A Fokker-Planck control framework for multidimensional stochastic processes

Alfio Borzì

Institute for Mathematics, Universität Würzburg, Germany
Joint work with Mario Annunziato (U Salerno)

A multidimensional stochastic process

We consider continuous-time stochastic processes described by the following multidimensional model

$$\begin{cases} dX_t = b(X_t, t; u)\, dt + \sigma(X_t, t)\, dW_t \\ X_{t_0} = X_0, \end{cases}$$

where $X_t \in \mathbb{R}^n$ is the state variable and $dW_t \in \mathbb{R}^m$ is a multi-dimensional Wiener process, with stochastically independent components.

We consider the action of a time-dependent control $u(t) \in \mathbb{R}^\ell$ in the drift term $b(X_t, t; u)$ that allows to drive the vector random process.

The average objective

Since X_t is random, a deterministic objective will result into a random variable for which an averaging step is required. Therefore, the following objective is usually considered

$$J(X,u) = \mathbb{E}[\int_0^T L(t, X_t, u(t))\, dt + \Psi[X_T]].$$

With this formulation it is supposed that the controller knows (all) the state of the system at each instant of time!

The average $\mathbb{E}[\cdot]$ of functionals of X_t is omnipresent in almost all stochastic optimal control problems considered in the literature.

Alternative approaches with deterministic objective

The state of a stochastic process can be characterized by the shape of its statistical distribution represented by the probability density function (PDF).

In some works, control schemes were proposed, where the deterministic objective depends on the PDF of the stochastic state variable and no average is needed. Examples are objectives defined by the Kullback-Leibler distance or the square distance between the state PDF and a desired one.

Nevertheless, stochastic governing models are used and the state PDF is obtained by averaging or by an interpolation strategy.

The Fokker-Planck (-Kolmogorov) equation (step 1)

Consider a particle at x at time t. Let $\pi_{\Delta x}^{(+)}(x)$ and $\pi_{\Delta x}^{(-)}(x)$ be the probabilities that the particle will be at $x + \Delta x$ and $x - \Delta x$, at $t + \Delta t$.

Let $p(x_0, x; t)\Delta x$ be the conditional probability that the particle arrives at x at time t starting from x_0 at $t = 0$ following a random path. We have

$$
\begin{aligned}
p(x_0, x; t)\Delta x &= p(x_0, x - \Delta x; t - \Delta t)\pi_{\Delta x}^{(+)}(x - \Delta x)\Delta x \\
&+ p(x_0, x + \Delta x; t - \Delta t)\pi_{\Delta x}^{(-)}(x + \Delta x)\Delta x \\
&+ p(x_0, x; t - \Delta t)(1 - \pi_{\Delta x}^{(+)}(x) - \pi_{\Delta x}^{(-)}(x))\Delta x.
\end{aligned}
$$

From this discrete model of a stochastic process, we build one with infinitesimal increments for $\Delta x, \Delta t \to 0$.

For a meaningful statistical limiting process, the probabilities $\pi_{\Delta x}^{(+)}$ and $\pi_{\Delta x}^{(-)}$ must be subject to some constraints.

The Fokker-Planck (-Kolmogorov) equation (step 2)

Consider the mean of change of particle position $X(t)$, conditional on $X(t) = x$,

$$\beta(x) = \lim_{\Delta t \to 0} \frac{E[X(t+\Delta t) - X(t) | X(t) = x]}{\Delta t}$$

and the corresponding variance is given by

$$\alpha(x) = \lim_{\Delta t \to 0} \frac{V[X(t+\Delta t) - X(t) | X(t) = x]}{\Delta t}.$$

On the other hand, given the particle at x at time t, then at time $t + \Delta t$ the mean value of change of position is as follows

$$\Delta x (\pi_{\Delta x}^{(+)}(x) - \pi_{\Delta x}^{(-)}(x))$$

and the corresponding variance is given by

$$\Delta x^2 (\pi_{\Delta x}^{(+)}(x) + \pi_{\Delta x}^{(-)}(x) - (\pi_{\Delta x}^{(+)}(x) - \pi_{\Delta x}^{(-)}(x))^2).$$

