

Risk-Minimizing Two-Player Zero-Sum Stochastic Differential Game via Path Integral Control

Apurva Patil¹

Yujing Zhou²

David Fridovich-Keil³

Takashi Tanaka³

Abstract—This paper addresses a continuous-time risk-minimizing two-player zero-sum stochastic differential game (SDG), in which each player aims to minimize its probability of failure. Failure occurs in the event when the state of the game enters into predefined undesirable domains, and one player’s failure is the other’s success. We derive a sufficient condition for this game to have a saddle-point equilibrium and show that it can be solved via a Hamilton-Jacobi-Isaacs (HJI) partial differential equation (PDE) with a Dirichlet boundary condition. Under certain assumptions on the system dynamics and cost function, we establish the existence and uniqueness of the saddle-point of the game. We provide explicit expressions for the saddle-point policies which can be numerically evaluated using path integral control. This allows us to solve the game online via Monte Carlo sampling of system trajectories. We implement our control synthesis framework on two classes of risk-minimizing zero-sum SDGs: a disturbance attenuation problem and a pursuit-evasion game. Simulation studies are presented to validate the proposed control synthesis framework.

I. INTRODUCTION

Interactions among multiple agents are prevalent in many fields such as economics, politics, and engineering. Game theory studies the collective decision-making process of multiple interacting agents [1]. Two-person zero-sum games involve two players with conflicting interests, and one player’s gain is the other’s loss. Pursuit-evasion games (competition between a pursuer and an evader) [2] and robust control (competition between a controller and the nature) [1] are some examples of two-player zero-sum games. In this paper, we consider a two-player stochastic differential game (SDG) in which the outcome of the game depends not only on the decision of both players but also on the stochastic input added by nature.

When the game dynamics and cost functions are known, the saddle-point equilibrium of a two-player zero-sum SDG can be characterized by the Hamilton-Jacobi-Isaacs (HJI) partial differential equation (PDE). The analytical solutions of the HJI PDEs are in general not available, and one needs to resort to numerical methods such as grid-based approaches [3], [4] to solve these PDEs approximately. However, the grid-based approaches suffer from *curse of dimensionality*, making them computationally intractable for systems with large dimensions [5]. Moreover, in general, the solutions

can not be computed in real-time using these methods; they need to be precomputed as lookup tables and recalled for use in an online setting [4]. Several reinforcement learning algorithms have also been proposed to find approximate solutions to game problems. A reinforcement-learning-based adaptive dynamic programming algorithm is proposed in [6] to determine online a saddle-point solution of linear continuous-time two-player zero-sum differential games. A deep reinforcement learning algorithm based on updating players’ policies simultaneously is proposed in [7] to solve two-player zero-sum games. These methods assume deterministic game dynamics and do not consider system uncertainties. In the presence of system uncertainties, the performance and safety of both players are affected unpredictably, if the uncertainties are not accommodated while designing policies. An effective uncertainty evaluation method, the multivariate probabilistic collocation, was used in [8] with integral reinforcement learning to solve multi-player SDGs for linear system dynamics online. A two-person zero-sum stochastic game with discrete states and actions is solved in [9] using Bayesian inverse reinforcement learning. Common challenges in the learning-based methods include training efficiency and rigorous theoretical guarantees on convergence and optimality. Moreover, these approaches do not explicitly take into account the players’ failure probabilities while synthesizing their policies.

In our work, we formulate a continuous-time, nonlinear, two-player zero-sum SDG on a state space modeled by an Itô stochastic differential equation. Since the stochastic uncertainties in our model are unbounded, both players have nonzero probabilities of failure. Failure occurs when the state of the game enters into predefined undesirable domains, and one player’s failure is the other’s success. Our objective is to solve a game in which each player seeks to minimize its risk of failure (failure probability)¹ along with its control cost; hence, the name *risk-minimizing zero-sum SDG*. We explain the risk-minimizing zero-sum SDG via the following example:

Example 1: Consider a pursuit-evasion game in which the pursuer catches the evader if they are less than a certain distance ρ away from each other. In this setting, the evader wishes to minimize its probability of entering the ball of radius ρ centered at the pursuer’s location. Whereas, the pursuer wishes to minimize the probability of staying out of the ball of radius ρ centered at the evader’s location. The

¹Throughout this paper the word “risk” simply means the probability of failure. It is not our intention to discuss various risk measures existing in the literature (e.g. [10], [11]).

This work was supported by AFOSR Grant FA9550-20-1-0101.

¹Walker Department of Mechanical Engineering, University of Texas at Austin, apurvapatil@utexas.edu. ²Department of Mechanical and Aerospace Engineering, Princeton University, yz1324@princeton.edu. ³Department of Aerospace Engineering and Engineering Mechanics, University of Texas at Austin, dfk@utexas.edu, ttanaka@utexas.edu.

goal of each player is to balance the trade-off between the above probabilities (probabilities of failure) and the control cost (e.g., their energy consumption). This problem can be formulated as a risk-minimizing zero-sum SDG.

We derive a sufficient condition for this game to have a saddle-point equilibrium and show that it can be solved via an HJI PDE with *Dirichlet boundary condition*. Under certain assumptions on the system dynamics and cost function, we establish the existence and uniqueness of the saddle-point equilibrium of the formulated risk-minimizing zero-sum SDG. Furthermore, explicit expressions for the saddle-point policies are derived which can be numerically evaluated using path integral control. The idea behind the path integral control is to use the Feynman-Kac lemma [12] and solve a linear PDE via Monte Carlo samples of system trajectories [13], [14]. The Monte Carlo simulations can be massively parallelized through the use of graphics processing units (GPUs); hence this approach is less susceptible to the curse of dimensionality. The use of the path integral technique to solve stochastic games was proposed in [15]. In this paper, we generalize their work and develop a path integral formulation to solve HJI PDEs with Dirichlet boundary conditions and find saddle-point equilibria of risk-minimizing zero-sum SDGs. The proposed framework allows us to solve the game online using Monte Carlo simulations of system trajectories, without the need for any offline training or precomputations.

