence of collective organization [6]. These observations motivate the model, but do not establish collision avoidance or predator evasion for the computed fields. Figure 5 illustrates the biological context.

(i) Notation and background

Let $\Omega \subseteq \mathbb{R}^n$ denote the state domain and $\boldsymbol{v} : [0, T] \times \Omega \to \mathbb{R}$ a scalar value field. A subgroup occupies a region $\mathcal{S}_i \subseteq \Omega$, and the collection occupies $\mathcal{C} = \bigcup_i \mathcal{S}_i$. When regions are disjoint,

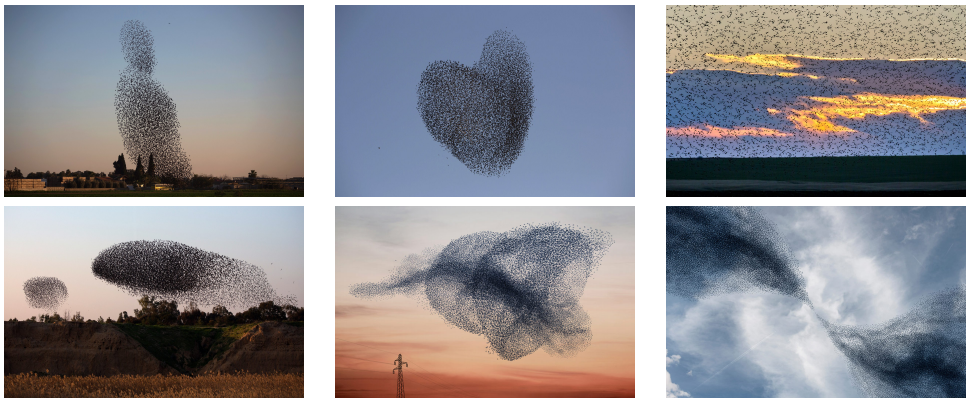


Figure 5: Photographs illustrate collective behaviors of European starlings (*Sturnus vulgaris*); they are not outputs of the numerical solver. Reading across each row. (i) A flock of starlings rises into the air, in a dense structure (Reuters/Amir Cohen). (ii) Starlings migrating over an Israeli village (AP Photo/Oded Balilty). (iii) Starlings feeding on laid seeds on the ground in Romania. (iv) Two flocks of migrating starlings (Menahem Kahana/AFP/Getty Images). (v) A concentric conical formation of starlings (Courtesy of The Gathering Site.). (vi) Splitting and joining of a flock of starlings.

$\mathcal{S}_i \cap \mathcal{S}_j = \emptyset$ for $i \neq j$. The boundary of a region is denoted by $\partial\mathcal{S}_i$. This geometric organization does not itself define control coupling between birds.

An individual bird has position and heading state $\boldsymbol{x}^{(i)} = (x^{(i)}, y^{(i)}, z^{(i)}, \theta^{(i)})$. The evader and pursuer maximize and minimize the payoff, respectively.

(ii) Murmuration as evolving level sets

The zero sublevel set of $\boldsymbol{v}_j$ represents a region, and its zero level set represents the interface. For prescribed normal speed a_j, $\partial_t \boldsymbol{v}_j + a_j |D\boldsymbol{v}_j| = 0$. A game replaces this prescribed-speed Hamiltonian by one obtained from the controlled vector field. Level-set methods accommodate connectivity changes without explicitly reconnecting tracked surfaces. Related representations include front tracking, Voronoi implicit interface methods, and volume-of-fluid methods [21, 22, 34, 37]. Their interface descriptions provide context for the diagnostics below.

(iii) Individual kinematics and neighborhood rules

Each topological neighborhood contains the n_c nearest agents, ranked by spatial distance with a fixed tie rule, independently of their labels [3]. This neighborhood model is motivated by local coordination [20].

Definition C.1 (Flock)**.** A flock F is a finite collection of agents together with a specified neighborhood rule and a law governing their interactions.

For context, a candidate kinematic model with constant forward speed $v > 0$, vertical control $u_z^{(i)}$, and steering rate $w^{(i)}$ is

$$\dot{\boldsymbol{x}}^{(i)} = \left[v\cos\theta^{(i)}, v\sin\theta^{(i)}, u_z^{(i)}, \overline{w}^{(i)} \right]^\top, \qquad \overline{w}^{(i)} = \frac{w^{(i)} + \sum_{j\in\mathcal{N}_i(t)} w^{(j)}}{1 + |\mathcal{N}_i(t)|}. \tag{C.13}$$

Here $\mathcal{N}_i(t)$ is the chosen neighbor set. This expression averages steering rates, not heading angles. It specifies a possible coupling law; the numerical experiment below instead uses independently bounded rate controls in an individual relative-state game.

