Quantitative Resilience of Driftless Linear Systems

Jean-Baptiste Bouvier, Kathleen Xu and Melkior Ornik

Department of Aerospace Engineering and Coordinated Science Laboratory,
University of Illinois Urbana-Champaign, USA
Email: bouvier3@illinois.edu

SIAM Conference on Control and Its Applications, July 2021

This work was supported by an Early Stage Innovations grant from NASA's Space Technology Research Grants Program, grant no. 80NSSC19K0209, and by the United States Air Force AFRL/SBRK under contract no. FA864921P0123.

Motivating Example: Opinion Dynamics

The opinion x of n agents without interactions, but watching $m+p$ TV channels is modeled by $\dot{x}(t) = \bar{B}\bar{u}(t)$. Matrix $\bar{B} \in \mathbb{R}^{n\times(m+p)}$ represents the trust of the agents with respect to the TV channels.

The controller, e.g., a world-wide media conglomerate, uses its channels inputs $\bar{u} \in \bar{U}$ to steer the opinion x of agents towards a target x_{goal}.

After a loss of control authority over p channels, they produce uncontrolled and undesirable inputs $w \in W$. The remaining controls are $u \in U$ and we split $\bar{B} = \big[B\ C\big]$, so that

$$\dot{x}(t) = Bu(t) + Cw(t) \qquad x(0) = x_0.$$

Motivating Example: Opinion Dynamics

Consider $n = 3$ agents watching 6 TV channels with

$$\bar{B} = \begin{bmatrix} 0.8 & -0.9 & 0.4 & -0.4 & -0.7 & 0 \\ 1 & -1 & 0.3 & 0.2 & 0.7 & -0.1 \\ 0.9 & -0.8 & -0.4 & -0.4 & 0.7 & 0.1 \end{bmatrix} \quad \text{and} \quad x_0 = \begin{bmatrix} 0 \\ 0 \\ 0 \end{bmatrix}.$$

The agents strongly trust channel 1 but not channel 2, they have diverging moderate opinions on channels 3 and 4, strongly diverging trust of channel 5, while they are barely influenced by channel 6.

Can any target still be reached after the loss of control authority over any single channel? How is the performance affected?

Problem Statement

Assumption: Thanks to sensors, the undesirable inputs w are measured in real-time, so the control law u can depend on their current value.

Definition 1

The system is resilient *to the loss of p of its actuators (matrix C) if for all undesirable input $w \in W$, there is a control law $u_w \in U$ and a time T such that $x(T) = x_{goal}$ for $\dot{x}(t) = Bu_w\big(t, w(t)\big) + Cw(t)$.*

Problem: How to quantify the resilience of our system?

Definition of Quantitative Resilience

The *nominal reach time* T_N^* is the shortest time for the nominal system to reach x_{goal}: $T_N^*(d) := \inf\limits_{\bar{u} \in \bar{U}} \Big\{ T \geq 0 : \int_0^T \bar{B}\bar{u}(t)\, dt = d \Big\}$, $d = x_{goal} - x_0$.

The *malfunctioning reach time* T_M^* is the shortest time for the malfunctioning system to reach x_{goal} for the worst undesirable input:

$$T_M^*(d) := \sup_{w \in W} \Big\{ \inf_{u \in U} \Big\{ T \geq 0 : \int_0^T Bu(t) + Cw(t)\, dt = d \Big\}\Big\}.$$

The *ratio of reach times* in the direction $d \in \mathbb{R}^n$ is $t(d) := \frac{T_M^*(d)}{T_N^*(d)}$.

The *quantitative resilience* of a system is $r_q := \frac{1}{\sup\limits_{d \in \mathbb{R}^n} t(d)} = \inf\limits_{d \in \mathbb{R}^n} \frac{T_N^*(d)}{T_M^*(d)}$.

Problem: How to compute r_q efficiently?

Limitations of other Approaches

- *Fault-tolerant control* typically studies either actuators locking in place, losing effectiveness but remaining controllable, or a mix of both.
- *Robust control* only works for small undesirable inputs, and is too conservative because it has no real-time reading of the undesirable inputs.
- *Quantitative resilience* has only been developed for specific applications like water infrastructures or nuclear power plants.

Framework

Reminder: the nominal system is $\dot{x}(t) = \bar{B}\bar{u}(t)$ with $\bar{u} \in \bar{U}$. After the loss of control over p actuators, the dynamics are $\dot{x}(t) = Bu(t) + Cw(t)$ with $u \in U$ the controls, and $w \in W$ the undesirable inputs.