The contributions of this work are as follows: 1) We formulate a continuous-time risk-minimizing zero-sum SDG in which players aim at balancing the trade-off between the failure probability and control cost. A sufficient condition for this game to have a saddle-point equilibrium is derived, and it is shown that this game can be solved via an HJI PDE with a Dirichlet boundary condition. 2) Under certain assumptions on the system dynamics and cost function, we establish the existence and uniqueness of the saddle-point solution. We also obtain explicit expressions for the saddle-point policies which can be numerically evaluated using path integral control. 3) The proposed control synthesis framework is validated by applying it to two classes of risk-minimizing zero-sum SDGs, namely a disturbance attenuation problem and a pursuit-evasion game.

Notation

Bold symbols such as \mathbf{x} represent random variables. If a stochastic process $\mathbf{x}(s)$, $s \geq t$ starts from x at time t , then let $P_{x,t}(\mathcal{E})$ denote the probability of event \mathcal{E} conditioned on $\mathbf{x}(t) = x$, and let $\mathbb{E}_{x,t}[F(\mathbf{x})]$ denote the expectation of a functional $F(\mathbf{x})$ conditioned on $\mathbf{x}(t) = x$. Let $\mathbf{1}_{\mathcal{E}}$ be an indicator function, that returns 1 when the condition \mathcal{E} holds and 0 otherwise. $\text{Tr}(A)$ denotes the trace of a matrix A .

II. PROBLEM FORMULATION

We consider a two-player zero-sum stochastic differential game (SDG) on a finite time horizon $t \in [t_0, T]$, $t_0 < T$. Consider a class of control-affine stochastic systems described by the following Itô stochastic differential equation (SDE):

$$d\mathbf{x}(t) = f(\mathbf{x}(t), t) dt + G_u(\mathbf{x}(t), t) u(\mathbf{x}(t), t) dt + G_v(\mathbf{x}(t), t) v(\mathbf{x}(t), t) dt + \Sigma(\mathbf{x}(t), t) d\mathbf{w}(t) \quad (1)$$

where $\mathbf{x}(t) \in \mathbb{R}^n$ is the state, $u(\mathbf{x}(t), t) \in \mathbb{R}^m$ is the control input of the first player (henceforth called the agent), and $v(\mathbf{x}(t), t) \in \mathbb{R}^l$ that of the second player (called the adversary). $\mathbf{w}(t) \in \mathbb{R}^k$ is a k -dimensional standard Wiener process on a suitable probability space (Ω, \mathcal{F}, P) . We assume sufficient regularity in the functions $f(\mathbf{x}(t), t) \in \mathbb{R}^n$, $G_u(\mathbf{x}(t), t) \in \mathbb{R}^{n \times m}$, $G_v(\mathbf{x}(t), t) \in \mathbb{R}^{n \times l}$ and $\Sigma(\mathbf{x}(t), t) \in \mathbb{R}^{n \times k}$ so that a unique strong solution of (1) exists [12]. Both the control inputs u , and v are assumed to be square integrable (i.e., of finite energy). In the rest of the paper, for notational compactness, the functional dependencies on x and t are dropped whenever it is unambiguous.

Let $\mathcal{X}_s \subseteq \mathbb{R}^n$ be a bounded open set representing a safe region, $\partial\mathcal{X}_s$ be its boundary, and closure $\overline{\mathcal{X}_s} = \mathcal{X}_s \cup \partial\mathcal{X}_s$. Suppose that the agent tries to keep the system (1) in the safe set \mathcal{X}_s for the entire time horizon $[t_0, T]$ of the game, whereas the adversary seeks the opposite. For example, in pursuit-evasion games, the safe set \mathcal{X}_s could be a region outside the ball of radius ρ , centered at the adversary's location. Or, in the disturbance rejection problems, if an agent wishes to navigate through obstacles in the presence of adversarial disturbances, then the region outside obstacles could be considered as a safe set. Suppose, when the game starts at t_0 , the system is in the safe set i.e., $\mathbf{x}(t_0) = x_0 \in \mathcal{X}_s$. If the system leaves the region \mathcal{X}_s at any time $t \in (t_0, T]$, we say that the agent fails. On the other hand, the adversary fails if the system stays in \mathcal{X}_s for all $t \in [t_0, T]$. Therefore, we define the agent's probability of failure $P_{\text{fail}}^{\text{ag}}$ as

$$P_{\text{fail}}^{\text{ag}} := P_{x_0, t_0} \left(\bigvee_{t \in (t_0, T]} \mathbf{x}(t) \notin \mathcal{X}_s \right) \quad (2)$$

and the adversary's probability of failure $P_{\text{fail}}^{\text{ad}} := 1 - P_{\text{fail}}^{\text{ag}}$. We define the terminal time \mathbf{t}_f of the game as

$$\mathbf{t}_f := \begin{cases} T, & \text{if } \mathbf{x}(t) \in \mathcal{X}_s, \forall t \in (t_0, T), \\ \inf \{t \in (t_0, T) : \mathbf{x}(t) \notin \mathcal{X}_s\}, & \text{otherwise.} \end{cases} \quad (3)$$

Alternatively, \mathbf{t}_f can be defined as

$$\mathbf{t}_f := \inf\{t > t_0 : (\mathbf{x}(t), t) \notin \mathcal{Q}\} \quad (4)$$

where $\mathcal{Q} = \mathcal{X}_s \times [t_0, T]$ is a bounded set with the boundary $\partial\mathcal{Q} = (\partial\mathcal{X}_s \times [t_0, T]) \cup (\mathcal{X}_s \times \{T\})$, and closure $\overline{\mathcal{Q}} = \mathcal{Q} \cup \partial\mathcal{Q} = \overline{\mathcal{X}_s} \times [t_0, T]$. Note that by the above definitions, $(\mathbf{x}(\mathbf{t}_f), \mathbf{t}_f) \in \partial\mathcal{Q}$ and the agent's failure probability $P_{\text{fail}}^{\text{ag}}$ in (2) can be written in terms of \mathbf{t}_f as

$$P_{x_0, t_0} \left(\bigvee_{t \in (t_0, T]} \mathbf{x}(t) \notin \mathcal{X}_s \right) = \mathbb{E}_{x_0, t_0} [\mathbf{1}_{\mathbf{x}(\mathbf{t}_f) \in \partial\mathcal{X}_s}]. \quad (5)$$

Since the stochastic uncertainty of the system (1) is modeled with an unbounded distribution, both the agent and the adversary have nonzero probabilities of failure. In our two-player