The Fokker-Planck (-Kolmogorov) equation (step 3)

For the limiting process, we require

$$\beta(x) = \lim_{\Delta x, \Delta t \to 0} (\pi^{(+)}_{\Delta x}(x) - \pi^{(-)}_{\Delta x}(x)) \frac{\Delta x}{\Delta t}$$

and

$$\alpha(x) = \lim_{\Delta x, \Delta t \to 0} (\pi^{(+)}_{\Delta x}(x) - \pi^{(-)}_{\Delta x}(x) - (\pi^{(+)}_{\Delta x}(x) - \pi^{(-)}_{\Delta x}(x))^2) \frac{\Delta x^2}{\Delta t}.$$

These provide constraints for the form of $\pi^{(+)}_{\Delta x}(x)$ and $\pi^{(-)}_{\Delta x}(x)$. We suppose the scale law $(\Delta x)^2 = A \Delta t$ (Wiener or Gaussian white noise). The choices

$$\pi^{(+)}_{\Delta x}(x) = \frac{1}{2A} (\alpha(x) + \beta(x) \Delta x)$$

and

$$\pi^{(-)}_{\Delta x}(x) = \frac{1}{2A} (\alpha(x) - \beta(x) \Delta x)$$

satisfy the above constraints. We require $\alpha(x) \geq \beta(x) \Delta x$.

The Fokker-Planck (-Kolmogorov) equation (step 4)

By expanding in Taylor series (step 1) up to second order, we obtain

$$
\begin{aligned}
p \;&\simeq\; (p - p_x \Delta x + \tfrac{1}{2} p_{xx} \Delta x^2 - p_t \Delta t)(\pi_{\Delta x}^{(+)} - \pi_{\Delta x}^{(+)'} \Delta x + \tfrac{1}{2} \pi_{\Delta x}^{(+)''} \Delta x^2) \\
&+\; (p + p_x \Delta x + \tfrac{1}{2} p_{xx} \Delta x^2 - p_t \Delta t)(\pi_{\Delta x}^{(+)} + \pi_{\Delta x}^{(+)'} \Delta x + \tfrac{1}{2} \pi_{\Delta x}^{(+)''} \Delta x^2) \\
&+\; (p - p_t \Delta t)(1 - \pi_{\Delta x}^{(+)} - \pi_{\Delta x}^{(-)}).
\end{aligned}
$$

Finally, by using the constraints for α and β, and the scale law, we obtain the Fokker-Planck equation

$$
\partial_t p(x_0, x; t) = \frac{1}{2} \partial_{xx}^2 (\alpha(x) p(x_0, x; t)) - \partial_x (\beta(x) p(x_0, x; t)).
$$

A new approach based on the Fokker-Planck equation

The evolution of the PDF given by $f = f(x,t)$, $x \in \Omega \subset \mathbb{R}^n$, associated to the stochastic process is modelled by the Fokker-Planck (FP) equation.

$$\partial_t f - \frac{1}{2} \sum_{i,j=1}^{n} \partial^2_{x_i x_j} (a_{ij}\, f) + \sum_{i=1}^{n} \partial_{x_i} (b_i(u)\, f) = 0$$

$$f(t_0) = \rho$$

This is a partial differential equation of parabolic type with Cauchy data given by the initial PDF distribution.
The formulation of objectives with the PDF and the Fokker-Planck equation provide a consistent framework to the optimal control of stochastic processes.

The transition density probability and the PDF

Denote with $f(x,t)$ the probability density to find the process at $x = (x_1, \dots, x_n)$ at time t.

Let $\hat{f}(x,t;y,s)$ denotes the transition density probability distribution function for the stochastic process to move from $y \in \mathbb{R}^n$ at time s to $x \in \mathbb{R}^n$ at time t.

