2023 62nd IEEE Conference on Decision and Control (CDC)
December 13-15, 2023. Marina Bay Sands, Singapore

Feasibility detection for nested codesign of hypersonic vehicles

Chris van der Heide Peter Cudmore Ingo Jahn Viv Bone Peter M. Dower Chris Manzie

***Abstract*— Controllability and feasiblity measures are used to determine whether a given system can achieve its specified objective. However, for nonlinear systems with state constraints, the controllable and feasible sets may be highly sensitive to minor perturbations in the system's constraints, initial states and parameters. This becomes particularly important in codesign of hypersonic vehicles, where functions governing the dynamics must be estimated from expensive computational fluid dynamics simulations, and poor initialization can lead to significant waste of resources. By relaxation of the constraints and introduction of a surrogate cost, we provide a method for detecting and quantifying which constraints are violated. To demonstrate the method in a concrete example, we apply the technique to simulation of hypersonic vehicle trajectories.**

I. INTRODUCTION

Codesign is a multidisciplinary optimization approach to design of parameterized dynamical systems with control inputs [2], [10], [1]. The two most successful approaches are simultaneous and nested optimization strategies [12], [23]. In the simultaneous approach, an objective function is jointly optimized in both the plant parameters and control inputs. While this approach is effective and efficient in some situations, in many settings it is more effective to exploit more structure in the the problem, as is the case when high-fidelity simulations (typically via computational fluid dynamics or finite element analysis) are required in the optimization loop. In contrast, the nested approach decomposes the problem as a bi-level optimization problem, where the inner optimization loop computes the optimal control inputs for fixed plant parameters, and the outer loop optimizes the plant parameters. However, for tightly constrained problems, a poor choice of the initial design parameters or a parameter update in the optimization step can preclude the feasibility of the inner-loop control problem. To facilitate routine application of nested codesign, we require algorithms for efficiently detecting feasibility of the inner-loop control problem.

First achieved in the 1950s in the X-15 program, hypersonic flight represents a significant human technological achievement and continues to be an active area of aerospace research [9]. Hypersonic vehicles (HSVs) fly aggressive trajectories in extreme environments and are thus required to operate close to many system constraints, such as thermal load limitations, maximum dynamic pressure limits, actuator limits, and aerodynamic survivability constraints. Moreover, safety, reliability, reproducibility and economic factors are all paramount to ensure feasibility of commercial and military applications. Each of these factors are important and they often compete with one another; for example, a lighter airframe may be capable of flying more complex trajectories, but may be more vulnerable to maximum dynamic pressure limits. Due to the presence of interacting constraints and competing design objectives, codesign approaches are a powerful tool for improving the performance of hypersonic vehicles. However, codesign implementations typically rely on ad-hoc methods, and tools for reliable black-box algorighms for end-to-end design of these systems are still being developed.

Even for a fixed vehicle design, control system design for airbreathing HSVs is a challenging problem due to the nonlinearity of the vehicle's dynamics and the presence of interacting state and input constraints. Moreover, the dynamics themselves often have limited control authority, which renders feasibility of the control problem particularly sensitive to minor changes in the aircraft and mission design. Further, the difference in the magnitudes and time-scales of different variables results in trajectories that are highly sensitive to minor perturbations of the optimal control. As such, judicious choice of feedback controller becomes crucial, since the computational intensity required can be prohibitive for online control [17]. Similarly, approaches that directly incentivize some level of robustness to perturbations of the optimal trajectory have shown promise [22].

The characteristics of the HSV design problem impose further challenges on the codesign procedure [13], [16]. To enable good performance, vehicle geometries are often highly parameterized, resulting in high-dimensional outer-loop optimization problems. For each set of parameter values, we must generate the aerodynamic coefficient (such as lift, drag, and pitching moment) functions that appear in the vehicle dynamics. When high accuracy is required, these functions are typically modeled by approximating results sampled from computational fluid dynamics (CFD) simulations, and using lower resolution simulations can have profound effects on the accuracy of the optimal trajectories [8], [19]. Time-scale separation means that steady-state CFD simulations provide good approximations to dynamic behaviour, but even these simpler simulations still impose a formidable computational bottleneck on the codesign procedure [14], [15]. Thus, to make codesign tractable, we require an efficient method that can (i) iterate the initial design parameters towards a region where the inner-loop control problem is feasible and (ii) during the codesign problem, ensure feasibility of the inner-