SDG setting, we assume that both players aim to design an optimal policy against the *worst* possible opponent's policy such that their own *risk of failure* is minimized. Therefore, we define the following *risk-minimizing* cost function:

$$C(x_0, t_0; u, v) := \eta \mathbb{E}_{x_0, t_0} [\mathbb{1}_{\mathbf{x}(\mathbf{t}_f) \in \partial \mathcal{X}_s}] + \mathbb{E}_{x_0, t_0} \left[\psi(\mathbf{x}(\mathbf{t}_f)) \cdot \mathbb{1}_{\mathbf{x}(\mathbf{t}_f) \in \mathcal{X}_s} + \int_{t_0}^{\mathbf{t}_f} L(\mathbf{x}(t), \mathbf{u}(t), \mathbf{v}(t), t) dt \right]. \quad (6)$$

The first term indicates the penalty associated with the agent's failure with the weight parameter $\eta > 0$. $\psi(\mathbf{x}(\mathbf{t}_f))$ and $L(\mathbf{x}(t), \mathbf{u}(t), \mathbf{v}(t), t)$ denote the terminal and running costs, respectively. Note that the game ends at \mathbf{t}_f and the system doesn't evolve after that (this is motivated by the applications where "collision" or "capture" ends the game). Therefore, the running cost is integrated over the time horizon $[t_0, \mathbf{t}_f]$. The agent tries to minimize C by controlling u , whereas the adversary tries to maximize it by controlling v . The weight parameter η balances the trade-off between the control cost and the failure probability.

Notice that if we define $\phi : \bar{\mathcal{X}}_s \rightarrow \mathbb{R}$ as:

$$\phi(x) := \psi(x) \cdot \mathbb{1}_{x \in \mathcal{X}_s} + \eta \cdot \mathbb{1}_{x \in \partial \mathcal{X}_s}, \quad (7)$$

then, the first term in (6) can be absorbed in a new terminal cost function ϕ as follows:

$$C(x_0, t_0; u, v) = \mathbb{E}_{x_0, t_0} \left[\phi(\mathbf{x}(\mathbf{t}_f)) + \int_{t_0}^{\mathbf{t}_f} L(\mathbf{x}, \mathbf{u}, \mathbf{v}, t) dt \right]. \quad (8)$$

In this paper, we consider the following running cost that is quadratic in u and v :

$$L(\mathbf{x}, \mathbf{u}, \mathbf{v}, t) = V(\mathbf{x}, t) + \frac{1}{2} \mathbf{u}^T R_u(\mathbf{x}, t) \mathbf{u} - \frac{1}{2} \mathbf{v}^T R_v(\mathbf{x}, t) \mathbf{v} \quad (9)$$

where $V(\mathbf{x}, t)$ denotes a state dependent cost, and $R_u(\mathbf{x}, t) \in \mathbb{R}^{m \times m}$ and $R_v(\mathbf{x}, t) \in \mathbb{R}^{l \times l}$ are given positive definite matrices (for all values of \mathbf{x} and t). Now, we formulate our risk-minimizing zero-sum SDG as follows:

Problem 1 (Risk-Minimizing Zero-Sum SDG):

$$\begin{aligned} & \min_u \max_v \mathbb{E}_{x_0, t_0} \left[\phi(\mathbf{x}(\mathbf{t}_f)) + \int_{t_0}^{\mathbf{t}_f} \left(\frac{1}{2} \mathbf{u}^T R_u \mathbf{u} - \frac{1}{2} \mathbf{v}^T R_v \mathbf{v} + V \right) dt \right] \\ & \text{s.t. } d\mathbf{x} = f dt + G_u u dt + G_v v dt + \Sigma d\mathbf{w}, \\ & \mathbf{x}(t_0) = x_0. \end{aligned} \quad (10)$$

where the admissible policies u, v are measurable with respect to the σ -algebra generated by $\mathbf{x}(s), t_0 \leq s \leq t$.

Note that Problem 1 is a *variable-terminal-time* zero-sum SDG where the terminal time is determined by (4).

III. SYNTHESIS OF MINIMAX POLICIES

In Section III-A, we show that Problem 1 can be solved via an HJI PDE with appropriate Dirichlet boundary condition. In Section III-B, we find a solution of a class of risk-minimizing zero-sum SDGs via path integral control.

A. HJI PDE with Dirichlet Boundary Condition

Problem 1 can be solved by utilizing the principle of dynamic programming. For each $(x, t) \in \bar{\mathcal{Q}}$, and admissible policies u, v over $[t, T]$, define the cost-to-go function:

$$C(x, t; u, v) = \mathbb{E}_{x, t} [\phi(\mathbf{x}(\mathbf{t}_f))] + \mathbb{E}_{x, t} \left[\int_t^{\mathbf{t}_f} \left(\frac{1}{2} \mathbf{u}^T R_u \mathbf{u} - \frac{1}{2} \mathbf{v}^T R_v \mathbf{v} + V \right) dt \right]. \quad (11)$$

Definition 1 (Saddle-point solution): [16, Chapter 2]: Given a two-player zero-sum differential game, a pair of admissible policies (u^*, v^*) over $[t, T]$ constitutes a saddle-point solution, if for each $(x, t) \in \bar{\mathcal{Q}}$, and admissible policies (u, v) over $[t, T]$,

$$C(x, t; u^*, v) \leq C^* := C(x, t; u^*, v^*) \leq C(x, t; u, v^*).$$

The quantity C^* is the *value* of the game. The value of the game is defined if it satisfies the following relation

$$C^* = \min_u \max_v C(x, t; u, v) = \max_v \min_u C(x, t; u, v).$$

The following theorem provides the sufficient condition for a saddle-point solution of Problem 1 to exist.