(iv) Framework for separated payoffs

Definition C.2 (Union payoff)**.** If continuous agent payoffs Φ_j describe target sets by $\{\Phi_j \le 0\}$, their union is represented by the scalar payoff $\Phi_{\text{union}} = \min_j \Phi_j$.

The minimum represents a union target, but taking it need not commute with optimization under shared controls. The bird experiment evolves one initial condition formed from seven disks, rather than combining independently optimized solutions.

(v) Local geometry and individual pursuit

Local neighbor interactions motivate descriptions of subgroup geometry, but coordinating individual pursuit games into a jointly controlled flock requires a coupling model. Equal-speed vehicles with matched headings can preserve their separation by applying identical admissible steering signals. That kinematic observation alone does not characterize a game barrier [25]. The following relative-coordinate model isolates the individual pursuit problem used in the computation.

For the numerical individual model, write $\boldsymbol{x} = (\boldsymbol{x}_1, \boldsymbol{x}_2, z, \theta)$. The physical-time dynamics are

$$\dot{\boldsymbol{x}} = \left(-1 + \cos\theta + w_e \boldsymbol{x}_2,\ \sin\theta - w_e \boldsymbol{x}_1,\ u_p^z - u_e^z,\ w_p - w_e\right)^\top, \tag{C.14}$$

with $w_e, w_p \in [-1, 1]$ and $u_e^z, u_p^z \in [-\gamma, \gamma]$, where $\gamma = 0.5$. The altitude state satisfies $z \in \mathbb{R}$, and $\theta \in \mathbb{S}^1$. The numerical model uses an independent admissible rate w_e as a surrogate for neighbor-dependent averaged steering. The population-query experiment evaluates this individual controlled field; it does not simulate the interacting neighbor law.

For an evader maximizing and a pursuer minimizing the payoff, the vertical contribution to the physical Hamiltonian is $\gamma_e|p_z| - \gamma_p|p_z|$. Equal vertical bounds cancel it. The altitude-independent initial condition then gives an altitude-independent solution, whose three-dimensional elapsed Hamiltonian is (C.6). The reversal $\boldsymbol{H} = -\boldsymbol{H}_{\text{phys}}$ is applied once when forming the elapsed-time Cauchy equation. Sampling altitude in $[-1, 1]$ specifies the query distribution and imposes no altitude state constraint.

(vi) Aggregation of value fields

For scalar payoffs, define a union representation by

$$\Phi_1 \cup \Phi_2 := \min\{\Phi_1, \Phi_2\}. \tag{C.15}$$

The symbol denotes an operation on scalar representations, whose zero sublevel sets satisfy the usual set-union identity.

Remark C.1 (Proposed aggregation of flock values)**.** A pointwise minimum combines scalar functions whose zero sublevel sets represent a union. For separately computed fields, the proposed display aggregate is

$$\Phi_{\text{union}}(\boldsymbol{x}, t) = \min_j \Phi_j(\boldsymbol{x}, t), \qquad \{\Phi_{\text{union}} \le 0\} = \bigcup_j \{\Phi_j \le 0\}. \tag{C.16}$$

The set identity follows from the definition of a minimum. It does not imply that a union or minimum of Hamiltonians governs a coupled flock game. Optimizing separately and then taking a minimum need not commute with optimizing a shared strategy. A joint model requires a specified payoff, state, control coupling, and dynamic-programming argument.

For each declared individual model, its physical-time game Hamiltonian has the control form

$$\boldsymbol{H}^{(i)}_{\text{phys}}(t, \boldsymbol{x}, \boldsymbol{p}) = \max_{\boldsymbol{u} \in \mathcal{U}_i} \min_{\boldsymbol{w} \in \mathcal{W}_i} \boldsymbol{p} \cdot f^{(i)}(t, \boldsymbol{x}, \boldsymbol{u}, \boldsymbol{w}). \tag{C.17}$$

The preceding dynamics specify candidate individual vector fields. Their coordinate conventions and control bounds must agree with the chosen numerical model before this formula defines a

computed value. An individual Hamiltonian is not automatically a mechanical energy, and a set union of its scalar values is not a scalar differential operator.

Remark C.2 (Scope of the population model)**.** The illustrations combine low-dimensional agent or flock calculations. They do not establish a Hamilton–Jacobi equation on the full joint population state. Identifying a population game with those aggregates would require a decomposition theorem or a separately verified approximation bound. Neither follows from signed-distance initialization or equal agent speeds.