CH, PC, VB, PD and CM are with the Department of Electrical and Electronic Engineering, The University of Melbourne, VIC, 3010 Australia. IJ is with the School of Engineering, The University of Southern Queensland, Springfield, QLD, 4300 Australia. Email: chris.vanderheide@.unimelb.edu.au; peter.cudmore@uqconnect.edu.au; ingo.jahn@usq.edu.au; viv.bone@unimelb.edu.au; pdower@unimelb.edu.au; manziec@unimelb.edu.au.

loop when iterating already-feasible values of the design parameters.

In order to facilitate tractable black-box codesign for HSVs, we propose a procedure with an additional surrogate inner-loop optimization problem that quantifies infeasibility by measuring constraint violations along the 'most feasible' trajectory. When the problem is feasible — i.e., the design parameters admit an optimal control trajectory that respects all constraints — the codesign problem proceeds as usual. Otherwise, we use sensitivity tools in the NLP solver to generate gradient information that can then be used to iterate the design parameters back towards a region where the inner-loop is feasible. This process significantly reduces the computational cost of the entire codesign problem. The feasibility detection step adds only small computational cost itself, but avoids the need for unnecessary design parameter iterations, which would require expensive CFD simulations.

The layout of the paper is as follows. In Section II we begin with a review of the codesign procedure and introduce a simplified version of the HSV's dynamics that retain the pathologies characteristic of more sophisticated models. We then extend the codesign framework to incorporate feasibility of the inner-loop optimal control problem, and discuss evaluation of the feasibility sub-problem. Numerical simulations using the HSV example are given in Section III. Finally, we provide discussion and conclude in Section IV.

II. CODESIGN FOR HYPERSONIC VEHICLES

The nested optimal codesign framework is a decomposition-based approach to design of engineered systems. In its simplest formulation in the current context, it casts the design problem as a bi-level optimization with a single objective function, where the outer-loop optimizes the design parameters θ of the physical system and the inner-loop computes the optimal control inputs $u(\cdot)$ and state trajectories $x(\cdot)$. A simple version of the codesign procedure is as follows: for fixed values of the design parameters, we first solve an optimal control problem to compute the input trajectory that minimizes the objective function. Next, we update the parameters. While heuristic methods can be deployed, we consider the commonly used gradient-informed approach. This is done by differentiating the objective function J^{sys} evaluated with the optimal control, with respect to the design parameters, and then updating the parameters, typically via gradient descent. This process is repeated until some convergence criteria is met and a final (locally optimal) pair of design/control parameters is selected. Of course, this description is heavily simplified; more sophisticated systems require, for example, choice of an inner-loop feedback controller, closed-loop simulations with disturbances, and state estimation.

In general, both the inner-loop and outer-loop optimization problems in the codesign formulation are highly non-convex and are not guaranteed to converge to a globally optimal design. The dynamics themselves are typically highly non-linear and hard constraints on both the control inputs and states are present. Moreover, for certain values of the design parameters, the inner-loop control problem may be infeasible. Due to the inherent nonlinearities in the aerodynamics models, HSV codesign is affected by these non-convexity and infeasibility problems, even when considering simplified vehicle dynamics and only a small number of design parameters.

A schematic of a codesign framework for HSVs appears in Figure I.

A. Hypersonic Vehicle Dynamics

The most commonly used control oriented model of an airbreathing HSV contains planar rigid body aircraft dynamics, as given by Parker et al [21]. For clarity we consider a trimmed version of this model with three degrees of freedom, which still captures the pathological behaviour of the full-order system. Under this model, the dynamics evolve via

$$\begin{aligned} \dot{v} &= \frac{1}{m}\left(T_r\cos(\alpha) - D\right) - g\sin(\gamma), \\ v\dot{\gamma} &= \frac{1}{m}\left(T_r\sin(\alpha) + L\right) - g\cos(\gamma), \\ \dot{y} &= v\sin(\gamma), \end{aligned}$$

