7th IEEE Conference on Control Technology and Applications (CCTA) 2023

Tailoring stochastic predictive schemes for on-orbit control of space platforms

Martina Mammarella[1], Fabrizio Dabbene[1]

[1] Institute of Electronics, Computer and Telecommunication Engineering, National Research Council of Italy, Torino, Italy

Workshop on *Controls for aerospace applications: current solutions to future challenges*
August 15th 2023, Bridgetown, Barbados

Outline

Model Predictive Control

MPC is like playing CHESS

- The choice of a move (**control action**) is realized by projecting in the future the game scenery (**dynamical process model**) and trying to predict how the opponent will answer to our moves (**output**).
- If in the next move the opponent answers in an unexpected way (**measurements**), we need to re-plan our move again in order to counteract the effect of the opponent move (**feedback**).

Classic MPC

- Let us consider a discrete-time, linear system

$$x_{k+1} = Ax_k + Bu_k,$$

with state and input constrained into convex sets, i.e. $x_k \in \mathbb{X}$ and $u_k \in \mathbb{U}$, such that

$$\mathbb{X} = \{x \in \mathbb{R}^n \,|\, H_x x \leq h_x\}, \mathbb{U} = \{u \in \mathbb{R}^m \,|\, H_u x \leq h_u\}.$$

- As standard in MPC, we *repeatedly* solve a finite-horizon, receding optimal control problem

$$\begin{aligned}
&\min_{\mathbf{u}_k} \; J_T(x_k, \mathbf{u}_k) \\
&s.t. \; x_{\ell+1|k} = Ax_{\ell|k} + Bu_{\ell|k}, \quad x_{0|k} = x_k, \quad && \ell \in [0, T-1] \\
&\qquad x_{\ell|k} \in \mathbb{X}, \; u_{\ell|k} \in \mathbb{U}, && \ell \in [0, T-1] \\
&\qquad x_{T|k} \in \mathbb{X}_T
\end{aligned}$$

obtaining $\mathbf{u}_k^* = [u_{0|k}^*, \ldots, u_{T-1|k}^*]$ but implementing only the *first* control action $u_{0|k}^*$.

Expectation vs reality

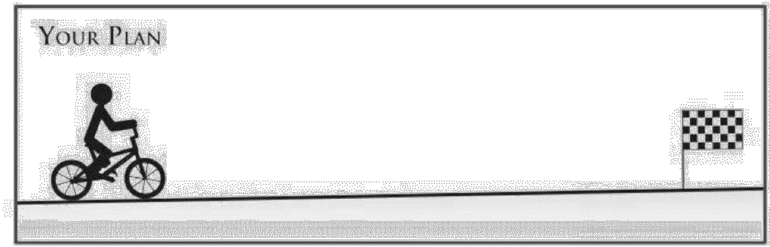

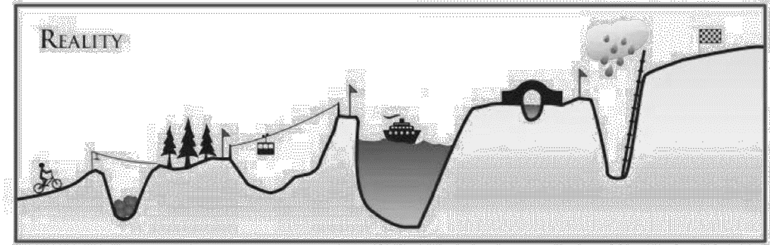

Uncertain world

- In real-world applications, the *complexity* of the phenomena and the *random* nature of data makes dealing with **uncertainty** essential.

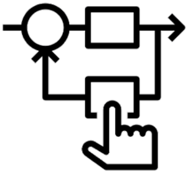

A. modeling

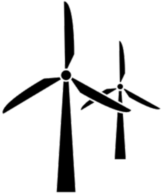

B. system

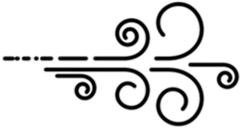

C. environment

- Coping in an *efficient* way with uncertainty is a **key** issue.
- Control engineers have started dealing *explicitly* with uncertainty.

Design under uncertainty

"Solum certum nihil esse certi"

"The one thing that is certain is that there is nothing certain"

Plinius the old

Naturalis Historiæ, 77 A.D.

Outline

From robust to stochastic

- The idea of **robust** control has been playing a fundamental role. The robust controller:

  1. guarantees performance satisfaction for *all* uncertainties,
  2. is inevitably *conservative*,
  3. provides a *pessimistic* view-point.