Theorem 1: Suppose there exists a function $J : \bar{\mathcal{Q}} \rightarrow \mathbb{R}$ such that

- (a) $J(x, t)$ is continuously differentiable in t and twice continuously differentiable in x in the domain \mathcal{Q} ;
- (b) $J(x, t)$ solves the following stochastic HJI PDE:

$$\begin{cases} -\partial_t J = V + f^T \partial_x J + \frac{1}{2} \text{Tr}(\Sigma \Sigma^T \partial_x^2 J) & \forall (x, t) \in \mathcal{Q}, \\ + \frac{1}{2} (\partial_x J)^T (G_v R_v^{-1} G_v^T - G_u R_u^{-1} G_u^T) \partial_x J, \\ \lim_{\substack{(x, t) \rightarrow (y, s) \\ (x, t) \in \mathcal{Q}}} J(x, t) = \phi(y), & \forall (y, s) \in \partial \mathcal{Q}. \end{cases} \quad (12)$$

Then, the following statements hold:

- (i) $J(x, t)$ is the value of the game formulated in Problem 1. That is,

$$\begin{aligned} J(x, t) &= \min_u \max_v C(x, t; u, v) \\ &= \max_v \min_u C(x, t; u, v), \quad \forall (x, t) \in \bar{\mathcal{Q}}. \end{aligned} \quad (13)$$

- (ii) The optimal solution to Problem 1 is given by

$$u^*(x, t) = -R_u^{-1}(x, t) G_u^T(x, t) \partial_x J(x, t), \quad (14)$$

$$v^*(x, t) = R_v^{-1}(x, t) G_v^T(x, t) \partial_x J(x, t). \quad (15)$$

Proof: Refer to the supplementary material [17, Appendix A] ■

Remark 1: Theorem 1 does not say anything about the existence of a function $J(x, t)$ satisfying statements (a) and (b), and it is not in the scope of this paper. However, in Section III-B, we focus on a special case in which (12) can be linearized where the existence and uniqueness of such a function are guaranteed.

B. Path Integral Formulation

In this section, we derive a path integral formulation to solve a class of risk-minimizing zero-sum SDGs that satisfy certain assumptions on the system dynamics and cost function. Let $\xi(x, t)$ be the logarithmic transformation (known as Cole-Hopf transformation in the PDE literature) of the value function $J(x, t)$ defined as

$$J(x, t) = -\lambda \log(\xi(x, t)) \quad (16)$$

where λ is a proportionality constant to be defined. Applying the transformation in (16) to (12) yields

$$\begin{cases} \partial_t \xi = \frac{V\xi}{\lambda} - \frac{1}{2} \text{Tr}(\Sigma \Sigma^\top \partial_x^2 \xi) + \frac{1}{2\xi} (\partial_x \xi)^\top \Sigma \Sigma^\top \partial_x \xi \\ + \frac{\lambda}{2\xi} (\partial_x \xi)^\top (G_v R_v^{-1} G_v^\top - G_u R_u^{-1} G_u^\top) \partial_x \xi - f^\top \partial_x \xi, \\ \forall (x, t) \in \mathcal{Q}, \\ \lim_{\substack{(x,t) \rightarrow (y,s) \\ (x,t) \in \mathcal{Q}}} \xi(x, t) = \exp\left(-\frac{\phi(y)}{\lambda}\right), \quad \forall (y, s) \in \partial \mathcal{Q}. \end{cases} \quad (17)$$

Now, we make the following assumption:

Assumption 1: For all $(x, t) \in \overline{\mathcal{Q}}$, there exists a constant $\lambda > 0$ such that

$$\begin{aligned} \Sigma(x, t) \Sigma^\top(x, t) &= \lambda G_u(x, t) R_u^{-1}(x, t) G_u^\top(x, t) \\ &\quad - \lambda G_v(x, t) R_v^{-1}(x, t) G_v^\top(x, t). \end{aligned} \quad (18)$$

Assumption 1 is similar to the assumption required in the path integral formulation of a single agent stochastic control problem [18]. A possible interpretation of condition (18) is that in a direction with high noise variance, the agent's control cost has to be low whereas that of the adversary has to be high. Therefore, the weights of the control cost R_u and R_v need to be tuned appropriately for the given diffusion coefficient $\Sigma(x, t)$ and the control gains $G_u(x, t)$ and $G_v(x, t)$ in the system dynamics (1). See [13], [14] for further discussion on a similar condition in the single agent setting. Assumption 1 also implies that the stochastic noise has to enter the system dynamics via the control channels. Therefore, in what follows, we assume that system (1) can be partitioned into subsystems that are directly and non-directly driven by the noise as:

$$\begin{aligned} \begin{bmatrix} d\mathbf{x}^{(1)} \\ d\mathbf{x}^{(2)} \end{bmatrix} &= \begin{bmatrix} f^{(1)}(\mathbf{x}, t) \\ f^{(2)}(\mathbf{x}, t) \end{bmatrix} dt + \begin{bmatrix} \mathbf{0} \\ G_u^{(2)}(\mathbf{x}, t) \end{bmatrix} u(\mathbf{x}, t) dt \\ &\quad + \begin{bmatrix} \mathbf{0} \\ G_v^{(2)}(\mathbf{x}, t) \end{bmatrix} v(\mathbf{x}, t) dt + \begin{bmatrix} \mathbf{0} \\ \Sigma^{(2)}(\mathbf{x}, t) \end{bmatrix} d\mathbf{w} \end{aligned} \quad (19)$$

where $\mathbf{0}$ denotes a zero matrix of appropriate dimensions. By assuming a λ satisfying Assumption 1 holds in (17), we obtain the linear PDE in ξ with Dirichlet boundary condition:

$$\begin{cases} \partial_t \xi = \frac{V\xi}{\lambda} - f^\top \partial_x \xi - \frac{1}{2} \text{Tr}(\Sigma \Sigma^\top \partial_x^2 \xi), \quad \forall (x, t) \in \mathcal{Q}, \\ \lim_{\substack{(x,t) \rightarrow (y,s) \\ (x,t) \in \mathcal{Q}}} \xi(x, t) = \exp\left(-\frac{\phi(y)}{\lambda}\right), \quad \forall (y, s) \in \partial \mathcal{Q}. \end{cases} \quad (20)$$

The solution of a linear Dirichlet boundary value problem of the form (20) exists under a sufficiently regular boundary

condition, and it is unique [19, Chapter 6]. Furthermore, the solution admits the Feynman-Kac representation [20]. Suppose $\hat{\mathbf{x}}(t) \in \mathbb{R}^n$ is an uncontrolled process driven by the following SDE:

$$d\hat{\mathbf{x}}(t) = f(\hat{\mathbf{x}}(t), t) dt + \Sigma(\hat{\mathbf{x}}(t), t) d\mathbf{w}(t) \quad (21)$$

and let $\hat{t}_f := \inf\{t > t_0 : (\hat{\mathbf{x}}(t), t) \notin \mathcal{Q}\}$. Then, the solution of the PDE (20) is given as

$$\xi(x, t) = \mathbb{E}_{x,t} \left[\exp\left(-\frac{1}{\lambda} S(\tau)\right) \right] \quad (22)$$

where $S(\tau)$ denotes the cost-to-go of a trajectory τ of the uncontrolled system (21) starting at (x, t) :

$$S(\tau) = \phi(\hat{\mathbf{x}}(\hat{t}_f)) + \int_t^{\hat{t}_f} V(\hat{\mathbf{x}}(t), t) dt. \quad (23)$$