subject to the initial conditions $(v(0), \gamma(0), y(0)) = (v_0, \gamma_0, y_0)$. Here, the coordinates y, v, and γ respectively model the altitude, velocity, and flight path angle. The control input α represents angle of attack. The constants m and g are the vehicle mass and the gravitational acceleration, while the coefficient maps T_r, L, D are the vehicle thrust, aerodynamic lift and drag forces. For simplicity, we consider a constant thrust model, and model the drag and lift as

$$\begin{aligned} D(v,\, \alpha,\, y,\, \theta) &= \frac{1}{2}v^2\rho(y)S(\theta)C_D(\alpha, \theta), \\ L(v,\, \alpha,\, y,\, \theta) &= \frac{1}{2}v^2\rho(y)S(\theta)C_L(\alpha, \theta), \end{aligned}$$

where S is the reference area. The air density $\rho(y) \doteq \rho_0 \exp\big(-c(y-y_0)\big)$ is modeled via exponential decay as from a reference altitude with decay rate c, with reference density ρ_0 at y_0. The coefficient functions C_D and C_L are typically modeled as low-order polynomials in α, although more sophisticated methods have been proposed [20], [24]. Even though this is a reductive assumption, for a fixed vehicle design, estimates on the coefficients must be re-calculated via regression. The data used in this regression is obtained from CFD sweeps across the operating envelope for each design point. This data acquisition is by far the most computationally expensive component in the overall model. The issue of tractability is compounded in the codesign setting, where the outer loop's parameter updates invalidate the previous iteration's estimates on these coefficients. In our trimmed model, we assume that α is controlled directly (in contrast to the full-order model, where α varies according to some setpoint value and higher-order dynamics).

We also impose the following altitude, flight path angle, angle of attack bounds, as well as constraints on the dynamic

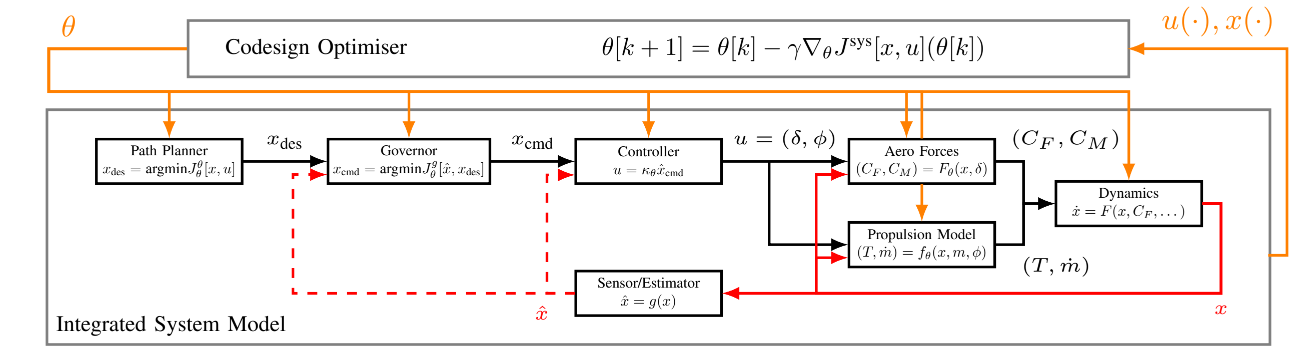

Fig. 1. Schematic of codesign framework. Note the different modules in the integrated system model may use different fidelity process models with resulting different uncertainty bounds.

pressure and terminal altitude:

$$\begin{aligned}
\underline{y} < y < \overline{y}, &\qquad \text{(altitude)}\\
\underline{\gamma} < \gamma < \overline{\gamma}, &\qquad \text{(flight path angle)}\\
\underline{\alpha} < \alpha < \overline{\alpha}, &\qquad \text{(angle of attack)}\\
\underline{q} < \frac{1}{2}\rho(y)v^2 < \overline{q}, &\qquad \text{(dynamic pressure)}\\
\underline{y}_T \le y(T) \le \overline{y}_T, &\qquad \text{(terminal altitude).}
\end{aligned}$$

These types of hard constraints are typically required to ensure the vehicle survivability over the mission. The feasibility of the inner-loop control problem with respect to these constraints is highly sensitive to design parameters changes, which occur in the outer loop. This makes finding feasible design parameters challenging, precluding the use of naïve nested approaches. To address this challenge, we propose an auxiliary test for feasibility. This test not only detects infeasibility, but provides a method to update the outer-loop design parameters to improve feasibility of the inner-loop control problem.

