Domain Predictive Control Under Uncertain Numerical State Information

Johannes Löhr and **Patrick Eyerich**
University of Freiburg, Germany
{loehr, eyerich}@tf.uni-freiburg.de

Stefan Winkler
Astrium GmbH - Satellites
stefan.winkler.sat@astrium.eads.net

Bernhard Nebel
University of Freiburg, Germany
nebel@tf.uni-freiburg.de

Abstract

In planning, hybrid system states consisting of logical and numerical variables are usually assumed to be completely known. In particular, for numerical state variables full knowledge of their exact values is assumed. However, in real world applications states are results of noisy measurements and imperfect actuators. Therefore, a planned sequence of state transitions might fail to lead a hybrid system to the desired goal. We show how to propagate and reason about uncertain state information directly in the planning process, enabling hybrid systems to find plans that satisfy numerical goals with predefined confidence.

Introduction

Planning is decision-making by reasoning about predicted future states which emerge from a current state, called *initial state*. Hybrid systems have numerical and logical state variables. Which confidence has a plan in the case of uncertainties within the numerical part of the initial state?

Domain Predictive Control (DPC) is a planning-based control approach applicable to hybrid systems with finite sets of discretizable inputs and modes, e.g., actuator, sensor, or controller configurations. It suits best to systems that involve many logical dependencies (Löhr et al. 2012). The inherent convergence of numerical states (stability) of switching systems cannot be guaranteed in general (Liberzon 2003). Therefore, DPC generates interleaved sequences of control signals and mode switches by utilizing plannig methods that explicitly calculate and heuristically evaluate future states to avoid instable evolutions.

DPC has been applied in numerical time simulations to safely land spacecraft on extraterrestrial surfaces, reacting on changing environment information and system failures by continuous replanning (Löhr, Nebel, and Winkler 2012).

Handling of uncertainties is very common in control of dynamic systems (Gelb 1974). Also in path planning the consideration of uncertainties leads to more robust results (Blackmore, Li, and Williams 2006). Especially when estimating system states by using stochastic filters like Kalman-Filters, uncertainties play an important role (Kalman 1960).

The contribution of this paper is an extension of DPC that allows it to guarantee predefined confidences of generated plans by propagating uncertainties in the initial state to predicted future states. While this extension does not directly guarantee deployable systems based on DPC, it brings DPC more close to applicability for real-world scenarios.

The remainder of the paper is structured as follows. After giving some background information we show how to consider uncertainties and noise in planning actions. Then the need of planning with uncertainties in real world applications is shown by exemplifying the approach to orbital maneuver planning of spacecraft. Finally, we conclude.

Planning Hybrid System Transitions

Domain Predictive Control is an approach that utilizes action planning for the control of dynamic systems. It generates logically consistent control strategies for hybrid systems by planning sequences $\pi = \langle a_1, a_2, \ldots, a_z \rangle$ of *actions* a_i, called *plans*, that lead dynamic systems from an *initial state* s_0 to a state satisfying the *goal conditions* $s_z \in s^{\star}$. An action a_i consists of numerical input signals defined over the duration δ_i of a_i, and specifies which sensors, actuators and controllers are active.

Modeling Disturbed Systems

A *hybrid system state* $[\mathbf{x}_n^T, \mathbf{x}_l^T]^T$ is composed of a numerical part $\mathbf{x}_n$ and a logical part $\mathbf{x}_l$. The numerical states $\mathbf{x}_n \in \mathbb{R}^p$ evolve according to the *plant dynamics*

$$\dot{\mathbf{x}}_n = A\mathbf{x}_n + B_i\mathbf{u} + G_i\mathbf{w}, \tag{1}$$

where $A \in \mathbb{R}^{p \times p}$ is the *state matrix*, $B \in \mathbb{R}^{p \times m}$ is the *input matrix*, and $G \in \mathbb{R}^{p \times g}$ is the *noise input matrix*. The actuator noise $\mathbf{w} \in \mathbb{R}^g$ is a zero-mean Gaussian white signal.