- In the **stochastic** framework, the goal is to find a solution that is *feasible* for *almost* all possible uncertainty realization.
- This translates in accepting some **risk** up to a certain safe level that the performance may be *violated*.

Optimization under uncertainty

- Consider an *uncertain* optimization problem

A. Robust optimization

$$\min_{\theta} c^\top \theta$$
$$\text{s.t. } f(\theta, w) \leq 0, \ \forall w \in \mathbb{W}$$

B. Chance-constrained optimization

$$\min_{\theta} c^\top \theta$$
$$\text{s.t. } \mathsf{Pr}_{\mathbb{W}}\{f(\theta, w) > 0\} \leq \varepsilon$$

- Constraints involving stochastic parameters that are required to be satisfied with a *pre-specified probability threshold* are called **chance constraints**.
- Efficiently solving or approximating CC problems remains an important task requiring a critical trade-off between complexity of the approach **VS** goodness of the approximation.
- To deal with chance constraints, we propose a **sample-based** approach.

Stochastic MPC

- Let us consider a discrete-time, linear system of the form

$$x_{k+1} = A(q_k)x_k + B(q_k)u_k + w_k,$$

  - $(w_k)_{k\in\mathbb{N}_{\geq 0}}$: realization of a stochastic process, where the disturbances are iid, zero-mean random variables with bounded and convex support $\mathbb{W}$;
  - $q_k \in \mathbb{Q} \subset \mathbb{R}^{n_q}$: realization of iid, multivariate real random variables Q_k.
- Let's assume that the system is subject to p_x individual chance constraints on the state x_k

$$\text{Pr}_{\mathbb{W}}\{x_k \in \mathbb{R}^n \,|\, H_{x,j}x_k \leq h_{x,j}\} \geq 1 - \varepsilon_j, \quad \forall k \in \mathbb{N}_{\geq 0},\ \forall j \in [1, p_x],$$

and m hard constraints on the input u_k

$$u_k \in \mathbb{U} = \{u_k \in \mathbb{R}^m \,|\, H_{x,j}x_k \leq h_{x,j}\}, \quad \forall k \in \mathbb{N}_{\geq 0}.$$

Stochastic MPC

- Let's define the following dual-mode, parametrized feedback policy

$$u_{\ell|k} = v_{\ell|k} + K x_{\ell|k},$$

with K quadratically stabilizing for the closed-loop system.

- The correction term sequence $\mathbf{v}_k\{v_{\ell|k}\}_{\ell=0}^{T}$ is the *minimizer* of the expected cost

$$J_T(x_k, \mathbf{v}_k) \doteq \mathbb{E}\left\{ \sum_{\ell=0}^{T-1} \left(\|x_{\ell|k}\|_Q^2 + \|u_{\ell|k}\|_R^2 \right) + \|x_{T|k}\|_P^2 \right\} = \mathsf{x}^\top \tilde{S} \mathsf{x}$$

subject to state chance constraints and input hard constraints.

- The expected values of the finite-horizon cost matrix $\tilde{S}$ can be evaluated **offline**, taking *random samples* of uncertainty and disturbance sequences.

Stochastic MPC

- Following the sampling-based approach in [1], we derive a convex, *inner* approximation $\tilde{\mathbb{X}}$ of the chance-constrained set in $x_k, \mathbf{v}_k, \mathbf{w}_k$, using results from SLT that define the *minimum* number N_ℓ of iid samples to draw to guarantee that, with probability high probability,

$$\tilde{\mathbb{X}} \subseteq \mathbb{X}_P, \quad \mathbb{X}_P \doteq \{x_k, | \mathbb{P}\{x_k \in \mathbb{X}\} \geq 1 - \varepsilon\}.$$

- To ensure *robust recursive feasibility*, we add a first-step constraint of the form

$$\mathbb{D}_R = \{x_k, \mathbf{v}_k \,|\, H_\infty (A + BK)^j x_{0|k} + H_\infty B^j v_{0|k} \leq h_\infty - H_\infty w_{0|k}\},$$

such that the final constraint set is defined as

$$\mathbb{D} \doteq \tilde{\mathbb{X}} \times \mathbb{U} \cap \mathbb{D}_R$$

[1] Lorenzen, M., Dabbene, F., Tempo, R., Allgöwer, F. (2017). Stochastic MPC with offline uncertainty sampling. Automatica, 81, 176-183.