The components of each input have the same bound u_{max} at all times:

$$\begin{aligned} \bar{U} &:= \big\{\bar{u} : \mathbb{R}^+ \to \mathbb{R}^{m+p},\ \bar{u}(t) \in \bar{U}_c\big\} \qquad & \bar{U}_c &:= \big\{\bar{u} \in \mathbb{R}^{m+p},\ |\bar{u}_i| \leq u_{max}\big\}, \\ U &:= \big\{u : \mathbb{R}^+ \to \mathbb{R}^{m},\ u(t) \in U_c\big\} & U_c &:= \big\{u \in \mathbb{R}^{m},\ |u_i| \leq u_{max}\big\}, \\ W &:= \big\{w : \mathbb{R}^+ \to \mathbb{R}^{p},\ w(t) \in W_c\big\} & W_c &:= \big\{w \in \mathbb{R}^{p},\ |w_j| \leq u_{max}\big\}. \end{aligned}$$

The set of vertices of W_c is denoted as V_c.

The unit circle in $\mathbb{R}^n$ is denoted by $\mathbb{S} := \big\{x \in \mathbb{R}^n,\ \|x\| = 1\big\}$.

Dynamics of the Nominal System

Reminder: $$T_N^*(d) = \inf_{\bar{u} \in \bar{U}} \Big\{ T \geq 0 : \int_0^T \bar{B}\bar{u}(t)\, dt = d \Big\}.$$

Proposition 1 (Constant optimal control)

If the system $\dot{x}(t) = \bar{B}\bar{u}(t)$ is controllable and $d = x_{goal} - x_0 \in \mathbb{R}^n$, the infimum $T_N^*(d)$ is achieved with a constant control input $\bar{u}^* \in \bar{U}_c$.

$$\frac{1}{T_N^*(d)} = \max_{\bar{u} \in \bar{U}_c} \big\{ \lambda : \bar{B}\bar{u} = \lambda d \big\}.$$

Proposition 2 (Absolute homogeneity)

The nominal reach time verifies $T_N^*(\lambda d) = |\lambda| T_N^*(d)$ for $d \in \mathbb{R}^n$, $\lambda \in \mathbb{R}$.

Dynamics of the Malfunctioning System

Reminder: $$T_M^*(d) = \sup_{w \in W} \Big\{ \inf_{u \in U} \Big\{ T \geq 0 : \int_0^T Bu(t) + Cw(t)\, dt = d \Big\} \Big\}.$$

Proposition 3 (Constant optimal inputs)

If the system $\dot{x}(t) = Bu(t) + Cw(t)$ is resilient and $d = x_{goal} - x_0$, the supremum and infimum of $T_M^*(d)$ are both achieved with constant inputs in V_c and U_c respectively.

$$\frac{1}{T_M^*(d)} = \min_{w \in V_c} \Big\{ \max_{u \in U_c} \big\{ \lambda : Bu + Cw = \lambda d \big\} \Big\}.$$

Proposition 4 (Absolute homogeneity)

The malfunctioning reach time also verifies $T_M^*(\lambda d) = |\lambda| T_M^*(d)$ for $d \in \mathbb{R}^n$, $\lambda \in \mathbb{R}$.

Application to Opinion Dynamics

Reminder:

$$\bar{B} = \begin{bmatrix} 0.8 & -0.9 & 0.4 & -0.4 & -0.7 & 0 \\ 1 & -1 & 0.3 & 0.2 & 0.7 & -0.1 \\ 0.9 & -0.8 & -0.4 & -0.4 & 0.7 & 0.1 \end{bmatrix}, \quad t(d) = \frac{T_M^*(d)}{T_N^*(d)}.$$

For $x_{goal} = d = (1,\ 1,\ 1)$, we have $t(d) = \begin{bmatrix} 4.8 & 5.4 & 1.6 & 1.6 & \infty & 1.0 \end{bmatrix}$.

For $d = (-1,\ -1,\ 1)$, we have $t(d) = \begin{bmatrix} 1.7 & 1.7 & 6.5 & 5.4 & \infty & 1.3 \end{bmatrix}$.

Quantitative Resilience

With the absolute homogeneity of T_N^* and T_M^*, $\max\limits_{d \in \mathbb{R}^n} t(d) = \max\limits_{d \in \mathbb{S}} t(d)$.

Theorem 1 (Direction Maximizing t)

For a resilient system with C a single column matrix, $\max\limits_{d \in \mathbb{S}} t(d) = t(C)$.

Let $X := \big\{Cw_c : w_c \in W_c\big\}$ and $Y := \big\{Bu_c : u_c \in U_c\big\}$.