A vector of q *measurements* $\mathbf{y} \in \mathbb{R}^q$, provided by sensors, is given by

$$\mathbf{y} = C_i\mathbf{x}_n + \mathbf{n}, \tag{2}$$

where $C \in \mathbb{R}^{q \times p}$ is the *measurement matrix* and the sensor noise $\mathbf{n} \in \mathbb{R}^q$ is a zero-mean Gaussian white signal.

The *input signal*

$$\mathbf{u} = \mathbf{u}^{ff} + \mathbf{u}^{cl} + \mathbf{u}^{ef} \in \mathbb{R}^m \tag{3}$$

consists of a *feed-forward input* $\mathbf{u}^{ff}$, an error state feedback $\mathbf{u}^{ef}$, introduced in the paragraph on reference tracking below, and a *closed-loop input*

$$\mathbf{u}^{cl} = -K_i\mathbf{y} \tag{4}$$

due to measurement feedback via controller gain $K \in \mathbb{R}^{m \times q}$.

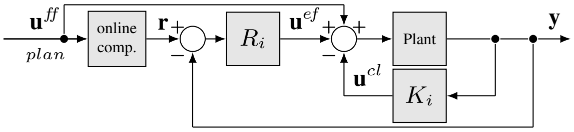

Figure 1: Tracking the computed trajectory of the plan.

True State, Error and Estimate The *true state* vector

$$\mathbf{x}_n = \hat{\mathbf{x}}_n + \mathbf{e} \quad (5)$$

is composed of the *estimate* $\hat{\mathbf{x}}_n$ and the *error* $\mathbf{e} = [e_1, e_2, \ldots, e_{\mathtt{p}}]^T$, where e_i are random variables defining a normally distributed uncertainty on the estimates. Inserting Eqs. 2 to 5 into Eq. 1 leads to the dynamics of the true state in dependency of the estimate and the error

$$\dot{\mathbf{x}}_n = \dot{\hat{\mathbf{x}}}_n + \dot{\mathbf{e}} = A_{cl,i}\hat{\mathbf{x}}_n + A_{cl,i}\mathbf{e} - B_i K_i \mathbf{n} + G_i \mathbf{w} + B_i(\mathbf{u}^{ff} + \mathbf{u}^{ef}) \quad (6)$$

where $A_{cl,i} = A - B_i K_i C_i$ is the *closed loop state matrix* switched between actions a_i, but invariant during an action.

Reference Tracking During the execution of z planned actions the system can continuously compare its measurement with the expected measurements called *reference* $\mathbf{r}(t) \equiv \hat{\mathbf{y}}(t) = C_i \hat{\mathbf{x}}_n(t), t \in [t_0, t_0 + \sum_{i=1}^{z} \delta_i]$ that can be obtained by online trajectory computation of the plan. This enables (future) generation of the *error feedback signal*

$$\mathbf{u}^{ef} = R_i\ (\mathbf{r} - \mathbf{y}) = -R_i(C_i \mathbf{e} + \mathbf{n}) \quad (7)$$

reducing the errors $\mathbf{e}$ using the control matrix $R_i \in \mathbb{R}^{m \times q}$, as depicted in Figure 1. Effects of future error state feedback can be accounted for directly in the planning process.

Inserting Eq. 7 into Eq. 6, the state dynamics becomes

$$\dot{\hat{\mathbf{x}}}_n + \dot{\mathbf{e}} = A_{cl,i}\hat{\mathbf{x}}_n + F_{cl,i}\mathbf{e} - B_i(K_i + R_i)\mathbf{n} + G_i\mathbf{w} + B_i\mathbf{u}^{ff}, \quad (8)$$

where $F_{cl,i} = A - B_i(K_i + R_i)C_i$ is the *closed loop error state matrix* . The error dynamics is called *stable* if all eigenvalues of $F_{cl,i}$ have negative real parts resulting in error states reducing with time (Bellman 1953). It is worth mentioning that the evolution of $\hat{\mathbf{x}}_n$ is not affected by R_i. In the following the evolution of the estimated states and the error dynamics are treated separately.