B. Codesign Problem

The simplified codesign problem for the vehicle model described above has the general form of minimizing a cost functional

$$C[u](\theta) \doteq \int_0^T R(t, x(t), u(t), \theta)\, dt, \tag{1a}$$

where the constraints are given by

$$\begin{aligned}
\dot{x} &= f(x(t), u(t), \theta), & t \in [0, T], \quad & \text{(1b)}\\
0 &\le c_j(x(t), u(t), \theta), & j = 1, \dots, J, \quad & \text{(1c)}\\
0 &\le \chi_k(x(0), x(T), \theta), & k = 1, \dots, K. \quad & \text{(1d)}
\end{aligned}$$

Here t denotes the time, $x : [0, T] \to \mathbb{R}^n$ is the state of the system, and $u \in \mathcal{U} \subset L^\infty([0, T]; \mathbb{R}^m)$ is the control. The design parameters θ are constrained in some set $\Theta \subset \mathbb{R}^p$. The function $f \in C^1(\mathbb{R}^n \times \mathbb{R}^m \times \Theta; \mathbb{R}^n)$ governs the state dynamics, and $R : [0, T] \times \mathbb{R}^n \times \mathbb{R}^m \times \Theta \to \mathbb{R}$ is the running cost. The functions $c_j : \mathbb{R}^n \times \mathbb{R}^m \times \Theta \to \mathbb{R}$ and $\chi_k : \mathbb{R}^n \times \mathbb{R}^n \times \Theta \to \mathbb{R}$ constrain the state and control of the system within the finite time horizon and at the endpoints, respectively. Here, we have rewritten equality constraints as pairs of inequalities. These conditions typically constrain the solution and its controls to lie in compact sets $X \subset \mathbb{R}^n$ and $U \subset \mathbb{R}^m$. The admissible class of controls then becomes $\mathcal{U} \doteq L^\infty([0, T]; U)$. If a solution to the constrained minimization problem (1a)-(1d) exists, then nested codesign gives a locally optimal design $\theta^\star$, as well as corresponding optimal trajectory and controls $x^\star(\cdot)$ and $u^\star(\cdot)$.

Remark 1: In the case where f is linear in its arguments and only the control is constrained, i.e. $c_j(x, u, \theta) = c_j(u)$, determining the feasibility of the problem is trivial.

For nonlinear systems with constrained states, the set of inputs for which the control problem is feasible can be very sensitive to minor perturbations of the design parameters. This sensitivity can make the problem difficult to initialize, and parameter updates of the outer loop can lead to failure of existence of solutions to the constrained dynamical system (1b)-(1d), or otherwise infeasibility of the optimal control problem.

We will see that this is the case for the HSV models, including the simplified model outlined in the previous section. Consequently, a reasonable codesign approach requires detection of when the parameters are in a region where there is no feasible solution, and correction towards a feasible design. In order to facilitate our approach, we will consider the problem of minimizing (1a) for a fixed design, by writing

$$W(\theta) \doteq \inf_{u \in \mathcal{U}} \{C[u](\theta) \;:\; \text{(1b)-(1d) hold}\}.$$

C. Feasibility Detection

Our proposed solution relies on the observation that the existence of a feasible trajectory to the constrained dynamical system (1b)-(1d) is independent of the cost function (1a) appearing in the minimization problem. Inspired by tools from sensitivity analysis [7], we propose the following optimal control problem with auxiliary cost whose parameters allow for relaxation of the constraints. This perturbation is designed to detect which, if any, of the constraints in (1c)-(1d) are involved in the failure of feasibility of the optimal control problem (1a)-(1d).