Equation (22) provides a path integral form for the exponentiated value function $\xi(x, t)$, which can be numerically evaluated using Monte Carlo sampling of trajectories generated by the uncontrolled SDE (21). We now obtain the expressions for the saddle-point policies via the following theorem:

Theorem 2: Suppose Assumption 1 holds and the system (1) can be partitioned as (19). Then, a saddle-point solution of the risk-minimizing zero-sum SDG (10) exists, is unique, and is given by

$$u^*(x, t) dt = \mathcal{G}_u(x, t) \frac{\mathbb{E}_{x,t}[\exp(-\frac{1}{\lambda} S(\tau)) \Sigma^{(2)}(x, t) d\mathbf{w}]}{\mathbb{E}_{x,t}[\exp(-\frac{1}{\lambda} S(\tau))]}, \quad (24)$$

where

$$\mathcal{G}_u = R_u^{-1} G_u^{(2)\top} \left(G_u^{(2)} R_u^{-1} G_u^{(2)\top} - G_v^{(2)} R_v^{-1} G_v^{(2)\top} \right)^{-1}$$

$$v^*(x, t) dt = \mathcal{G}_v(x, t) \frac{\mathbb{E}_{x,t}[\exp(-\frac{1}{\lambda} S(\tau)) \Sigma^{(2)}(x, t) d\mathbf{w}]}{\mathbb{E}_{x,t}[\exp(-\frac{1}{\lambda} S(\tau))]}, \quad (25)$$

where

$$\mathcal{G}_v = -R_v^{-1} G_v^{(2)\top} \left(G_u^{(2)} R_u^{-1} G_u^{(2)\top} - G_v^{(2)} R_v^{-1} G_v^{(2)\top} \right)^{-1}.$$

Proof: The existence and uniqueness of the saddle-point solution follow from the existence and uniqueness of the linear Dirichlet boundary value problem (20) [19, Chapter 6] and from Theorem 1. The saddle-point solution $u^*(x, t)$ (14) and $v^*(x, t)$ (15) can be computed by taking the gradient of (22) with respect to x and using the condition (18). (The derivation of (14) and (15) is in the same vein as the derivation of optimal controls in single agent settings [18], [21]; not presented here for brevity.) ■

Equations (24) and (25) provide the path integral forms for the saddle-point equilibrium. Similar to (22), the expectations in (24) and (25) can be numerically evaluated in real-time via the Monte Carlo sampling of the trajectories generated by the uncontrolled SDE (21). The path integral framework evaluates the solution locally without requiring knowledge of the solution nearby so that there is no need for a (global) discretization of the computational domain. This allows us to solve the game online without requiring any offline training

or precomputations. Even though Monte Carlo simulations must be performed in real-time in order to evaluate (24, 25) for the current (x, t) , these simulations can be massively parallelized through the use of GPUs.

IV. EXAMPLES

In this section, we apply the path integral framework to two classes of risk-minimizing zero-sum SDGs (10): a disturbance attenuation problem and a pursuit-evasion game.

A. Disturbance Attenuation Problem

Consider a special class of systems (1):

$$d\mathbf{x} = f(\mathbf{x}, t)dt + G_u(\mathbf{x}, t)(u(\mathbf{x}, t)dt + v(\mathbf{x}, t)dt + d\mathbf{w}) \quad (26)$$

where $u(\mathbf{x}, t) \in \mathbb{R}^m$ is the control input, $v(\mathbf{x}, t) \in \mathbb{R}^m$ is the bounded disturbance and $\mathbf{w}(t) \in \mathbb{R}^m$ is a Wiener process. Here, we have two sources of noise that corrupt the system's control input u : the bounded noise v whose statistics are unknown, and the white noise $d\mathbf{w}$. In the disturbance attenuation problem, the objective is to design a policy u in the presence of stochastic noise and bounded disturbance v such that the system's control performance $\mathbb{E}_{x_0, t_0} \left[\phi(\mathbf{x}(t_f)) + \int_{t_0}^{t_f} \left(\frac{1}{2} \mathbf{u}^\top \mathbf{u} + V \right) dt \right]$ is minimized. This problem can be solved using the following zero-sum SDG, where u is considered as a control input of the first player (agent) and v that of the second player (adversary):

$$\min_u \max_v \mathbb{E}_{x_0, t_0} \left[\phi(\mathbf{x}(t_f)) + \int_{t_0}^{t_f} \left(\frac{1}{2} \mathbf{u}^\top \mathbf{u} - \frac{\gamma^2}{2} \mathbf{v}^\top \mathbf{v} + V \right) dt \right]. \quad (27)$$

γ is a given positive constant that determines the level of disturbance attenuation. Theorem 3 provides an upper bound on the system's control performance (in the presence of a bounded disturbance v) that can be obtained by solving the game (27).

Theorem 3: Suppose (u_γ^*, v_γ^*) represent the saddle-point policies of the SDG (27) for any γ , and let

$$\delta_\gamma := \mathbb{E}_{x_0, t_0}^{u_\gamma^*, v_\gamma^*} \left[\int_{t_0}^{t_f} \mathbf{v}^\top \mathbf{v} dt \right] \quad (28)$$

where the superscript on \mathbb{E} denotes the policies under which the expectation is computed. Then, for all adversarial policies v such that $\mathbb{E}_{x_0, t_0}^{u_\gamma^*, v} \left[\int_{t_0}^{t_f} \mathbf{v}^\top \mathbf{v} dt \right] \leq \delta$ (for any $\delta > 0$), we get the following upper bound on the system's control performance in the presence of disturbance v :

$$\begin{aligned} & \mathbb{E}_{x_0, t_0}^{u_\gamma^*, v_\gamma^*} \left[\phi(\mathbf{x}(t_f)) + \int_{t_0}^{t_f} \left(\frac{1}{2} \mathbf{u}^\top \mathbf{u} + V \right) dt \right] + \frac{\gamma^2}{2} (\delta - \delta_\gamma) \\ & \geq \mathbb{E}_{x_0, t_0}^{u_\gamma^*, v} \left[\phi(\mathbf{x}(t_f)) + \int_{t_0}^{t_f} \left(\frac{1}{2} \mathbf{u}^\top \mathbf{u} + V \right) dt \right]. \end{aligned} \quad (29)$$