Both $f(x,t)$ and $\hat{f}(x,t;y,s)$ are nonnegative functions and the following holds

$$\hat{f}(x,t|y,s) \geq 0, \qquad \int_\Omega \hat{f}(x,t|y,s)\,dx = 1 \quad \text{for all } t \geq s.$$

Given an initial PDF $\rho(y,s)$ at time s, we have the following

$$f(x,t) = \int_\Omega \hat{f}(x,t|y,s)\rho(y,s)\,dy, \qquad t > s.$$

Also ρ should be nonnegative and $\int_\Omega \rho(y,s)dy = 1$.

A tracking objective

We consider the control problem formulated in the time window $(t_k,\ t_{k+1})$ with known initial value at time t_k.

We formulate the problem to determine a piecewise constant control $u(t) \in \mathbb{R}^\ell$ such that the process evolves towards a desired target probability density $f_d(x,t)$ at time $t = t_{k+1}$.

This objective can be formulated by the the following tracking functional

$$J(f,u) := \frac{1}{2}\|f(\cdot,t_{k+1}) - f_d(\cdot,t_{k+1})\|^2_{L^2(\Omega)} + \frac{\nu}{2}|u|^2.$$

where $|u|^2 = u_1^2 + \ldots + u_\ell^2$.

A Fokker-Planck optimal control problem

The optimal control problem to find u that minimizes the objective J subject to the constraint given by the FP equation is formulated by the following

$$\min J(f,u) := \frac{1}{2}\|f(\cdot,t_{k+1}) - f_d(\cdot,t_{k+1})\|^2_{L^2(\Omega)} + \frac{\nu}{2}|u|^2$$

$$\partial_t f - \frac{1}{2}\sum_{i,j=1}^{n} \partial^2_{x_i x_j}(a_{ij}\, f) + \sum_{i=1}^{n} \partial_{x_i}(b_i(u)\, f) = 0$$

$$f(t_k) = \rho.$$

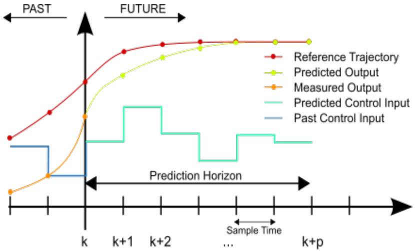

A Fokker-Planck optimality system

This first-order necessary optimality condition is characterized as the solution of the following optimality system

$$
\begin{aligned}
\partial_t f - \tfrac{1}{2} \textstyle\sum_{i,j=1}^n \partial^2_{x_i x_j} (a_{ij} f) + \sum_{i=1}^n \partial_{x_i} (b_i(u) f) &= 0 \quad \text{in } Q_k \\
f(x, t_k) &= \rho(x) \quad \text{in } \Omega \\
-\partial_t p - \tfrac{1}{2} \textstyle\sum_{i,j=1}^n a_{ij} \partial^2_{x_i x_j} p - \sum_{i=1}^n b_i(u) \partial_{x_i} p &= 0 \quad \text{in } Q_k \\
p(x, t_{k+1}) = f(x, t_{k+1}) - f_d(x, t_{k+1}) & \quad\ \ \text{in } \Omega \\
f = 0,\ p &= 0 \quad \text{on } \Sigma_k \\
\nu\, u_l + \left(\textstyle\sum_{i=1}^n \partial_{x_i} \big(\frac{\partial b_i}{\partial u_l} f \big), p \right) = 0 \quad \text{in } Q_k \quad l &= 1, \dots, \ell
\end{aligned}
$$

where $Q_k = \Omega \times (t_k, t_{k+1})$ and $\Sigma_k = \partial\Omega \times (t_k, t_{k+1})$.

The reduced gradient

In the optimality equation, we have used the following inner product

$$(\phi, \psi) = \int_{t_k}^{t_{k+1}} \int_{\Omega} \phi(x,t)\, \psi(x,t)\, dx\, dt.$$

The lth component of the reduced gradient $\nabla \hat{J}$ is given by

$$(\nabla \hat{J})_l = \nu\, u_l + \left(\sum_{i=1}^{n} \partial_{x_i} \left(\frac{\partial b_i}{\partial u_l} f \right), p \right), \qquad l = 1, \ldots, \ell,$$

where $p = p(u)$ is the solution of the adjoint equation for given $f(u)$.