Offline Sampling-based Stochastic MPC

OFFLINE STEP. *Before running the online control algorithm:*

1. *Compute the expected value of* $\tilde{S}$;
2. *Compute the inner approximation* $\tilde{\mathbb{X}}$ *of the chance-constrained set;*
3. *Compute* $\mathbb{D}$ *and remove redundant constraints.*

ONLINE IMPLEMENTATION. *At each time step* k:

1. *Measure the current state* x_k;
2. *Determine the minimizer of* J *subject to* $\tilde{\mathbb{X}}$, $\mathbb{U}$, *and* $\mathbb{D}_R$, *i.e.,*

$$\mathbf{v}_k^* = \arg \min_{\mathbf{v}_k} J(x_k, \mathbf{v}_k)$$

$$\text{s.t. } (x_k, \mathbf{v}_k) \in \mathbb{D};$$

3. *Apply the control input* $u_k = Kx_k + v^*_{0|k}$.

Outline

Earth-observation mission

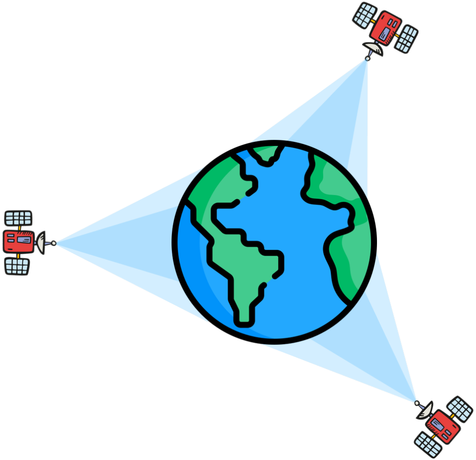

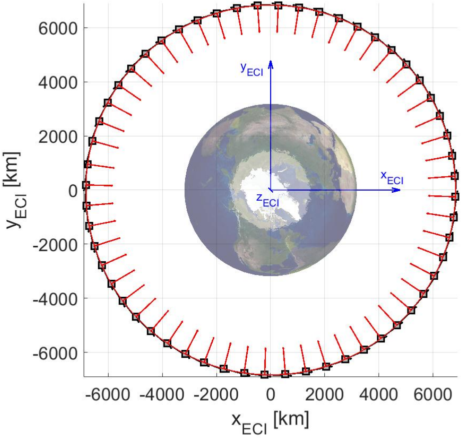

Earth-observation mission

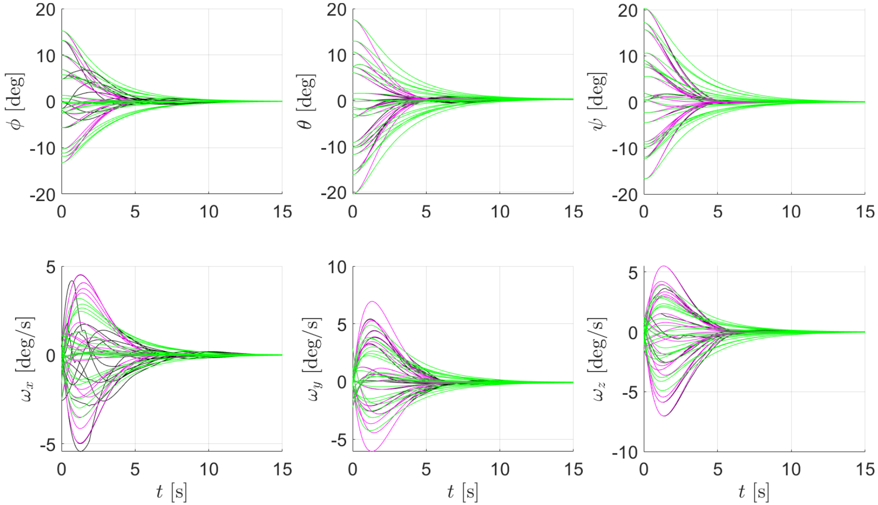

Earth-observation mission

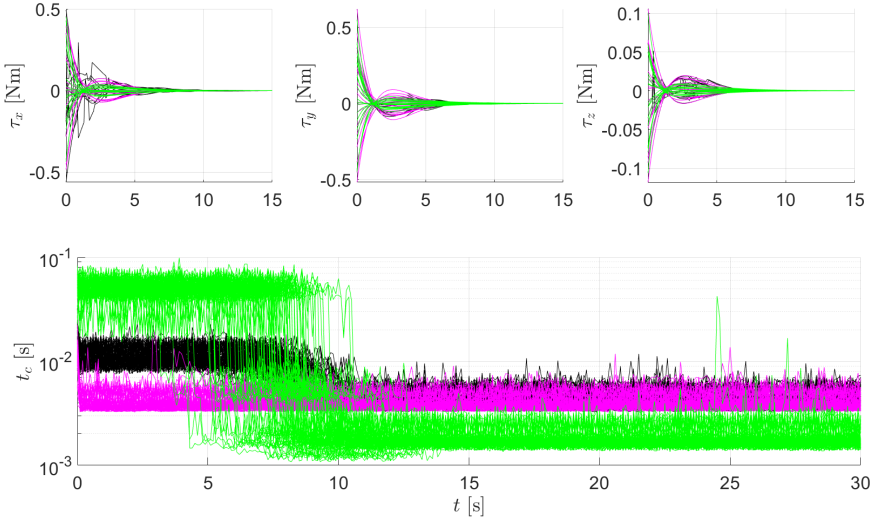

Earth-observation mission

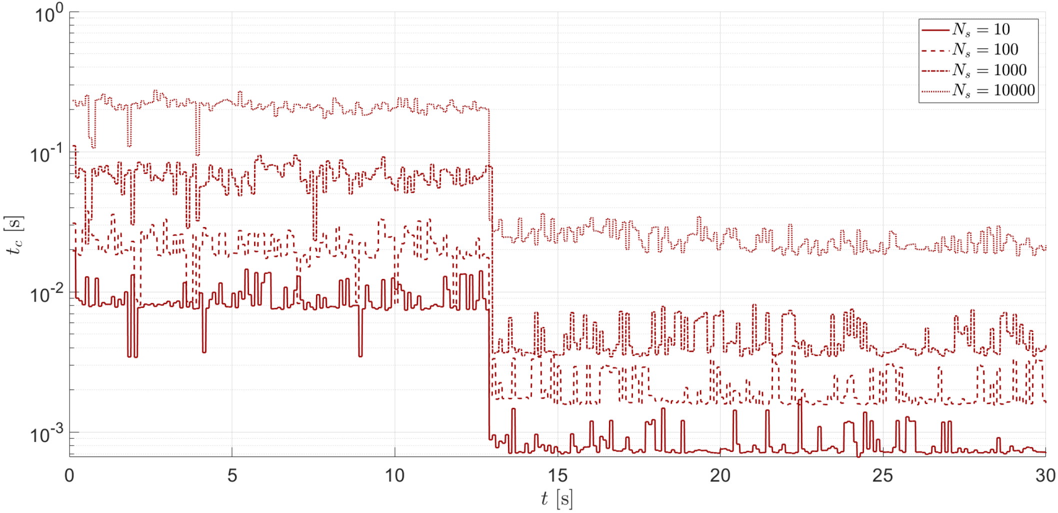

Outline

Autonomous rendezvous and docking

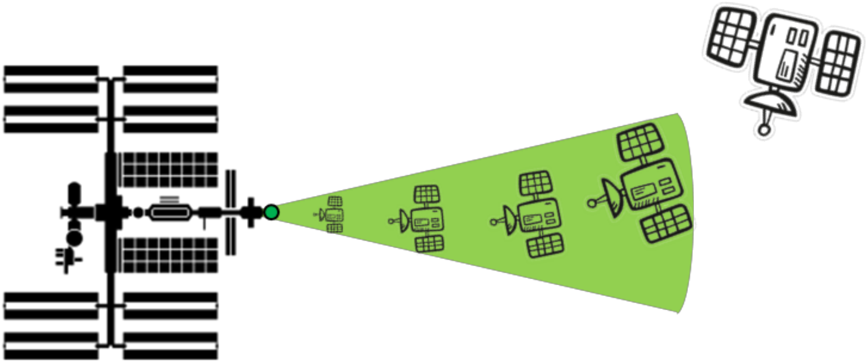

Experimental results @ NPS (Monterey, CA)

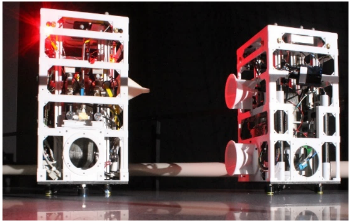

- mass: 9.882 ± 0.001 kg
- thrust: 0.15 ± 0.01 N
- OBC: Intel Atom 1.6 GHz 32bit
- memory: 2GB RAM, 8GB SSD
- OS: RTAI-patched ubuntu 14.04