Proof: Please refer to the supplementary material [17, Appendix B] ■

In order to solve the HJI PDE associated with the game (27) via the path integral framework described in Section III-B, it is necessary to find a constant $\lambda > 0$ (by Assumption 1) such that $\lambda \left(1 - \frac{1}{\gamma^2} \right) = 1$. Therefore, for all $\gamma > 1$, Assumption

1 is satisfied and as a consequence, the zero-sum SDG (27) admits a unique saddle-point solution.

We now present a simulation study of the disturbance attenuation problem using a unicycle navigation example. Consider the following unicycle dynamics model:

$$\begin{aligned} \begin{bmatrix} d\mathbf{p}_x \\ d\mathbf{p}_y \\ ds \\ d\theta \end{bmatrix} &= -k \begin{bmatrix} \mathbf{p}_x \\ \mathbf{p}_y \\ \mathbf{s} \\ \theta \end{bmatrix} dt + \begin{bmatrix} \mathbf{s} \cos \theta \\ \mathbf{s} \sin \theta \\ 0 \\ 0 \end{bmatrix} dt \\ &+ \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix} \left(\begin{bmatrix} a \\ \omega \end{bmatrix} dt + \begin{bmatrix} \Delta a \\ \Delta \omega \end{bmatrix} dt + \begin{bmatrix} \sigma & 0 \\ 0 & \nu \end{bmatrix} d\mathbf{w} \right), \end{aligned} \quad (30)$$

where $(\mathbf{p}_x, \mathbf{p}_y)$, \mathbf{s} and θ denote the position, speed, and the heading angle of the unicycle, respectively. The control input $u := [a \ \omega]^\top$ consists of acceleration a and angular speed ω . $v := [\Delta a \ \Delta \omega]^\top$ is the bounded disturbance acting on the system's control input, and $d\mathbf{w} \in \mathbb{R}^2$ is the white noise with σ and ν being the noise level parameters. As illustrated in Figure 1, the unicycle aims to navigate in a two-dimensional space from its initial position (represented by the yellow star) to the origin (represented by the magenta star), in finite time, while avoiding the red obstacles and the outer boundary. The white region that lies between the outer boundary and the obstacles is the safe region \mathcal{X}_s . This is a disturbance attenuation problem since the unicycle aims to design its control policy u in order to minimize the control performance and risk of failure (collision with the obstacles or the outer boundary) under worst-case disturbance v . Therefore, we can formulate this problem as the risk-minimizing zero-sum SDG (27). In the simulation, we set $\sigma = \nu = 0.1$, $k = 0.2$, $t_0 = 0$, $T = 10$, $x_0 = [-0.4 \ -0.4 \ 0 \ 0]^\top$, $V(\mathbf{x}) = \mathbf{p}_x^2 + \mathbf{p}_y^2$ and $\psi(\mathbf{x}(T)) = \mathbf{p}_x^2(T) + \mathbf{p}_y^2(T)$. In order to evaluate the optimal policies (24) and (25) via Monte Carlo sampling, 10^4 trajectories and a step size equal to 0.01 are used. We demonstrate two experiments.

1) *Experiment 1:* In this experiment, we set $\eta = 0.67$ and plot in Figure 1 100 sample trajectories generated using synthesized saddle-point policies (u^*, v^*) for two values of γ . The trajectories are color-coded; the blue paths collide with the obstacles, while the green paths converge in the neighborhood of the origin (the target position). The figure shows that for a higher value of γ i.e. when the adversary becomes less powerful, the failure probability of the agent $P_{\text{fail}}^{\text{ag}}$ reduces.

2) *Experiment 2:* In this experiment, we set $\gamma^2 = 3$, $\eta = 1$, and study the effect of ignoring the adversary. First, we compute saddle-point policies (u^*, v^*) for the game (27) same as Experiment 1 and plot in Figure 2-(a) 100 sample trajectories generated using (u^*, v^*) . In this case, the agent is aware of the adversary and designs its policy u^* cautiously. The probability of failure is 23%. In the second case, the agent is not aware of the presence of the adversary and computes its policy (say) \tilde{u}^* by solving a single agent

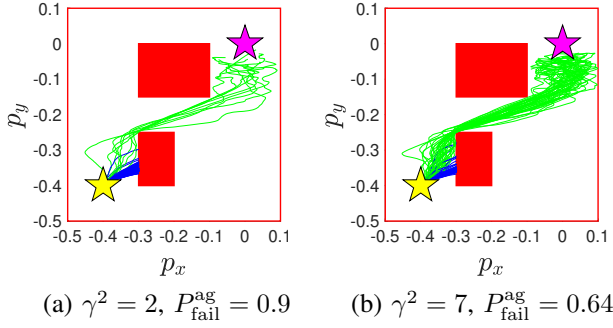


Fig. 1. Unicycle navigation in the presence of bounded and stochastic disturbances. The start position is shown by a yellow star and the target position (the origin) by a magenta star. 100 sample trajectories generated using saddle-point policies (u^*, v^*) for two values of γ are shown. The trajectories are color-coded; blue paths collide with the red obstacles or the outer boundary, while the green paths converge in the neighborhood of the magenta star. The failure probabilities of the agent $P_{\text{fail}}^{\text{ag}}$ are noted.