(vii) Vacuole nucleation diagnostics

A hole in a displayed zero sublevel set can serve as a geometric marker of a vacuole. Predator entry and a local Hamiltonian threshold alone do not force that topology. Let $\mathcal{L}_h(t)$ denote a finite planar cubical approximation of the displayed sublevel set.

Definition C.3 (Euler characteristic of a displayed sublevel set)**.** For consistent vertex, edge, and face counts V, E, F of $\mathcal{L}_h(t)$, define

$$\chi_h(t) = V - E + F = \beta_0(t) - \beta_1(t). \tag{C.18}$$

The counts β_0 and β_1 denote connected components and independent holes of that planar complex. Grid connectivity and boundary treatment form part of the definition [14].

Remark C.3 (Local thresholds and topology)**.** A condition such as $\boldsymbol{H}_f(\boldsymbol{x}^*, \boldsymbol{p}^*) = 0$ does not identify a critical point of $\boldsymbol{v}$, a Morse saddle, or a change in a global sublevel set. A smooth Morse-theoretic argument would need these additional hypotheses and a specified parameter crossing. Even an index-one handle can merge two components rather than create a hole. A vacuole interpretation therefore requires spatial evidence of an enclosed complementary region, rather than a local Hamiltonian test alone.

(viii) Flock splitting diagnostics

Remark C.4 (Local motion and global disconnection)**.** A negative attacked-agent Hamiltonian can describe local motion under the model's sign convention. It is neither a necessary nor a sufficient condition for a displayed reachable set to become disconnected. Global disconnection depends on the entire field and domain. A rise in the measured component count is a proposed fragmentation marker; linking it to a particular control threshold or game equilibrium requires a separate argument.

(ix) Cordon geometry

Definition C.4 (Annular display region)**.** For $0 < r_{\text{in}} < r_{\text{out}}$, define the horizontal annular region

$$\mathcal{C} = \{\boldsymbol{x} : r_{\text{in}} \le \|\boldsymbol{x}_{1:2}\| \le r_{\text{out}}\}. \tag{C.19}$$

Remark C.5 (Topology does not certify a barrier)**.** An annular displayed component has one hole, but that geometry alone gives no invariant barrier or finite-time capture guarantee. Such conclusions require the appropriate dynamical boundary condition and value-error control. A cordon interpretation must also verify which agents or regions the displayed loop encloses.

(x) Flash expansion diagnostics

Remark C.6 (Trajectory speed and sublevel-set motion)**.** A horizontal speed bound v_e bounds an individual free trajectory's horizontal displacement by $v_e \Delta t$ over an interval of length Δt. It does not by itself bound the motion of a relative-game sublevel set, which also depends on pursuer motion, controls, the target, and time orientation. A reachable-tube rate requires an additional argument. Changes in the displayed radius remain descriptive diagnostics.

(xi) Markers in numerical slices

Proposed geometric diagnostics use a two-dimensional slice of the four-dimensional value approximation. Fix the remaining state coordinates, a grid, a sign threshold, and a connectivity convention. Compute the resulting cubical complex, its connected components, and its enclosed complementary components. Record χ_h, β_1, and n_{comp} for each sampled time.

Remark C.7 (Observational event labels)**.** An increase in β_1 suggests a new displayed hole, and an increase in n_{comp} suggests fragmentation of the displayed set. A change in χ_h alone cannot distinguish those mechanisms. Calling a hole a cordon also requires an enclosure test. These labels describe the chosen slice and threshold; they neither classify every flock behavior nor establish a disjoint partition of the physical state space.

Grid resolution, slicing, smoothing, boundary conditions, and numerical value error can change these counts. Robust event identification requires sensitivity analysis and evidence that the relevant signs and connections persist under the numerical error budget. No such topological certification follows from the retained-source Cauchy theorem alone.

(xii) Numerical geometry and its interpretation

Figure 6 compares the displayed zero sublevel geometry with the independent reference. On the 193×193 planar grid at $\theta = 0$, both solvers give the same component, hole, and Euler counts at all 11 sampled times. The displayed set has one connected component throughout. Its hole persists through $t = 1.4$ and is absent at $t = 1.6$, placing the observed change between these sampled times. This agreement supports the geometric description of this slice; it does not locate a continuum transition time or certify an invariant cordon.

Figure 7 displays a reproducible subset of 2,000 states from the 100,000-state population. Color encodes each state's heading. The overlaid zero contour is the $\theta = 0$ slice at $t = 2$; each population value instead uses that state's own heading. The contour supplies spatial context and cannot classify every displayed state. The paired value and sign comparisons in Section 3(b) quantify the numerical agreement relevant to that distinction.