Experimental results @ NPS (Monterey, CA)

- Uncertain dicrete-time FSS dynamics

$$x_{k+1} = \begin{bmatrix} q_1 & 0 & 1 & 0 \\ 0 & q_1 & 0 & 1 \\ 0 & 2q_2 & 0 & 0 \\ 0 & 3q_3 & -2q_2 & 0 \end{bmatrix} x_k + \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ \frac{1}{m} + q_4 & 0 \\ 0 & \frac{1}{m} + q_4 \end{bmatrix} u_k + w_k$$

- LOS approaching cone defined as a convex polytope with vertexes χ_1, χ_2, χ_3.

- Saturation constraints for thrusters:

$$u_k \in \mathbb{U} = \{\|u_k\|_\infty \leq 0.3\};$$

- Sample times:

  i) FSS: 0.01 s; ii) nav: 0.02 s; iii) MPC: 3-5 s.

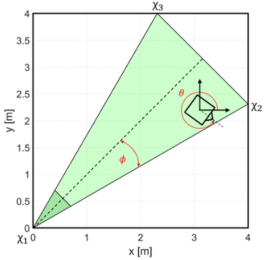

OS-SMPC: experimental validation

- 20 runs for each initial condition in terms of position and orientation;
- maneuver duration: 120-200 s (faster than classical/robust MPC);
- max computation time: 4 s (over 5s allocated for the control task);
- avg control effort: 4.69 Ns (lower than classic/robust MPC).

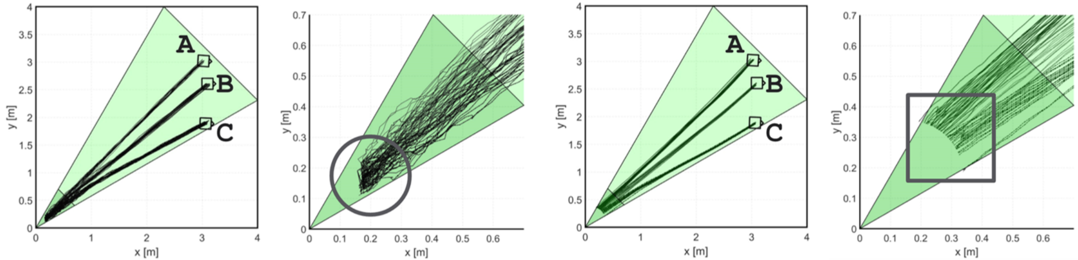

Outline

A new tool for design under uncertainty

Basic idea of *probabilistic scaling*.

Probabilistic scaling approach

- The probabilistic scaling approach leads to approximations of **user-chosen complexity**.
- We introduce **Scalable Simple Approximating Sets (SAS)**.

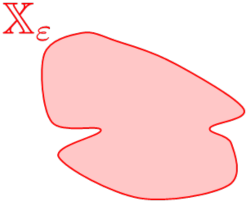

ε-CSS $\mathbb{X}_\varepsilon$

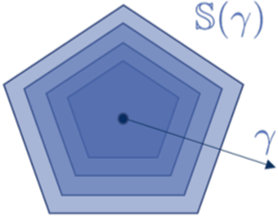

$\mathbb{S}(\gamma) = \theta_c \oplus \gamma\mathbb{S}$

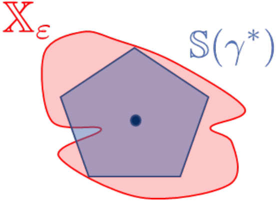

$\mathsf{Pr}_{\mathbb{W}}\{\mathbb{S}(\gamma^*) \subseteq \mathbb{X}_\varepsilon\} \geq 1 - \delta$

Probabilistic SAS scaling

Algorithm 1 Probabilistic SAS Scaling

Given a candidate Scalable SAS $\mathbb{S}(\gamma)$, and probability levels ε and δ, choose

$$N_\gamma \geq \frac{7.47}{\varepsilon} \ln \frac{1}{\delta} \quad \text{and} \quad r = \left\lceil \frac{\varepsilon N_\gamma}{2} \right\rceil .$$

Draw N_γ samples of the uncertainty $w^{(1)}, \ldots, w^{(N_\gamma)}$.