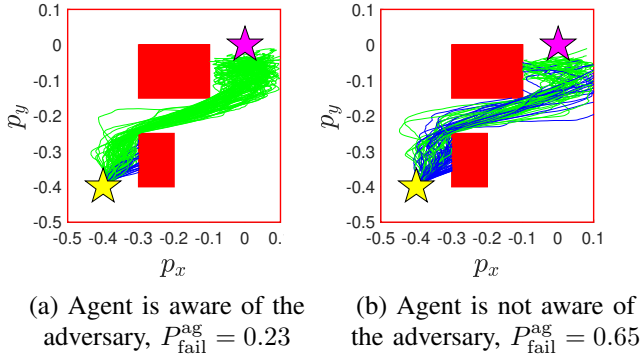


Fig. 2. Unicycle navigation in the presence of bounded and stochastic disturbances. (a) The agent is aware of the presence of an adversary. 100 sample trajectories generated using saddle-point policies (u^*, v^*) . (b) The agent is not aware of the presence of an adversary. 100 sample trajectories generated using (\tilde{u}^*, v^*) . The failure probabilities of the agent $P_{\text{fail}}^{\text{ag}}$ are noted for each case.

optimization problem. However, in reality, the adversary is present and suppose it follows the policy v^* . Figure 2-(b) shows 100 sample trajectories generated using (\tilde{u}^*, v^*) . In this case, the agent's performance is poor, it fails 65% of the time. The color-coding of the trajectories is the same as in Experiment 1.

B. Pursuit-Evasion Game

Consider a two-player zero-sum SDG on a finite time horizon $[t_0, T]$, in which the adversary is chasing the agent and the agent is trying to escape from the adversary. We will call the adversary as a pursuer and the agent as an evader. Suppose the evader and the pursuer are moving in a two-dimensional plane according to

$$\begin{aligned} dp_x^E &= u_x dt + \sigma_x^E d\mathbf{w}_x^E, & dp_x^P &= v_x dt + \sigma_x^P d\mathbf{w}_x^P, \\ dp_y^E &= u_y dt + \sigma_y^E d\mathbf{w}_y^E, & dp_y^P &= v_y dt + \sigma_y^P d\mathbf{w}_y^P, \end{aligned} \quad (31)$$

where $\mathbf{x}_E := [\mathbf{p}_x^E \ \mathbf{p}_y^E]^\top$ is the position and $u := [u_x \ u_y]^\top$ is the control input of the evader. Similarly, $\mathbf{x}_P := [\mathbf{p}_x^P \ \mathbf{p}_y^P]^\top$ and $v := [v_x \ v_y]^\top$ are the position

and control input of the pursuer. $\mathbf{w}_x^E, \mathbf{w}_y^E, \mathbf{w}_x^P, \mathbf{w}_y^P$ are independent one-dimensional standard Brownian motions. If at any time $t \in (t_0, T]$, the pursuer gets within a distance ρ of the evader, then it catches the evader and the evader fails. On the other hand, if the evader avoids getting within a distance ρ of the pursuer for the entire time horizon $[t_0, T]$, then that's a failure for the pursuer. The pursuer aims at designing its control policy v in order to maximize the probability of catching the evader, whereas the evader seeks the opposite by designing u . For this two-player differential game, it is the relative position of the pursuer and evader that is important (and relevant), rather than their absolute positions. Let $\mathbf{x} := [\mathbf{p}_x \ \mathbf{p}_y]$ be the evader's position with respect to the pursuer where $\mathbf{p}_x = \mathbf{p}_x^E - \mathbf{p}_x^P$, $\mathbf{p}_y = \mathbf{p}_y^E - \mathbf{p}_y^P$ and the origin coincides with the pursuer's position. Thus, the coordinate system is attached to the pursuer and is not fixed in space. The system \mathbf{x} follows the SDE

$$d\mathbf{x} = \begin{bmatrix} d\mathbf{p}_x \\ d\mathbf{p}_y \end{bmatrix} = \begin{bmatrix} u_x \\ u_y \end{bmatrix} dt - \begin{bmatrix} v_x \\ v_y \end{bmatrix} dt + \begin{bmatrix} \sigma_x & 0 \\ 0 & \sigma_y \end{bmatrix} d\mathbf{w} \quad (32)$$

where $\sigma_x = \sqrt{(\sigma_x^E)^2 + (\sigma_x^P)^2}$, $\sigma_y = \sqrt{(\sigma_y^E)^2 + (\sigma_y^P)^2}$ and \mathbf{w} is a two-dimensional standard Brownian motion. In this game, the safe set \mathcal{X}_s can be defined as $\mathcal{X}_s := \{x \in \mathbb{R}^2 : \|x\| > \rho\}$. Suppose the control cost matrix R_u of the evader is unity and that of the pursuer $R_v = r_v^2$, where r_v is a given positive scalar constant. Therefore, the risk-minimizing zero-sum SDG takes the form:

$$\min_u \max_v \mathbb{E}_{x_0, t_0} \left[\phi(\mathbf{x}(t_f)) + \int_{t_0}^{t_f} \left(\frac{1}{2} \mathbf{u}^T \mathbf{u} - \frac{r_v^2}{2} \mathbf{v}^T \mathbf{v} + V \right) dt \right]. \quad (33)$$

In order to solve the associated HJI equation of this game via the path integral framework, it is necessary to find a constant $\lambda > 0$ (by Assumption 1) such that $\lambda \left(1 - \frac{1}{r_v^2} \right) = 1$.

Therefore, for all $r_v > 1$, Assumption 1 is satisfied and as a consequence, the zero-sum SDG (33) admits a unique saddle-point solution.

In the simulation, we set $\sigma_x^E = \sigma_y^E = \sigma_x^P = \sigma_y^P = \sqrt{0.1}$, $\rho = 0.1$, $t_0 = 0$, $T = 2$, $x_0 = [0.3 \ 0.3]^\top$, $V(\mathbf{x}) = \psi(\mathbf{x}(T)) = 0$, $\eta = 0.2$, $r_v^2 = 2$. Figure 3 shows a plot of two sample trajectories of the system (32) generated using synthesized saddle-point policies (u^*, v^*) . The trajectories start from x_0 shown by the yellow star. The red disc of radius $\rho = 0.1$, centered at the origin represents that the pursuer is within a distance ρ of the evader. The green trajectory never enters the red disc in the horizon $[t_0, T]$, thus, it represents a case when the evader escapes from the pursuer. The blue trajectory, on the other hand, enters the red disc and thus represents a case when the pursuer catches the evader. Figure 4 shows a plot of failure probabilities of the agent (i.e., evader) as a function of r_v , when the players follow the saddle-point policies (u^*, v^*) . These values are computed using naïve Monte Carlo sampling, with 400 sample trajectories. The plot shows that as the control cost weight r_v of the adversary (i.e., pursuer) increases, the chances of the evader getting caught reduces.