Planning the Evolution of the Estimate

Instead of planning with the unknown true numerical state vector $\mathbf{x}_n$ we plan with the estimate, which is the *expected value* $E\{\mathbf{x}_n\} = \hat{\mathbf{x}}_n$ of $\mathbf{x}_n$ that is usually available in real dynamic systems from state estimation. Since $E\{\mathbf{e}\} = \mathbf{0}$ and with (zero-mean) white noise assumption of $\mathbf{w}$ and $\mathbf{n}$, the expected evolution of Eqs. 6 and 8 is given by

$$\dot{\hat{\mathbf{x}}}_n(t) = A_{cl,i}\ \hat{\mathbf{x}}_n(t) + B_i \mathbf{u}^{ff}(t) \quad (9)$$

within the time interval $t \in [t_k, t_k + \delta_i]$ of an action a_i. The effect $E_{\dashv}$ of an action a_i on the estimated numerical state can be obtained by discretization of Eq. 9

$$E_{\dashv}(a_i) : \hat{\mathbf{x}}_n(t_k + \delta_i) = \Phi_{t_k+\delta_i, t_k}\ \hat{\mathbf{x}}_n(t_k) + \Psi_i, \quad (10)$$

$$\text{where } \Phi_{\eta,\xi} = e^{A_{cl,i}\ (\eta - \xi)},$$

$$\text{and } \Psi_i = \int_{t_k}^{t_k+\delta_i} \left(\Phi_{t_k+\delta_i,\tau} B_i \mathbf{u}^{ff}(\tau)\right) d\tau,$$

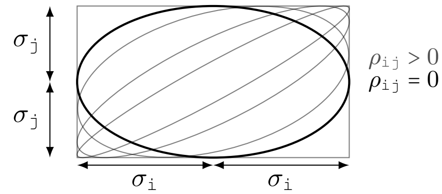

Figure 2: Upper border of covariance ellipses for two random variables $e_{\mathtt{i}}$ and $e_{\mathtt{j}}$ with standard deviations $\sigma_{\mathtt{i}}$ and $\sigma_{\mathtt{j}}$ and varying correlation coefficient $\rho_{\mathtt{ij}}$.

where $\Phi_{\eta,\xi}$ is the state transition matrix from time ξ to η denoted by Φ_i for actions a_i. This prediction of future numerical estimations expressible in PDDL 2.1 (Fox and Long 2003) and can be performed directly in the planning process, whereas Φ_i and Ψ_i can be precomputed in a way much as in Löhr et al. (2012).

Planning the Evolution of the Error

The evolution of the error $\mathbf{e}$ on the estimates in Eq. 9 is given by the remaining terms of Eq. 8:

$$\dot{\mathbf{e}}(t) = \underbrace{F_{cl,i}\ \mathbf{e}(t)}_{\text{error propagation}} \underbrace{- B_i(K_i + R_i)\mathbf{n}(t)}_{\text{sensor noise}} + \underbrace{G_i \mathbf{w}(t)}_{\text{actuator noise}}. \quad (11)$$

Error Representation The direct propagation of Eq. 11 is not possible, since both error $\mathbf{e}$ and the time series of the random signals $\mathbf{n}(t)$ and $\mathbf{w}(t)$ for $t > t_0$ are unknown. Therefore, we plan with the second central moment $E\{\mathbf{e}\,\mathbf{e}^T\}$ also known as the *covariance matrix* P. The covariance matrix defines the confidence of the current estimated state and should be part of the planning states. The errors can be cross correlated resulting in a symmetrically filled matrix