for $i \in \mathbb{N}_1^{N_\gamma}$ **do**

    Compute the N_γ scaling factors $\gamma_i \doteq \gamma(w^{(i)}),\ i \in \mathbb{N}_1^{N_\gamma}$

end for

Return $\bar{\gamma} = \gamma_r^- = \min^{(r)}(\Gamma)$, the r-th smallest value of $\Gamma = \{\gamma_i\}_{i=1}^{N_\gamma}$.

$$\Rightarrow \ \mathsf{Pr}_{\mathbb{W}}\{\mathbb{S}(\bar{\gamma}) \subseteq \mathbb{X}_\varepsilon\} \geq 1 - \delta.$$

Norm-based SAS

- We define norm-based SAS of the form

$$\mathbb{S}_{\ell_p}(\gamma) \doteq \theta_c \oplus \gamma H \mathbb{B}_p^s, \quad s \geq n_\theta,$$

where $\mathbb{B}_p^s$ is a ℓ_p-ball in $\mathbb{R}^s$ and $H \in \mathbb{R}^{n_\theta \times n_s}$ is the shape matrix.

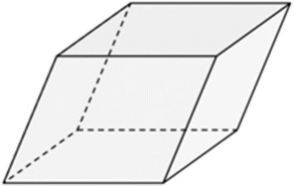

(a) $p = 1, n_\theta = s = 3$

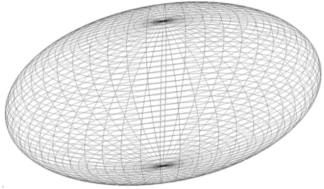

(b) $p = 2, n_\theta = s = 3$

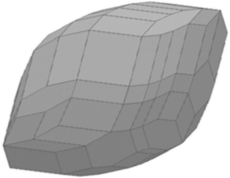

(c) $p = \infty, n_\theta = 3, s = 10$

Norm-based SAS

Theorem (Scaling factor for norm-based SAS)

Given a norm-based SAS $\mathbb{S}_{\ell_p}(\gamma) = \theta_c \oplus \gamma H \mathbb{B}_p^s$ *and a realization* $w \in \mathbb{W}$, *define* $\tau_\ell(w) \doteq g_\ell(w) - f_\ell^T(w)\theta_c$ *and* $\rho_\ell(w) \doteq \|H^T f_\ell(w)\|_{p^*}$. *The scaling factor* $\gamma(w)$ *can be computed as*

$$\gamma(w) = \min_{\ell \in \mathbb{N}_1^{n_\ell}} \gamma_\ell(w),$$

with $\gamma_\ell(w)$ *given by*

$$\gamma_\ell(w) = \begin{cases} 0 & \textit{if} \quad \tau_\ell(w) < 0, \ \textit{i.e.}\ \theta_c \notin \mathbb{X}(w) \\ \infty & \textit{if} \quad \tau_\ell(w) \geq 0, \rho_\ell(w) = 0, \\ \frac{\tau_\ell(w)}{\rho_\ell(w)} & \textit{if} \quad \tau_\ell(w) \geq 0, \rho_\ell(w) > 0. \end{cases}$$

Norm-based SAS

- First, we draw a fixed number N_S of design uncertainty samples $\{w^{(1)},\dots,w^{(N_S)}\}$ and we build $\mathbb{X}_{N_S} = \bigcap_{j=1}^{N_S} \mathbb{X}(w^{(i)})$.
- Then, we compute the largest (volume) set $\theta_c \oplus H\mathbb{B}_p^s$ included in $\mathbb{X}_{N_S}$, solving the following (relaxed) optimization problem

$$\min_{\theta_c, H, \eta_1,\dots,\eta_{N_S}} \quad -\mathrm{Vol}_p(H) + \xi \sum_{j=1}^{N_S} \max\{\eta_j, 0\}$$

$$\text{s.t.} \quad f_\ell^\top(w^{(j)})\theta_c + \|H^T f_\ell(w^{(j)})\|_{p^*} - g_\ell(w^{(j)}) \le 0, \quad \ell \in \mathbb{N}_1^{n_\ell},\ j \in \mathbb{N}_1^{N_S}.$$