Fig. 3. Two sample trajectories of the relative position of the players in a pursuit-evasion game. The start position of the trajectories is shown by a yellow star. The red disc of radius $\rho = 0.1$, centered at the origin represents that the pursuer is within the distance ρ of the evader. The green trajectory never enters the red disc in the horizon $[t_0, T]$, thus, it represents a case when the evader wins. The blue trajectory enters the red disc and thus represents a case when the pursuer wins.




Fig. 4. Failure probabilities of the agent (i.e., evader) as a function of r_v , when the players follow the saddle-point policies (u^*, v^*) .

The aim of the presented simulation studies is to validate the proposed theoretical formulation of the risk-minimizing zero-sum SDGs. Future work will emphasize scaling this framework to higher dimensional and more complex game dynamics.

V. CONCLUSION

The paper presented an HJI-PDE-based solution approach for a risk-minimizing two-player zero-sum stochastic differential game (SDG). This is a variable-end-time game in which each player tries to balance the trade-off between the probability of failure and the control cost. A sufficient condition for a saddle-point solution of this game was derived and it was shown that this game can be solved via an HJI PDE with the Dirichlet boundary condition. We developed a path integral framework to numerically solve a class of risk-minimizing zero-sum SDGs whose associated HJI PDE can be linearized and established the existence and uniqueness of the saddle-point solution. The presented approach allows the game to be solved online without the need for any offline training or precomputations. Application of our approach to two classes of risk-minimizing zero-sum SDGs: a disturbance attenuation problem and a pursuit-evasion game was presented and the framework was validated through simulation studies. In the future, we plan to conduct sample complexity analysis for path integral control in order to investigate how the accuracy of Monte Carlo sampling affects the solution of SDGs. The central challenge in using the path integral framework is the particular requirement on

the relationship between the cost function and the noise covariance. This requirement restricts the class of applicable system models and cost functions. In future work, we plan to find alternatives in order to get rid of this restrictive requirement (one such solution is provided in [18]). Another topic of future investigation could be chance-constrained stochastic games in which each player would aim to satisfy a hard bound on its failure probability.

REFERENCES

- [1] T. Başar and G. J. Olsder, *Dynamic noncooperative game theory*. SIAM, 1998.
- [2] P. J. Nahin, “Chases and escapes,” in *Chases and Escapes*. Princeton University Press, 2012.
- [3] M. Falcone, “Numerical methods for differential games based on partial differential equations,” *International Game Theory Review*, vol. 8, no. 02, pp. 231–272, 2006.
- [4] H. Huang, J. Ding, W. Zhang, and C. J. Tomlin, “Automation-assisted capture-the-flag: A differential game approach,” *IEEE Transactions on Control Systems Technology*, vol. 23, no. 3, pp. 1014–1028, 2014.
- [5] 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*, vol. 50, no. 7, pp. 947–957, 2005.
- [6] D. Vrabie and F. Lewis, “Adaptive dynamic programming for online solution of a zero-sum differential game,” *Journal of Control Theory and Applications*, vol. 9, no. 3, pp. 353–360, 2011.
- [7] M. Prajapat, K. Azizzadenesheli, A. Liniger, Y. Yue, and A. Anandkumar, “Competitive policy optimization,” in *Uncertainty in Artificial Intelligence*. PMLR, 2021, pp. 64–74.
- [8] M. Liu, Y. Wan, F. L. Lewis, and V. G. Lopez, “Adaptive optimal control for stochastic multiplayer differential games using on-policy and off-policy reinforcement learning,” *IEEE transactions on neural networks and learning systems*, vol. 31, no. 12, pp. 5522–5533, 2020.
- [9] X. Lin, P. A. Beling, and R. Cogill, “Multiagent inverse reinforcement learning for two-person zero-sum games,” *IEEE Transactions on Games*, vol. 10, no. 1, pp. 56–68, 2017.
- [10] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath, “Coherent measures of risk,” *Mathematical finance*, vol. 9, no. 3, pp. 203–228, 1999.
- [11] A. Dixit, M. Ahmadi, and J. W. Burdick, “Risk-averse receding horizon motion planning,” *arXiv preprint arXiv:2204.09596*, 2022.
- [12] B. Oksendal, *Stochastic differential equations: an introduction with applications*. Springer Science & Business Media, 2013.
- [13] H. J. Kappen, “Path integrals and symmetry breaking for optimal control theory,” *Journal of statistical mechanics: theory and experiment*, vol. 2005, no. 11, p. P11011, 2005.
- [14] G. Williams, A. Aldrich, and E. A. Theodorou, “Model predictive path integral control: From theory to parallel computation,” *Journal of Guidance, Control, and Dynamics*, vol. 40, no. 2, pp. 344–357, 2017.
- [15] D. Vrushabh, P. Akshay, K. Sonam, S. Wagh, and N. M. Singh, “Robust path integral control on stochastic differential games,” in *2020 28th Mediterranean Conference on Control and Automation (MED)*. IEEE, 2020, pp. 665–670.
- [16] T. Başar and P. Bernhard, *H-infinity optimal control and related minimax design problems: a dynamic game approach*. Springer Science & Business Media, 2008.
- [17] A. Patil, Y. Zhou, D. Fridovich-Keil, and T. Tanaka, “Risk-minimizing two-player zero-sum stochastic differential game via path integral control [with appendices],” *arXiv preprint arXiv:2308.11546*, 2023.
- [18] S. Satoh, H. J. Kappen, and M. Saeki, “An iterative method for nonlinear stochastic optimal control based on path integrals,” *IEEE Transactions on Automatic Control*, vol. 62, no. 1, pp. 262–276, 2016.
- [19] A. Friedman, *Stochastic differential equations and applications, vol. 1*. Academic Press, 1975.
- [20] A. Patil, A. Duarte, A. Smith, F. Bisetti, and T. Tanaka, “Chance-constrained stochastic optimal control via path integral and finite difference methods,” in *2022 IEEE 61st Conference on Decision and Control (CDC)*. IEEE, 2022, pp. 3598–3604.
- [21] E. Theodorou, J. Buchli, and S. Schaal, “A generalized path integral control approach to reinforcement learning,” *The Journal of Machine Learning Research*, vol. 11, pp. 3137–3181, 2010.