Notice that we are discussing a nonlinear control mechanism and thus the optimization problem is nonconvex.

Theory of the FP control problem (step 1)

Consider the following FP equation, $\mathcal{E}(f_0, u, g) = 0$, as follows

$$\begin{aligned} \partial_t f(x,t) - a\,\partial^2_{xx} f(x,t) + \partial_x \left(b(x;u)\, f(x,t) \right) &= g(x,t) \\ f(x,t) &= 0 \\ f(x,0) &= f_0(x) \end{aligned}$$

where $a > 0$, $u \in U$, and $f_0 \in L^2(\Omega)$. Consider the Ornstein-Uhlenbeck process with $b(x;u) = -\gamma x + u$. We introduce a source term $g \in L^2(Q)$.
Define $V = H^1_0(\Omega)$ and $V' = H^{-1}(\Omega)$ its dual with $(\cdot,\cdot)_{V'V}$ the duality pairing and $H = L^2(\Omega)$ the pivot space. We consider the space

$$W = \{ w \in L^2(0,T;V),\ \dot{w} = L^2(0,T;V') \}$$

with norm $\|w\|^2_W = \|w\|^2_{L^2(V)} + \|\dot{w}\|^2_{L^2(V')}$.

Theory of the FP control problem (step 2: lemma)

Assume that $b(x;u) = \gamma(x) + u$, $\gamma \in C^1(\Omega)$, $f_0 \in H$, $u \in U$, and $g \in L^2(V')$. Then if f is a solution to $\mathcal{E}(f_0, u, g) = 0$, the following inequalities hold.

$$\|f\|_{L^2(V)} \leq \frac{1}{s\sqrt{2}}\|f_0\|_H + \frac{1}{s^2}\|g\|_{L^2(V')}$$

$$\|f\|_{L^\infty(H)} \leq \|f_0\|_H + \alpha_1\|g\|_{L^2(V')}$$

$$\begin{aligned} \|\dot{f}\|_{L^2(V')} &\leq (\|u\|_U + \bar{\gamma})\left(\|f_0\|_H + \alpha_1\|g\|_{L^2(V')}\right) \\ &+ \alpha_2\left(\frac{1}{s\sqrt{2}}\|f_0\|_H + \frac{1}{s^2}\|g\|_{L^2(V')}\right) + \|g\|_{L^2(V')} \end{aligned}$$

where $s = \sqrt{\left(\frac{a}{1+c_{PF}}\right) - \bar{\gamma}}$, $\bar{\gamma} = \max_{x\in\Omega}(|\gamma(x)|, |\gamma'(x)|)$ is sufficiently small, c_{PF} is the Poincaré - Friedrichs constant corresponding to Ω, $\alpha_1 = \max\left(\frac{1}{\sqrt{2}}, \frac{\sqrt{2}}{\sqrt{s}}\right)$, and α_2 satisfies the following condition $a\|\partial_{xx}\varphi\|_{V'} \leq \alpha_2\|\varphi\|_V$, $\forall \varphi \in V$.

Theory of the FP control problem (step 3: propositions)

Proposition

Assume that $b(x; u) = \gamma(x) + u$, $\gamma \in C^1(\Omega)$ *and sufficiently small* $\bar{\gamma} = \max_{x\in\Omega}(|\gamma(x)|, |\gamma'(x)|)$, $f_0 \in H$, *and* $u \in U$. *Then the problem* $\mathcal{E}(f_0, u, 0) = 0$ *admits a unique solution* f *in* $L^2(V) \cap L^\infty(H)$ *with* $\dot{f} \in L^2(V')$. *In particular, we have* $f \in C([0, T]; H)$.