(e) Validation of the retained transformation source

An analytic advection–diffusion solution isolates the effect of retaining the source. On the periodic interval $[0, 2\pi)$, choose $\delta = 0.4$, $\boldsymbol{H}(\boldsymbol{p}) = 0.7\boldsymbol{p}$, initial condition $\boldsymbol{g}(\boldsymbol{x}) = \sin \boldsymbol{x}$, and prescribed coefficient $\boldsymbol{c}(t, \boldsymbol{x}) = 1 + 0.2 \cos \boldsymbol{x} + 0.1 \sin t$. The reference solution is

$$\boldsymbol{v}^{\delta}(t, \boldsymbol{x}) = e^{-\delta t/2} \sin(\boldsymbol{x} - 0.7t). \tag{C.20}$$

At $t = 0.5$, Fourier collocation uses 128 states and exponential-Euler integration uses 4 resolutions, from 100 to 800 steps. The full-source and homogeneous-heat comparisons share the coefficient, initial condition, grid, and time resolutions.

Figure 8 reports the maximum value and gradient errors at the evaluation states. With the full source, observed value-error orders range from 1.000 to 1.001. The finest value error is 1.003×10^{-4}, compared with 3.839×10^{-1} for the homogeneous heat solve. These measurements support source retention and temporal consistency on this smooth periodic problem. They do not establish a continuum error certificate, whole-space accuracy, a game benchmark, or high-dimensional computational scaling.

(i) Multidimensional analytic verification

The multidimensional Fourier implementation is checked against

$$\boldsymbol{v}_{\mathrm{ref}}(t, \boldsymbol{x}) = e^{-\delta d t/2} \prod_{j=1}^{d} \sin(\boldsymbol{x}_j - a_j t), \qquad a_j = 0.15j, \tag{C.21}$$

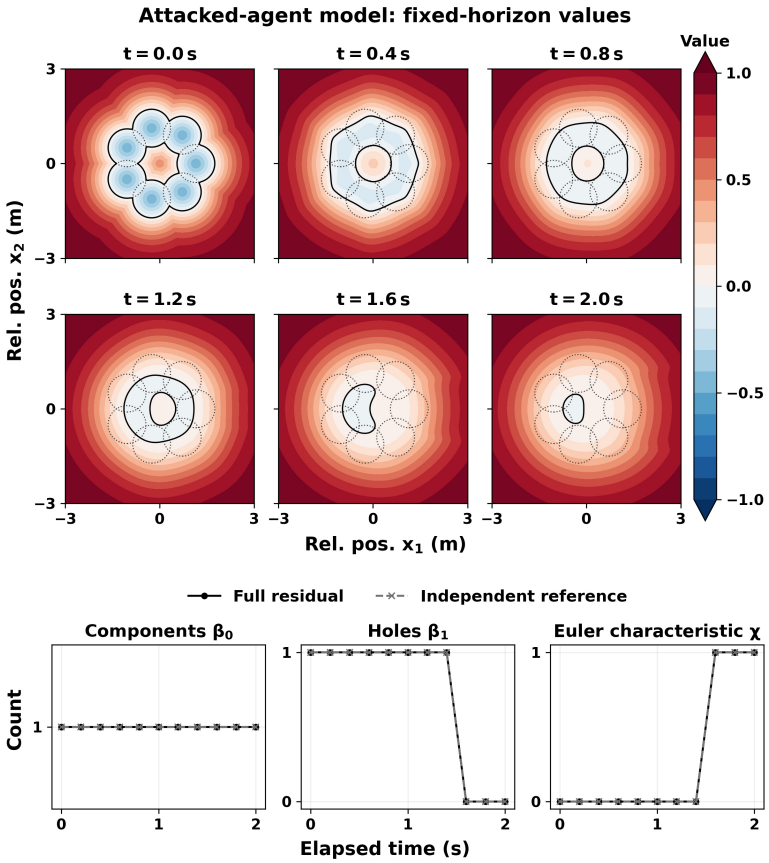


Figure 6: Numerical geometry of the flock value approximation. The upper panels show values at six elapsed horizons in seconds, with black zero contours and dotted initial target disks. Position axes are in meters under the reference scales in Section 3(c). The lower panel displays the component count, hole count, and Euler characteristic used to describe the slice topology. These are geometric diagnostics rather than a population-wide sign certificate.

on $[-\pi,\pi)^d$, for $d=2,3$, $\delta=0.3$, and horizon 0.4. The Hamiltonian is

$$\boldsymbol{H}(t,\boldsymbol{x},\boldsymbol{p})=\sum_{j=1}^{d}a_jp_j+0.1\big(|\boldsymbol{p}|^2-|D\boldsymbol{v}_{\mathrm{ref}}(t,\boldsymbol{x})|^2\big). \tag{C.22}$$

The prescribed coefficient is (3.2), with $L_j=2\pi$. Its time and spatial derivatives are nonzero. Grid counts increase through 16, 24, 32, and 48 per coordinate, while the step counts increase through 40, 80, 160, and 320. Separate fixed-grid time-halving tests check second-order temporal refinement of the exponential integrator.