To do so, we introduce the parameters $\mu \in \mathbb{R}^J$ and $\nu \in \mathbb{R}^K$ corresponding to the constraints c and χ, and consider

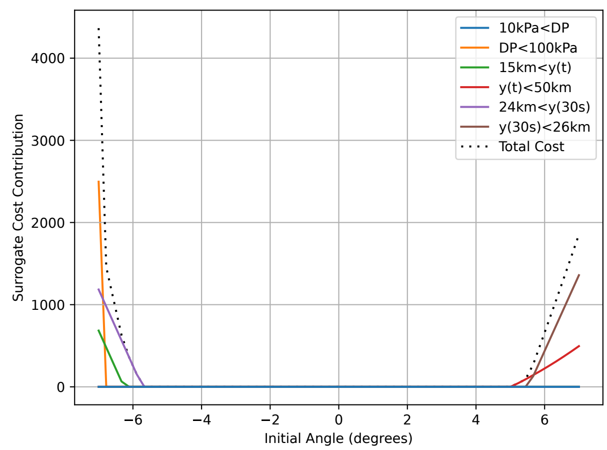

Fig. 2. Failure of feasibility of the trimmed HSV dynamics surrogate optimal control problem, measured by $W_{\rm fs}$ as a function of γ_0. Infeasibility results in non-zero surrogate cost.

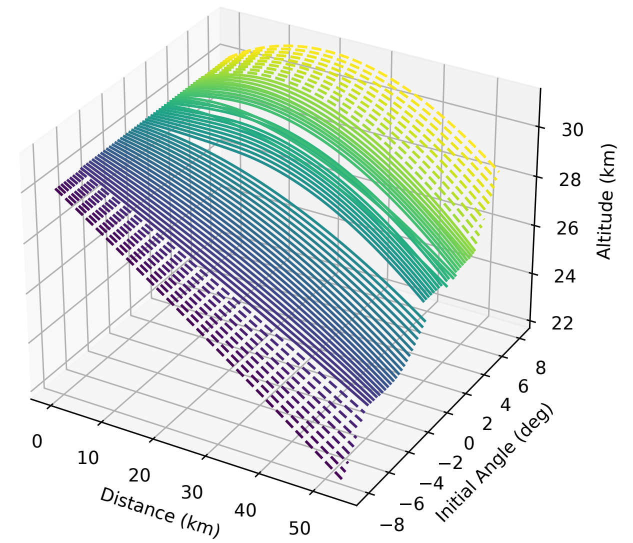

Fig. 3. Optimal trajectories with respect to the surrogate cost, as a function of γ_0. Trajectories violating the constraints are displayed as dotted lines. Infeasibility corresponds to trajectories that violate the constraints.

the surrogate cost

$$C_{\rm fs}[u,\mu,\nu](\theta) \doteq \mu^\top \mu + \nu^\top \nu. \tag{2}$$

For fixed design θ, this lets us solve the auxiliary optimal control problem

$$W_{\rm fs}(\theta) \doteq \inf_{\substack{u\in\mathcal{U},\\ (\mu,\nu)\in\mathbb{R}^{J+K}}} \{C_{\rm fs}[u,\mu,\nu](\theta) \,:\, \text{(3b)-(3d) hold}\}, \tag{3a}$$

where the constraints are given by

$$\begin{aligned} \dot{x}(t) &= f(x(t),u(t),\theta), && t\in[0,T], && (3b)\\ -\mu_j^2 &\le c_j(x(t),u(t),\theta), && j=1,\dots,J, && (3c)\\ -\nu_k^2 &\le \chi_k(x(0),x(T),\theta), && k=1,\dots,K. && (3d)\end{aligned}$$

Under standard conditions on f, c_i and χ_j given below, we have the following proposition. The result follows directly from the existence of solutions to the ODE appearing in (3b) with constant control, since the surrogate cost is constructed to always have a feasible solution.

Proposition 1: Let $X\subset\mathbb{R}^n$ and $U\subset\mathbb{R}^m$ be compact sets and $f\in C^1(\mathbb{R}^n\times\mathbb{R}^m\times\Theta;\mathbb{R})$. Further suppose that $c_j\in C^1(X\times U\times\Theta;\mathbb{R}), \chi_k\in C^1(X\times X\times\Theta;\mathbb{R})$ for $j=1,\dots,J$ and $k=1,\dots,K$. Then (3a)-(3d) is feasible.