Proposition

The mapping $\Lambda : U \to C([0, T]; H)$, $u \to f = \Lambda(u)$ *is the solution to* $\mathcal{E}(f_0, u, 0) = 0$, *is Fréchet differentiable and* $\Lambda'_{u^*} \cdot h$ *satisfies the equation*

$$\dot{e} + Ae = u^* Be + hBf^* + Ce$$
$$e(0) = 0,$$

where $f^* = \Lambda(u^*)$.

Theory of the FP control problem (step 4: propositions)

Proposition

The functional $\hat{J}(u)$ is differentiable and we have the derivative

$$d\hat{J}(u) \cdot h = \left(\nu\, u + \int_0^T (u_x f, p)_{V'V}\, dt, h \right)_U, \qquad \forall h \in U,$$

where p is the solution to the adjoint equation

$$-u_t p - a\, u_{xx}^2\, p - b(x;u)\, u_x\, p = 0, \quad p(x,T) = f(x,T) - f_d(x),$$

and f is the solution to $\mathcal{E}(f_0, u, 0) = 0$.

Theory of the FP control problem (step 5: propositions)

Proposition

Let f_1^* and f_2^* be the states corresponding to the optimal controls u_1^* and u_2^*, respectively. Further, let p_1^* and p_2^* be the adjoint states corresponding to the optimal controls u_1^* and u_2^*, respectively. Under the assumption of Lemma 1, the following inequalities hold

$$\|f_j^*\|_{L^2(V)} \leq \frac{1}{s\sqrt{2}}\|f_0\|_H, \qquad j = 1, 2,$$

$$\|f_j^*\|_{L^\infty(H)} \leq \|f_0\|_H, \qquad j = 1, 2,$$

$$\|p_j^*\|_{L^2(V)} \leq \frac{1}{s\sqrt{2}}\|f_j(T) - f_d\|_H, \qquad j = 1, 2,$$

$$\|p_j^*\|_{L^\infty(H)} \leq \|f_j(T) - f_d\|_H, \qquad j = 1, 2.$$

Theory of the FP control problem (step 6: uniqueness)

Proposition

Using previous estimates and for sufficiently small initial condition, i.e. small $\|f_0\|_H$, a unique optimal control exists.

$$\nu \|u_1 - u_2\|_U \leq \left(\frac{1}{s^2} \|f_1(T) - f_d\|_H + \frac{\alpha_1}{s\sqrt{2}} \|f_0\|_H + \frac{1}{s^2} \|f_2(T) - f_d\|_H \right) \|u_1 - u_2\|_U \|f_0\|_H$$

The discretization of the FP optimality system

The forward- and adjoint FP equations are discretized using the second-order backward time-differentiation formula (BDF2) as follows

$$\partial^-_{BD} y^m_{\mathbf{i}} := \frac{3y^m_{\mathbf{i}} - 4y^{m-1}_{\mathbf{i}} + y^{m-2}_{\mathbf{i}}}{2\delta t} \qquad \partial^+_{BD} p^m_{\mathbf{i}} := -\frac{3p^m_{\mathbf{i}} - 4p^{m+1}_{\mathbf{i}} + p^{m+2}_{\mathbf{i}}}{2\delta t}.$$

For spatial-discretization we use the Chang-Cooper (CC) scheme that is stable, second-order accurate, positive, and conservative.

The Chang-Cooper scheme

The FP equation can be written in flux form, $\partial_t f = \nabla \cdot F$, where

$$\nabla \cdot F \approx \frac{1}{h} \sum_{i=1}^{n} (F^i_{\mathbf{i}+\imath_i/2} - F^i_{\mathbf{i}-\imath_i/2}).$$