$$P = \begin{bmatrix} \sigma_{\mathtt{1}}^2 & \sigma_{\mathtt{12}}^2 & \ldots & \sigma_{\mathtt{1p}}^2 \\ \sigma_{\mathtt{21}}^2 & \sigma_{\mathtt{2}}^2 & & \vdots \\ \vdots & & \ddots & \\ \sigma_{\mathtt{p1}}^2 & \ldots & & \sigma_{\mathtt{p}}^2 \end{bmatrix}, \quad P \in \mathbb{R}^{p \times p}, \quad (12)$$

where the diagonal elements $\sigma_{\mathtt{i}}^2$ are also known as *variance* of the error $e_{\mathtt{i}}$ with standard deviation $\sigma_{\mathtt{i}}$ (Kamen and Su 1999). Two random variables $e_{\mathtt{i}}$ and $e_{\mathtt{j}}$ form elliptical regions of constant probability of presence, called *confidence region*. The ellipse has principal axis $\sigma_{\mathtt{i}}$ and $\sigma_{\mathtt{j}}$ and is compressed and rotated depending on the correlation coefficient $\rho_{\mathtt{ij}} = \frac{\sigma_{\mathtt{ij}}^2}{\sigma_{\mathtt{i}}\sigma_{\mathtt{j}}}$ of both variables. Eigenvalue decompositions for exact calculation of the elliptical confidence region[1] during the planning process can be avoided since $\sigma_{\mathtt{i}}$ and $\sigma_{\mathtt{j}}$ are upper borders as depicted in Figure 2.

We provide the planner with knowledge about the error variance *beside* the estimated state. The covariance matrix P is used as a stochastic entity to describe the error of the estimation requiring $\frac{p^2+p}{2}$ additional state variables $P_{\mathtt{ij}}$ in the domain model.

[1]The covariance ellipses can be generated by transforming Cartesian coordinates of the unit circle with the matrix $T = E^{\frac{1}{2}}\ V$, resulting from the eigenvalue decomposition of P, with eigenvalue matrix E and modal matrix V.

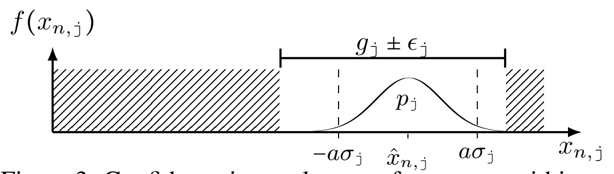

Figure 3: Confidence interval $\pm a\sigma_\mathtt{j}$ of state $x_{n,\mathtt{j}}$ within goal.

Error Propagation Using Eq. 11, we propagate the error covariance through an action a_i, accounting for the error dynamics and the additional noise terms. The error state transition matrix $\Xi_{\eta,\xi} = e^{F_{cl,i}(\eta-\xi)}$ from time ξ to η. Ξ_i is denoted as error transition matrix of a_i with duration δ_i.

- HOMOGENEOUS ERROR EVOLUTION using covariances is well known from e.g. *propagation steps* in Kalman filtering (Gelb 1974). The evolution of the covariance matrix during an action of duration δ_i is given by
$$P(t_k + \delta_i) = \Xi_i \, P(t_k) \, \Xi_i^T.$$
- ACTUATOR NOISE $\mathbf{w}_i$ induces more uncertainty into the system and increases the variance of the error state during execution of an action by
$$\Gamma_\mathbf{w} = E\{(G_i\mathbf{w}(t))\,(G_i\mathbf{w}(\tau))^T\}.$$
$$= \int_{t_k}^{t_k+\delta_i} \Xi_{t_k+\delta_i,\tau} G_i S_i G_i^T \Xi_{t_k+\delta_i,\tau}^T \, d\tau,$$
where S_i is the *power spectral density matrix* of the actuator noise acting during action a_i .
- SENSOR NOISE $\mathbf{n}_i$ induces more uncertainty into the system and increases the covariance by
$$\Gamma_\mathbf{n} = E\{(-B_i(K_i + R_i)\mathbf{n}(t))\,(-B_i(K_i + R_i)\mathbf{n}(\tau))^T\}$$
$$= \int_{t_k}^{t_k+\delta_i} \Xi_{t_k+\delta_i,\tau} B_i(K_i{+}R_i)T_i(K_i{+}R_i)^T B_i^T \Xi_{t_k+\delta_i,\tau}^T \, d\tau,$$
where T_i is the *power spectral density matrix* of the sensor noise while executing action a_i.