Norm-based SAS - Example

$\mathbb{S}_{\ell_1}$-SAS, $N_S = 100$

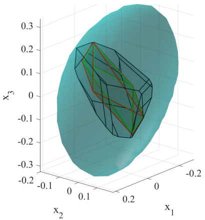

(a) $\gamma = 0.9701$

$\mathbb{S}_{\ell_1}$-SAS, $N_S = 1000$

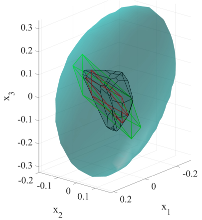

(b) $\gamma = 1.5995$

$\mathbb{S}_{\ell_\infty}$-SAS, $N_S = 100$

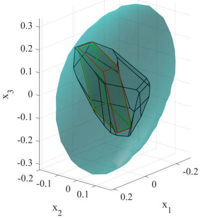

(c) $\gamma = 0.9696$

$\mathbb{S}_{\ell_\infty}$-SAS, $N_S = 1000$

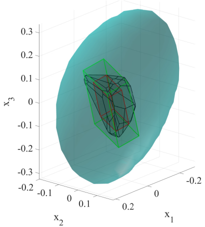

(d) $\gamma = 1.5736$

n.	$t_{c_{MAX_{OS}}}$	$t_{c_{MAX_{PS}}}$	$t_{c_{AVG_{OS}}}$	$t_{c_{AVG_{PS}}}$
1	2.0959	0.4178 (-80%)	0.0966	0.0087 (-91%)
2	2.9411	0.5626 (-81%)	0.7221	0.0190 (-97%)
3	2.1497	0.5434 (-75%)	0.2628	0.0086 (-96%)

Outline

Towards debris removal applications

- *Active debris removal* as mitigation strategy for non-functional satellites, abandoned vehicle stages, fragments, etc.
- Post-capture, *detumble and re-orient* the combined system.
- *Controllability* issues if ADR vehicle and debris have comparable inertia.
- *Model uncertainties* from partial/missing knowledge of debris mass/inertia.

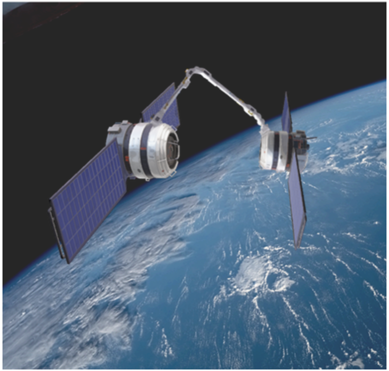

Disturbance probabilistic bounds

- Let's consider a discrete-time, linear system of the form

$$x_{k+1} = Ax_k + Bu_k + \zeta_k,$$

where *no assumption on neither independence nor Gaussianity* are made on the stochastic disturbance ζ_k, and with both state and input constrained in compact sets $\mathbb{X}$ and $\mathbb{U}$, defined in a mixed form as

$$Cx_k + Du_k \leq h, \quad \forall k \in \mathbb{N}_{\geq 0}.$$

GOAL: to achieve closed-loop guarantees by means of probabilistic validation techniques.

- The proposed approach is defined at two different levels:
  - a. *constraint tightening* using a sampling method to (upper) bound the effect of disturbances.
  - b. (offline) probabilistic design of *penalty function* to guarantee recursive feasibility.

Disturbance probabilistic bounds

- Being the disturbance persistent, the state $x_{\ell|k}$ can be split into

$$x_{\ell|k} = z_{\ell|k} + e_{\ell|k},$$

and selecting a parametrized feedback policy of the form $u_{\ell|k} = v_{\ell|k} + Kx_{\ell|k}$, we have

$$(C + DK)z_{\ell|k} + Dv_{\ell|k} \leq h - (C + DK)e_{\ell|k}.$$

- Since $\{e_{\ell|k}\}_{\ell=0}^{T-1}$ is completely determined by the disturbance sequence, we propose a sample-based method to obtain a **probabilistic upper bound** on $(C + DK)e_{\ell|k}$.

Theorem

Given $\varepsilon_q \in (0,1)$ *and* $\delta_q \in (0,1)$ *and a discarding parameter* r_q, *suppose that* S_q *i.i.d. samples* $\{\zeta^{(1)}, \dots, \zeta^{(S_q)}\}$ *are drawn according to* $\mathbb{P}_{\mathbb{D}}$, *and consequently we compute*

$$q_{\ell,j} = \left[\{C_{K,j} e_{\ell|0}(\zeta^{(i)})\}_{i=1}^{S_q} \right]_{r_q}, \forall \ell \in \mathbb{N}_0^{T-1}, \forall j \in \mathbb{N}_1^{n_h}.$$