Proof: By multiplication of f with smooth cut-off function $\phi\in C_0^\infty(\mathbb{R}^n;R)$ that vanishes outside of X in a controlled way if necessary, differentiability implies that this product is Lipschitz continuous, so that the existence of uniformly bounded solutions to $\dot{x}_0=\phi(x_0)f(x_0,u_0,\theta)$ for any fixed constant $u_0\in U$ is guaranteed for $t\in[0,T]$. Here, the initial data appear in χ. Since the c_i and χ_j are continuous they are themselves bounded on $B_{M(x_0)}\times U$ and $B_{M(x_0)}\times B_{M(x_0)}$ for each fixed θ, where $B_{M(x_0)}$ is a ball of radius $M(x_0)\doteq\max_{t\in[0,T]}|x_0(t)|$. This gives an upper bound on $W_{\rm fs}(\theta)$. ■

We remark that we have assumed slightly higher smoothness than is required for feasibility detection, in order to ensure the differentiability that will be used in the next section to generate parameter updates. By construction, the existence of a solution to (3a)-(3d) with $\mu=0$ and $\nu=0$ implies that the same trajectory gives a finite value of $W(\theta)$, and so the codesign problem $W(\theta)$ is feasible. In this situation, the codesign procedure can proceed by solving the optimal control problem (1a)-(1d). On the other hand, a positive value of $W_{\rm fs}(\theta)$ detects infeasibility of $W(\theta)$, and the sets $\{j\in 1,\dots,J:\mu_j\neq 0\}$ and $\{k\in 1,\dots,K:\nu_k\neq 0\}$ identify the violated constraints.

D. Evaluation and Iteration

Our proposed method is summarized as pseudocode in Algorithm 1. We use a direct collocation method to simulate the dynamics and control and thus evaluate the cost in (3a). Here, we first discretize the interval $0=t_0<\dots<t_{N_\alpha}$ into N_α segments, and on each interval $[t_{\alpha-1},t_\alpha]$ approximate x by a d degree polynomial x_α and u by a constant u_α. We ease notation by writing x and u for the collection of these approximations. This reduces problem (3a)-(3d) to the following nonlinear program for the infeasibility cost $W_{\rm fs}(\theta)$, which can be solved with off-the-shelf solvers,

$$W^\alpha_{\rm fs}(\theta) \doteq \inf_{\substack{u\in\mathcal{U},\\ (\mu,\nu)\in\mathbb{R}^{J+K}}} \{C_{\rm fs}[u,\mu,\nu](\theta) \,:\, \text{(4b)-(4d) hold}\}, \tag{4a}$$

subject to the constraints

$$\begin{aligned} \dot{x}_\alpha(t_\alpha) &= f(x_\alpha(t_\alpha),u_\alpha(t_\alpha),\theta), && \alpha=1,\dots,N_\alpha, && (4b)\\ x_\alpha(t_\alpha) &= x_{\alpha+1}(t_\alpha), && \alpha=1,\dots,N_\alpha\text{–}1, && (4c)\\ -\mu_j^2 &\le c_j(x_\alpha(t_\alpha),u_\alpha,\theta), && j=1,\dots,J, && (4d)\\ & && \alpha=1,\dots,N_\alpha, &&\\ -\nu_k^2 &\le \chi_k(x_1(0),x_{N_\alpha}(T),\theta), && k=1,\dots,K, && (4e)\\ & && \alpha=1,\dots,N_\alpha, &&\end{aligned}$$

with each $x_\alpha\in C^d([t_{\alpha-1},t_\alpha];X)$, $u_\alpha\in U$, and $(\mu,\nu)\in\mathbb{R}^{J+K}$.

Once the infeasibility cost has been evaluated, we can update the parameters θ in the direction of 'best feasibility

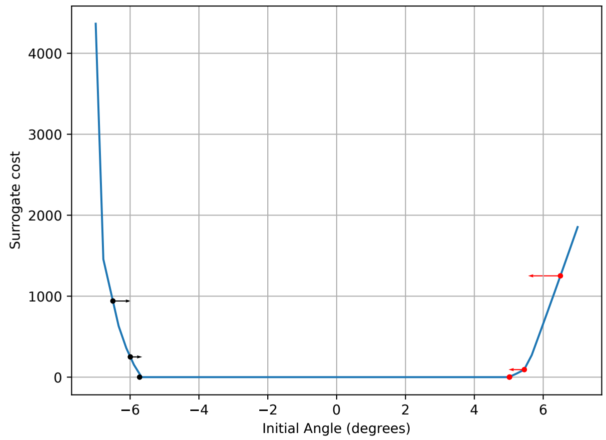

Fig. 4. Gradient-based parameter updates on the surrogate optimal control problem, initialized at ± 0.1134464 radians. Feasibility is achieved after 2 Newton iterations (black) and 3 iterations (red).