Finally, the effect $E_\dashv(a_i)$ of error propagation, sensor and actuator noise on element $P_\mathtt{ij}$ of the covariance matrix is

$$P_\mathtt{ij} = \sum_{m=1}^{p} \sum_{n=1}^{p} \left(\Xi_{i\,\mathtt{jm}} \, P_\mathtt{mn} \, \Xi_{i\,\mathtt{in}}\right) + \Gamma_{\mathbf{w}\,\mathtt{ij}} + \Gamma_{\mathbf{n}\,\mathtt{ij}}, \qquad (13)$$

if $\mathbf{n}, \mathbf{w}$ and $\mathbf{e}$ are zero cross-correlated. The prediction of future error covariances is expressible in PDDL 2.1 (Fox and Long 2003) and can be performed directly in the planning process, whereas Ξ_i, $\Gamma_\mathbf{n}$, and $\Gamma_\mathbf{w}$ can be precomputed.

Plan Confidence

We formulate numerical goals as a conjunction of intervals $x^\star_{n,\mathtt{j}} \in g_\mathtt{j} \pm \epsilon_\mathtt{j}$ usually specified by two conjunctive goals $x^\star_{n,\mathtt{j}} \le g_\mathtt{j} + \epsilon_\mathtt{j} \wedge x^\star_{n,\mathtt{j}} \ge g_\mathtt{j} - \epsilon_\mathtt{j}$, where $\epsilon_\mathtt{j}$ denotes the sufficient accuracy of the desired set point. Considering confidence intervals in goals is enabled by standard deviations

$$\hat{x}_{n,\mathtt{j}} + a\sigma_\mathtt{j} \le g_\mathtt{j} + \epsilon_\mathtt{j} \wedge \hat{x}_{n,\mathtt{j}} - a\sigma_\mathtt{j} \ge g_\mathtt{j} - \epsilon_\mathtt{j}, \qquad (14)$$

see Figure 3. It is convenient to reformulate Eq. 14 in terms of the available variances $P_\mathtt{jj}$ as follows:

$$a^2 \, P_\mathtt{jj} \le (g_\mathtt{j} - \hat{x}_{n,\mathtt{j}} \pm \epsilon_\mathtt{j})^2. \qquad (15)$$

A commonly used scaling factor is $a = 3$ for the confidence interval of Gaussian distributions with probability density function $f(x_{n,\mathtt{j}})$, meaning that the true state $x_{n,\mathtt{j}}$ lies with probability $p_\mathtt{j} = \int_{\hat{x}_{n,\mathtt{j}}-3\sigma_\mathtt{j}}^{\hat{x}_{n,\mathtt{j}}+3\sigma_\mathtt{j}} f(x_{n,\mathtt{j}})\, dx_{n,\mathtt{j}} \approx 99{,}73\%$ within the $3\sigma_\mathtt{j}$ bounds. Formulating N goals component-wise (as in Eq. 15) leads a probability of $\prod_{\mathtt{j}=1}^{N} p_\mathtt{j}$ that trajectories satisfy the confidence intervals and herewith the goals. This is called *confidence* of a plan.

Exemplary Simulation

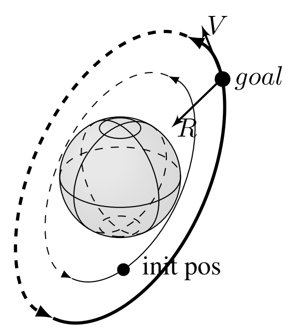

Figure 4: Nadir frame in the spacecraft orbit rendezvous problem.