Then, if $S_q \geq \frac{1}{\varepsilon_q}\left(r_q - 1 + \ln \frac{n_h T}{\delta_q} + \sqrt{2(r_q - 1)\ln \frac{n_h T}{\delta_q}} \right)$, *we have that, with probability no smaller than* $1 - \delta_q$,

$$\mathbb{P}_{\mathbb{D}}\{C_{K,j} e_{\ell|k}(\zeta) > q_{\ell,j}\} \leq \varepsilon_q, \forall \ell \in \mathbb{N}_0^{T-1}, \forall j \in \mathbb{N}_1^{n_h},$$

provided that $r_q \in \mathbb{N}_1^{S_q}$ *and* $\sum_{m=0}^{r_q-1} \left(\begin{array}{c} S_q \\ m \end{array} \right) \varepsilon_q^m (1-\varepsilon_q)^{Sq-m} \leq \frac{\delta_q}{n_h T}$.

Probabilistic bounds and penalty function

Tightened constraints

$$\begin{aligned}\min_{\mathbf{z}_k,\mathbf{v}_k} \quad & J(\mathbf{z}_k,\mathbf{v}_k)\\ s.t. \quad & z_{0|k} = x_k\\ & z_{\ell+1|k} = A_K z_{\ell|k} + B v_{\ell|k},\ \forall \ell \in \mathbb{N}_0^{T-2}\\ & z_{T-1|k} = A_K z_{T-1|k} + B v_{T-1|k}\\ & C_K z_{\ell|k} + D v_{\ell|k} \preceq h - q_\ell,\ \forall \ell \in \mathbb{N}_0^{T-1}\end{aligned}$$

Penalty function

$$\begin{aligned}\min_{\mathbf{z}_k,\mathbf{v}_k} \quad & J(\mathbf{z}_k,\mathbf{v}_k) + \rho \sum_{\ell=0}^{T-1} \langle C_K z_{\ell|k} + D v_{\ell|k} - h + q_\ell \rangle_+\\ s.t. \quad & z_{0|k} = x_k\\ & z_{\ell+1|k} = A_K z_{\ell|k} + B v_{\ell|k},\ \forall \ell \in \mathbb{N}_0^{T-2},\\ & z_{T-1|k} = A_K z_{T-1|k} + B v_{T-1|k}\\ & u_{0|k} = v_{0|k} + K x_{0|k} \in \mathbb{U}.\end{aligned}$$

- To design the penalty parameter ρ, we use a sample-based design technique to ensure given probabilistic guarantees over constraint satisfaction, obtaining a probabilistic upper bound on a *performance index*

$$g(w,\rho) \doteq \sum_{k=0}^{M} \langle C x_k(w,\rho) + D u_k(w,\rho) - h \rangle_+ \ \rightarrow\ g(w,\rho) = 0 \Leftrightarrow C x_k(w,\rho) + D u_k(w,\rho) \preceq h,\ \forall k \in \mathbb{N}_0^M,$$

for a finite-family of possible values of ρ.

Disturbance forecasting extension

- Let's consider again a discrete-time, linear system of the form

$$x_{k+1} = Ax_k + Bu_k + d_k, \quad d_k = \hat{d}_k + \zeta_k,$$

  affected by persistent, additive disturbance d_k, which can be split into two contributions with $\hat{d}_k$ the *estimated additive disturbance*, modelled from data using **disturbance forecasting techniques**, and w_k the unmodelled disturbance.
- In this way, we can *estimate* the contribution of modelled terms on one side and *probabilistically upper bound* the unmodelled one.
- For estimating the predicted evolution of the disturbance over the prediction horizon starting from past data, one can use different methods, e.g., *black-box AutoRegressive eXogenous* (ARX) or *Controlled Auto-Regressive and Integrated Moving-Average* (CARIMA).

LAST BUT NOT LEAST...

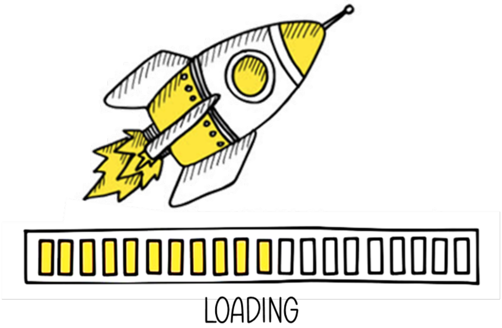


Buzz
aka VZLUSAT2
Woody
aka OS-SMPC
CZECH
AEROSPACE
RESEARCH
CENTRE

Tailoring stochastic predictive schemes for on-orbit control of space platforms

Thank you for your attention

Q&A

martina.mammarella@ieiit.cnr.it