The flux in the i-th direction is computed as follows

$$F^i_{\mathbf{i}+\imath_i/2} = \left[(1-\delta_i) B^{i,n}_{\mathbf{i}+\imath_i/2} + \frac{1}{h} C^{i,n}_{\mathbf{i}+\imath_i/2}\right] f^{n+1}_{\mathbf{i}+\imath_i} - \left(\frac{1}{h} C^{i,n}_{\mathbf{i}+\imath_i/2} - \delta_i B^{i,n}_{\mathbf{i}+\imath_i/2}\right) f^{n+1}_{\mathbf{i}}$$

where we set

$$B^i(x,t,u) = \frac{1}{2} \sum_{j=1}^{n} \partial_{x_j} a_{ij}(x,t) - b_i(x,t;u) \qquad C^i(x,t) = \frac{1}{2} a_{ii}(x,t)$$

and use the following (CC) linear spatial combination of f^{n+1}

$$f^{n+1}_{\mathbf{i}+\imath_i/2} = (1-\delta_i)\, f^{n+1}_{\mathbf{i}+\imath_i} + \delta_i\, f^{n+1}_{\mathbf{i}}, \qquad \delta_i \in [0, 1/2].$$

choosing $\delta_i = \frac{1}{w_i} - \frac{1}{\exp(w_i)-1}$ where $w_i = h\, B^{i,n}_{\mathbf{i}+\imath_i/2} / C^{i,n}_{\mathbf{i}+\imath_i/2}$.

A receding horizon model predictive control scheme

Let $(0, T)$ be the time interval where the process is considered. We assume time windows of size $\Delta t = T/N$ with N a positive integer. Let $t_k = k\Delta t$, $k = 0, 1, \ldots, N$. At time t_0, we have a given initial PDF denoted with ρ and with $f_d(\cdot, t_k)$, $k = 1, \ldots, N$, we denote the sequence of desired PDFs.

Algorithm (RH-MPC)

Set $k = 0$*; assign the initial PDF,* $f(x, t_k) = \rho(x)$ *and the targets* $f_d(\cdot, t_k)$, $k = 0, \ldots, N-1$*;*

1. *In* (t_k, t_{k+1})*, solve* $\min_u J(f(u), u)$.
2. *With the optimal solution* u *compute* $f(\cdot, t_{k+1})$.
3. *Assign this PDF as the initial condition for the FP problem in the next time window.*
4. *If* $t_{k+1} < T$*, set* $k := k + 1$*, go to 1. and repeat.*
5. *End.*

The solution of the optimization problem

In Step 1. of RH-MPC, we need to solve $\min_{u \in \mathbb{R}^\ell} J(f(u), u)$. For this purpose, we implement a nonlinear conjugate gradient (NCG) scheme with Dai and Yuan β and a robust bisection linesearch.

Algorithm (NCG Scheme)

- *Input: initial approx.* u_0, $d_0 = -\nabla \hat{J}(u_0)$, *index* $k = 0$, *maximum* k_{max}, *tolerance* tol.
  1. *While* $(k < k_{max}$ *&&* $\|g_k\|_{\mathbb{R}^\ell} > tol$ *) do*
  2. *Use Algorithm Bisection to search steplength* $\alpha_k > 0$ *along* d_k *satisfying the Armijo - Wolfe conditions;*
  3. *Set* $u_{k+1} = u_k + \alpha_k d_k$;
  4. *Compute* $g_{k+1} = \nabla \hat{J}(u_{k+1})$;
  5. *Compute* β_k^{DY}
  6. *Let* $d_{k+1} = -g_{k+1} + \beta_k^{DY} d_k$;
  7. *Set* $k = k + 1$;
  8. *End while*

Application to one-dimensional problems

A Ornstein-Uhlenbeck process with additive control: a massive particle immersed in a viscous fluid and subject to random Brownian fluctuations due to interaction with other particles.

$$b(X_t, t; u) = -\gamma X_t + u, \qquad \sigma(X_t, t) = \sigma$$

where X_t represents the velocity of the particle and u is the momentum induced by an external force field.

A geometric-Brownian process with additive drift control: The classical Merton's portfolio problem models the wealth and a wide variety of exotic options and other derivative contracts.

$$b(X_t, t; u) = (\mu + u)X_t \qquad \sigma(X_t, t) = \sigma X_t$$

where X_t is the wealth and u represents a fraction of the portfolio invested in a risk free and constant interest rate market.