The motivation to include handling of uncertainties into DPC arose from a spacecraft orbit rendezvous domain, where thrust and coast phases of a satellite have to be combined intelligently in order to reach a desired rendezvous position. As this domain can be discretized and linearized straight-forwardly, DPC is well suited to generate maneuver-based approach strategies in principle. However, because actions are relatively long, small uncertainties in the initial state typically accumulate to large errors in the goal state. Both uncertain initial conditions and noisy thrusters and sensors can affect the final position such that the found plan leads to a potential crash with the other spacecraft. The orbit maneuvers are planned within the *Nadir* frame with origin at the goal position (Figure 4), moving on a circular orbit around Earth. The R axis of the frame points towards the center of the Earth, while the V axis points in flight direction (Fehse 2003).

The orbital relative dynamics can be linearized at the nadir frame (Clohessy and Wiltshire 1960) and yield dynamic and input matrices of Eq. 1

$$A = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 2(\frac{2\pi}{T_P}) \\ 0 & 3(\frac{2\pi}{T_P})^2 & -2(\frac{2\pi}{T_P}) & 0 \end{bmatrix}, \quad B = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ \frac{1}{m} & 0 \\ 0 & \frac{1}{m} \end{bmatrix},$$

also known as Hill equations (Hill 1878), where T_P is the orbital period time and m is the mass of the spacecraft. The state $\mathbf{x}_n = [x, y, v_x, v_y]^T$ contains the position and velocity components along the R- and V-direction of the spacecraft. The measurement matrix C is the unit matrix $I \in \mathbb{R}^{4\times4}$ and the noise input matrix G equals B.

Planning Domain In this example a simplified domain model with the actions a_i defined in Table 1 is used. The thruster provides a force in flight direction and can either be switched on or off. A controller K is used to avoid movements of the spacecraft in R direction while thrusting in V direction. Additional possibilities of thrusting in $\pm R$ direction resulting in circular fly-around or $-V$ direction thrusts are not considered here for simplicity.

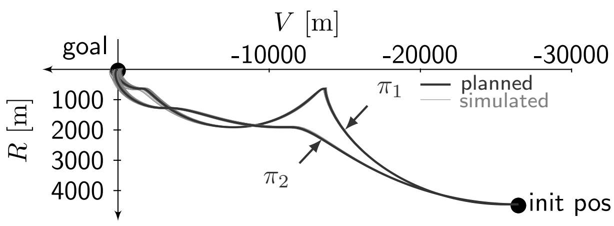

Figure 5: Trajectories of nominal and robust plans. Detailed view on Monte Carlo simulations (light grey) on the right.

The uncertainties due to noisy measurements and actuators are specified in Table 2. The noise of the thrusters $\mathbf{w}$ is assumed to be Gaussian white with power spectral density (PSD) matrix S and the frequency range $[-f_s^{act}/2, f_s^{act}/2]$. The sensor noise $\mathbf{n}$ is also Gaussian white and the frequencies of the power spectral density matrix T lie within the spectrum $[-f_s^{sens}/2, f_s^{sens}/2]$. It is worth to mention that the variance of Gaussian white noise and its PSD A are related by $\sigma^2 = A \int_{-f_s/2}^{f_s/2} df = A f_s$. The initial state knowledge of the spacecraft is given by the initial covariance matrix P_0.