Theorem 2 (Maximax Minimax Quotient Theorem)

If X and Y are two symmetric polytopes in $\mathbb{R}^n$ with $X \subset Y^\circ$, $\partial X = \{-x, x\}$ and $\dim Y = n$, then $\max\limits_{d \in \mathbb{S}} r_{X,Y}(d) = r_{X,Y}(x)$, *with*

$$r_{X,Y}(d) := \frac{\max\limits_{x \in X,\, y \in Y} \big\{\|x+y\| : x+y \in \mathbb{R}^+ d\big\}}{\min\limits_{x \in X} \big\{\max\limits_{y \in Y}\{\|x+y\| : x+y \in \mathbb{R}^+ d\}\big\}}.$$

Illustration of the Maximax Minimax Quotient Theorem

Reminder: $X = \{Cw_c : w_c \in W_c\}$ and $Y = \{Bu_c : u_c \in U_c\}$.

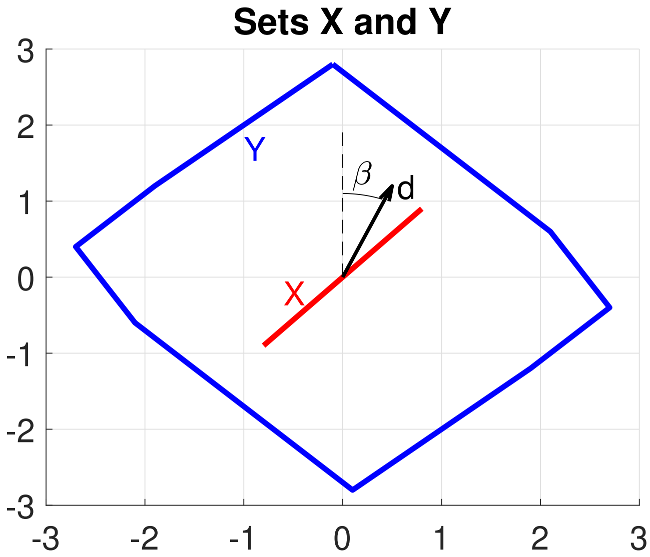

Figure 1: Polytopes X and Y in $\mathbb{R}^2$.

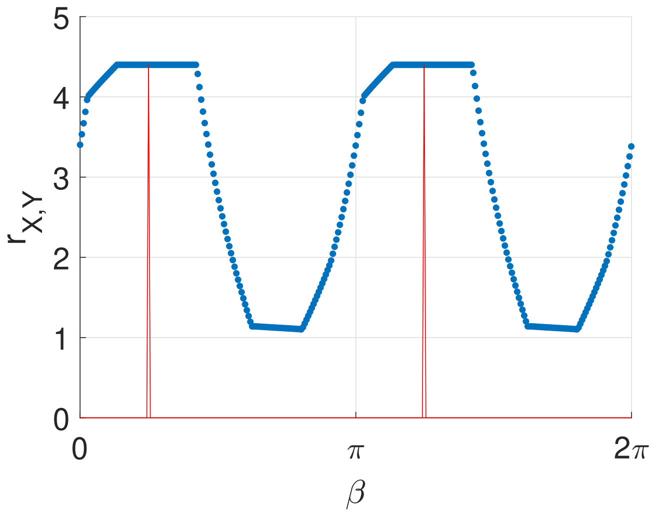

Figure 2: Ratio $r_{X,Y}\big(d(\beta)\big)$ with $\beta \in [0, 2\pi]$.

Quantitative Resilience

Theorem 3

For a resilient system with C a single column matrix, $r_q = r_{max}$ where

$$r_{max} := \frac{\lambda^* - u_{max}}{\lambda^* + u_{max}} \quad \textit{and} \quad \lambda^* = \max_{u_c \in U_c} \{\lambda : Bu_c = \lambda C\}.$$

Now we only need to solve a single linear optimization problem to determine r_q, instead of the initial four nested nonlinear problems.

Proposition 5 (Necessary and Sufficient condition for Resilience)

A system is resilient to the loss of a column C if and only if it is controllable and $r_{max} \in (0, 1]$.

Application to Opinion Dynamics

We compute r_{max} for the loss of control over each single channel:

$$r_{max} = \begin{bmatrix} 0.20 & 0.18 & 0.15 & 0.19 & -0.22 & 0.75 \end{bmatrix}.$$

Since $\bar{B}$ is controllable, the system is resilient to the loss of each single channel except channel 5. Then,

$$\frac{1}{r_q} = \begin{bmatrix} 4.9 & 5.5 & 6.5 & 5.4 & \infty & 1.3 \end{bmatrix}.$$

Thus, after the loss of control over channel 3, the time to reach some targets can be multiplied by up to 6.5 compared to the nominal reach time.

Conclusion

We introduced quantitative resilience for control systems losing control authority over some of their actuators and we established an efficient algorithm to verify the resilience and quantitative resilience of driftless systems.

Future work:

- Non-driftless systems.
- Asymmetric input sets.
- Simultaneous loss of several actuators.