A Ornstein-Uhlenbeck process with additive control

The initial distribution is a Gaussian with zero mean and variance $\sigma = 0.1$. The target is also Gaussian with mean value following the law $x(t) = 2\sin(\pi t/5)$ and variance $\sigma = 0.2$. We have time windows of $\Delta t = 0.5$ and $T = 5$. Parameter values $\gamma = 1$, $\nu = 0.1$.

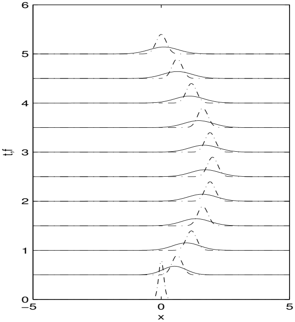

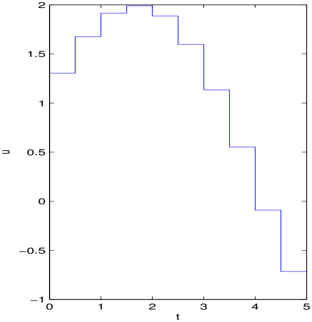

A geometric-Brownian process with additive drift control

The initial and target distributions are in the log-normal form

$$f_d(x,t) = \frac{1}{x\sqrt{2\pi\sigma^2}} \exp\left(-\frac{[\log(x)-\tilde{\mu}(t)]^2}{2\sigma^2}\right).$$

where for the initial distribution $\tilde{\mu}(t_0) = 0.8$, $\sigma = 0.1$, and for the target distribution $\tilde{\mu}(t) = 1 + \sin(\pi t/5)$ and $\sigma = 0.1$.
We have $\Delta t = 0.25$ and $T = 2.5$. Parameter values $\mu = 1, \sigma = 0.1$ and $\nu = 0.1$.

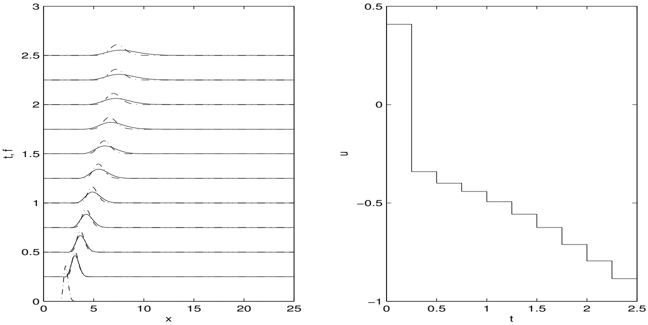

Application to multidimensional problems

A two-species generalized stochastic Lotka-Volterra prey-predator model

$$\begin{cases} dX_1 = b_1(X_1, X_2; u_1)dt + \sigma_1(X_1)dW_{1t} \\ dX_2 = b_2(X_1, X_2; u_2)dt + \sigma_2(X_2)dW_{2t} \end{cases}$$

where $X_1(t)$ and $X_2(t)$ represent populations of prey and predators, respectively.

The drift terms including the controllers u_1 and u_2 are as follows

$$\begin{cases} b_1(X_1, X_2; u_1) = a_1 X_1 - b_1 X_1^2 - c X_1 X_2 + u_1 \\ b_2(X_1, X_2; u_2) = a_2 X_2 - b_2 X_2^2 + c X_1 X_2 + u_2 \end{cases}$$

Here, u_1 and u_2 represent the rate of release of population species

The diffusion is $\sigma_1(X_1) = \sigma\sqrt{b_1 X_1^2}$ and $\sigma_2(X_2) = \sigma\sqrt{b_2 X_2^2}$.