Exemplary Planning Instance A sequence of thrust pulses shall be planned that moves the spacecraft from $\mathbf{x}_{n,0} = [-26483.3\text{m}, 4456.3\text{m}, 7.67\frac{\text{m}}{\text{s}}, 0\frac{\text{m}}{\text{s}}]^T$ to the target orbit with a period of $T_P = 6000\text{s}$ at $\mathbf{x}_n^\star \in [0\text{m}, 0\text{m}, 0\frac{\text{m}}{\text{s}}, 0\frac{\text{m}}{\text{s}}]^T \pm \epsilon$ corresponding to a region around the origin of the Nadir frame $\epsilon = [10\text{m}, 10\text{m}, 0.5\frac{\text{m}}{\text{s}}, 0.5\frac{\text{m}}{\text{s}}]^T$. The uncertainties of the final positions have to lie within the desired orbit area given by the intervals $x \in [-500\text{m}, 500\text{m}]$ and $y \in [-60\text{m}, 60\text{m}]$ due to the adjacent other spacecraft. This is accounted for by adding $N = 2$ constraints $a^2 p_{1,1}^\star \le (\hat{x}_1^\star \pm 500\text{m})^2$ and $a^2 p_{2,2}^\star \le (\hat{x}_1^\star \pm 60\text{m})^2$ to the goal conditions. Choosing $a = 3$ leads to a predefined confidence of at least $0.9973^N = 99.46\%$ of the plan, if the covariance is accounted for during planning.

Results The result of the planning process *without* propagation of the error state covariance using Temporal Fast Downward (Eyerich, Mattmüller, and Röger 2009) is $\pi_1 = \langle a_2, a_1, a_2, a_2, a_1, a_1, a_2, a_3, a_1, a_3 \rangle$. The planned trajectory is shown in Figure 5 and hits the designated goal position. However, uncertainties of the initial states and disturbances during plan execution let the spacecraft possibly crash with the target spacecraft (see Figure 6A).

Propagating the covariance matrix during planning yields the plan $\pi_2 = \langle a_1, a_2, a_2, a_3, a_1, a_2, a_3, a_1, a_2 \rangle$. It reaches the goal position and makes the 3σ confidence ellipse of e_1 and e_2 matching the desired orbit zone, see Figure 6B.

Allowing for *one* reference tracking phase with active error controller R (in action a_4) leads to plan $\pi_3 = \langle a_2, a_1, a_2, a_2, a_1, a_1, a_2, a_3, \mathbf{a_4}, a_3 \rangle$. The evolution of $\hat{\mathbf{x}}_n$ is identical to π_1. However, the planned reference tracking phase actively damps out the error states and leads to a final state satisfying the goal conditions, see Figure 6C.

In conclusion, plans π_2 and π_3, computed in a couple of seconds, ensure that the goal position is reached with the predefined confidence by either reordering actions intelligently (π_2) or by performing reference tracking (π_3). In contrast, planning without reasoning about errors and uncertainties (π_1) does not give any information about the probability to reach the goal position.

a_i	duration δ_i	dynamic controller	error controller	actuator PSD	sensor PSD	$\mathbf{u}^{ff}$
1	$\frac{T_P}{2}$	$\mathbf{0} \in \mathbb{R}^{2\times4}$	$\mathbf{0} \in \mathbb{R}^{2\times4}$	$\mathbf{0} \in \mathbb{R}^{4\times4}$	T	$[0,0]^T$ N
2	100 s	K	$\mathbf{0} \in \mathbb{R}^{2\times4}$	S	T	$[0,1]^T$ N
3	50 s	K	$\mathbf{0} \in \mathbb{R}^{2\times4}$	S	T	$[0,1]^T$ N
4	$\frac{T_P}{2}$	$\mathbf{0} \in \mathbb{R}^{2\times4}$	R	S	T	$[0,0]^T$ N

Table 1: Set of actions.

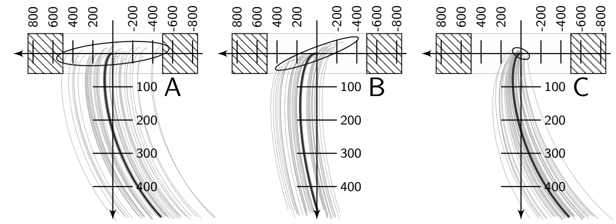

Figure 6: Trajectories close to the goal positions and final 3σ confidence ellipses: (A) *without* and (B) *with* covariance propagation, (C) with planned reference tracking phase.