On the finest Fourier grid, the maximum full residual value errors are 1.07×10^{-8} in two dimensions and 3.00×10^{-8} in three dimensions. The independent finite-difference errors are 1.68×10^{-6} and 2.08×10^{-6}, respectively. These deterministic comparisons check the implemented identity and discretizations against a known solution; they do not establish uniform error bounds for the game Hamiltonians.

At the finest grid, bird-model errors are 3.84×10^{-6} (2D) and 3.91×10^{-6} (3D) for FD4, and 6.14×10^{-9} (2D) and 7.53×10^{-9} (3D) for FD6, independently checking the transformation identity.

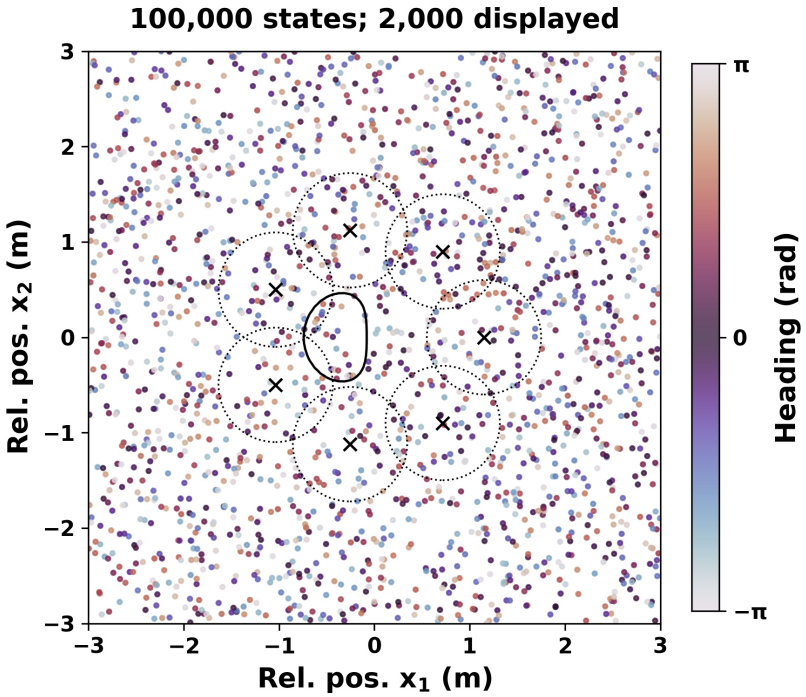


Figure 7: Planar projection of 2,000 displayed bird-model states, colored by heading. Crosses and dotted circles mark target centers and initial disks; positions are in meters. The $\theta = 0$ zero contour supplies context, while classification uses each state's heading and requires an error margin.

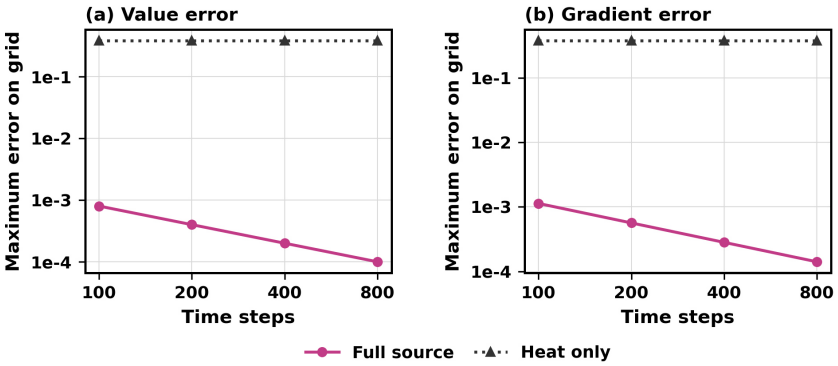


Figure 8: Temporal refinement at $t = 0.5$ on 128 periodic states. The panels show maximum value and gradient errors on logarithmic axes; circles denote the full source and triangles homogeneous heat flow. The advected damped sine is the deterministic reference; no Monte Carlo averaging enters the comparison.