Fast stabilization of the stochastic Lotka-Volterra model

The equilibrium PDF (at $t \to \infty$) is given by the following

$$f_d(x_1, x_2) = m \left[\frac{1}{x_1} \exp\left(\frac{2A_1}{\sigma^2} \log(x_1) - \frac{2}{\sigma^2}(x_1 - 1) \right) \right] \times \left[\frac{1}{x_2} \exp\left(\frac{2A_2}{\sigma^2} \log(x_2) - \frac{2}{\sigma^2}(x_2 - 1) \right) \right].$$

We choose $T = 10$ and time windows of size $\Delta t = 1$. Control weights $\nu = 0.1$ and $\nu = 0.001$. Dashed and dot-dashed lines are u_1, u_2. Solid line represents $\|f(\cdot, t_k) - f_d(\cdot, T)\|_\infty$ with controlled f.

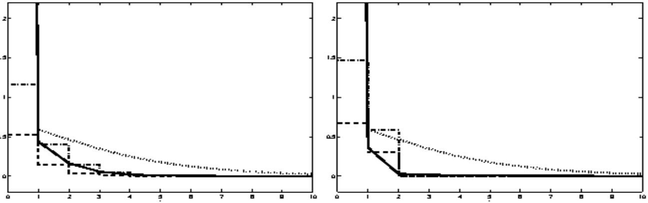

Tracking a trajectory with a limit-cycle model

Consider a noised limit cycle equation with control as follows

$$
\begin{aligned}
dX_1 &= (X_2 + (1 + u_1 - X_1^2 - X_2^2)X_1)\, dt \quad &+\sigma dW_{1t} \\
dX_2 &= (-X_1 + (1 + u_2 - X_1^2 - X_2^2)X_2)\, dt \quad &+\sigma dW_{2t}.
\end{aligned}
$$

The purpose of the control is to track the target given by a bi-modal multivariate Gaussian PDF

$$
f_d = \frac{1}{2}\frac{\exp\left(-\frac{(x_1-\mu_{11})^2}{2\sigma_{11}^2} - \frac{(x_2-\mu_{21})^2}{2\sigma_{21}^2}\right)}{2\pi\sigma_{11}\sigma_{21}} + \frac{1}{2}\frac{\exp\left(-\frac{(x_1-\mu_{12})^2}{2\sigma_{12}^2} - \frac{(x_2-\mu_{22})^2}{2\sigma_{22}^2}\right)}{2\pi\sigma_{12}\sigma_{22}}
$$

with peaks placed symmetrically with respect to the origin at the points $(\mu_{11}, \mu_{21}) = (-1.2, 0.8)$ and $(\mu_{12}, \mu_{22}) = (1.2, -0.8)$
We have $T = 30$ and the time-window size is $\Delta t = 5$.

A controlled noised limit-cycle model

The Fokker-Planck RH-MPC control strategy is able to drive the system to a bi-modal PDF configuration (!) starting from a initial approximate delta-Dirac PDF located at the point $(1.5, 1.5)$.

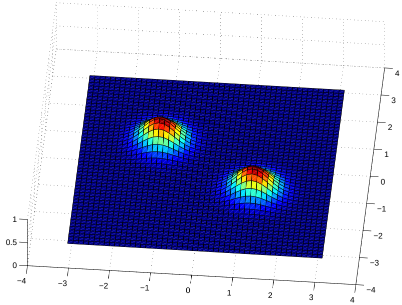

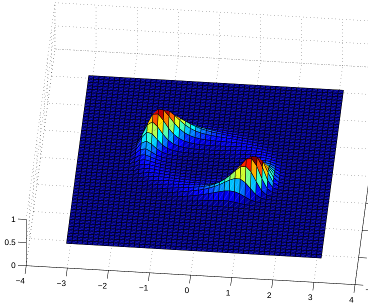

Conclusion and thanks

A novel Fokker-Planck optimization framework for determining controls of the PDF of multidimensional stochastic processes was presented.

The control strategy was based on a receding-horizon model predictive control scheme where optimal controls were obtained minimizing a deterministic PDF objective under the constraint given by the Fokker-Planck equation that models the evolution of the probability density function.

Thanks a lot for your attention