Parameter	Value
m	300 kg
f_s^{act}, f_s^{sens}	10 Hz, 1 Hz
P_0	$\text{diag}[4.5\text{m}^2, 4\text{m}^2, 3.125\ 10^{-6}\frac{\text{m}^2}{\text{s}^2}, 1.225\ 10^{-5}\frac{\text{m}^2}{\text{s}^2}]$
S	$\text{diag}[10^{-5}\frac{\text{N}^2}{\text{Hz}}, 10^{-5}\frac{\text{N}^2}{\text{Hz}}]$
T	$\text{diag}[4.5\frac{\text{m}^2}{\text{Hz}}, 4\frac{\text{m}^2}{\text{Hz}}, 3.125\ 10^{-6}\frac{\text{m}^2}{\text{s}^2\text{Hz}}, 1.225\ 10^{-5}\frac{\text{m}^2}{\text{s}^2\text{Hz}}]$
K	$[0, 0, 0, 2m(\frac{2\pi}{T_P}); 0, 3m(\frac{2\pi}{T_P})^2, -2m(\frac{2\pi}{T_P}), 0]$
R	$[0.0001, 0, 0.02, 0.6283; 0, 0.0011, -0.6283, 0.02]$

Table 2: Model and simulation parameters.

Conclusion

The success of planning-based guidance of hybrid systems is sensitive to the uncertainty of the initial state and the noise induced by sensors and actuators. Estimates of states and their covariances are usually available in autonomous systems and the power spectral density of induced noises are known properties of sensors and actuators. These are suitable entities for the enhancement of planning domains and the exploitation of this additional information leads to more robust results, as exemplary shown in the spacecraft rendezvous domain. Nevertheless, the propagation of covariances during planning is costly due to the additional numerical states. However, beside intelligent reordering of actions to fulfill a predefined confidence, the propagation of covariances enables the planner to reason about reference tracking phases which opens completely new fields of applications.

Acknowledgments

This work was supported by the German Aerospace Center (DLR) and EADS Astrium-Satellites as part of the project “Kontiplan” (50 RA 1221).

References

Bellman, R. E. 1953. *Stability theory of differential equations*, volume 2. McGraw-Hill New York.

Blackmore, L.; Li, H.; and Williams, B. 2006. A probabilistic approach to optimal robust path planning with obstacles. In *American Control Conference, 2006*, 7–pp. IEEE.

Clohessy, H., and Wiltshire, S. 1960. A terminal guidance system for satellite rendezvous, aero. *Astro Science* 27:563.

Eyerich, P.; Mattmüller, R.; and Röger, G. 2009. Using the context-enhanced additive heuristic for temporal and numeric planning. In *Nineteenth International Conference on Automated Planning and Scheduling*.

Fehse, W. 2003. *Automated rendezvous and docking of spacecraft*. Cambridge and and New York: Cambridge Univeristy Press.

Fox, M., and Long, D. 2003. Pddl2.1: An extension to pddl for expressing temporal planning domains. *J. Artif. Intell. Res. (JAIR)* 20:61–124.

Gelb, A. 1974. *Applied optimal estimation*. MIT press.

Hill, G. W. 1878. Researches in the lunar theory. *American Journal of Mathematics* 1(1):5–26.

Kalman, R. E. 1960. A new approach to linear filtering and prediction problems. *Journal of basic Engineering* 82(1):35–45.

Kamen, E. W., and Su, J. 1999. *Introduction to optimal estimation*. Springer.

Liberzon, D. 2003. *Switching in systems and control*. Springer.

Löhr, J.; Eyerich, P.; Keller, T.; and Nebel, B. 2012. A planning based framework for controlling hybrid systems. In *Proceedings of the International Conference on Automated Planning and Scheduling (ICAPS)*.

Löhr, J.; Nebel, B.; and Winkler, S. 2012. Planning-based autonomous lander control. In *AIAA/AAS Astrodynamics Specialist Conference*, Guidance, Navigation, and Control and Co-located Conferences. American Institute of Aeronautics and Astronautics.