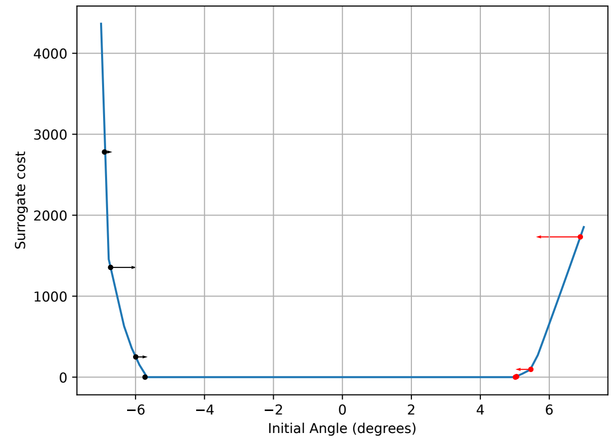

Fig. 5. Gradient-based parameter updates on the surrogate optimal control problem, initialized at ± 0.1204277 radians. Feasibility is achieved after 3 iterations in both cases.

Algorithm 1: Feasibility initialization via Newton's Method

Input: Initial design parameters θ; nonlinear program formulation (4a)-(4e) with objective cost $C[u](\theta)$ and surrogate cost W_{fs}; initialized trajectory and control discretizations x and u; initialized constraint violation vectors $\mu > 0 \in \mathbb{R}^J, \nu > 0 \in \mathbb{R}^K$, cost tolerance $\varepsilon > 0$.
while $\|\mu\| + \|\nu\| > \varepsilon$ **do**
  update $(u, \mu, \nu) \leftarrow \arg\min C_{\text{fs}}[u, \mu, \nu](\theta)$.
  compute $\nabla_\theta W^\alpha_{\text{fs}}(\theta)$.
  update $\theta \leftarrow \theta - W^\alpha_{\text{fs}}(\theta)/\nabla_\theta W^\alpha_{\text{fs}}(\theta)$.
end while

improvement'. We do this by first computing the gradient of the cost with respect to the design parameters $\nabla_\theta W^\alpha_{\text{fs}}(\theta)$. We then update θ using Newton's method [6], according to

$$\theta_{k+1} = \theta_k - W^\alpha_{\text{fs}}(\theta_k)/\nabla_\theta W^\alpha_{\text{fs}}(\theta_k),$$

where division of a scalar by a vector is interpreted element-wise. Note that the version of Newton's method we have used is for finding zeros of a function and is appropriate here, rather than the version for seeking critical points that may be more familiar to the reader.

Remark 2: While we have chosen to implement Newton's method for the gradient update, gradient descent or any number of its variants could also be used and may be more suitable in other applications. Using Newton's method to find zeros of the gradient may also be suitable in some situations, but non-convexity precluded its usefulness in our example.

This process of feasibility detection and improvement can be repeated until design parameters are found that admit a feasible solution. At this stage, the usual codesign optimization procedure can commence. Analysis of the convergence of this process is left for future work. We demonstrate it in a numerical example to follow.

III. NUMERICAL EXAMPLE

To illustrate Algorithm 1, we consider the trimmed HSV model from Section II. To this end, we set the running cost in (1a) to $R \doteq -v\cos(\gamma)$, so that C maximizes displacement in the x-axis. We used third-degree Lagrange polynomials with zeros at the zeros of the corresponding Legendre polynomials, between $N_\alpha - 1 = 64$ uniformly spaced knot points, although more sophisticated hp-adaptive pseudospectral methods could be used [11]. We implemented the direct collocation scheme using `CasADi` [4], and use its sensitivity tools to obtain the gradients $\nabla_\theta W^\alpha_{\text{fs}}(\theta)$ [5]. For simplicity, and to avoid unnecessary computation involving CFD, we consider problem initialization in the setting where the sole design parameter γ_0 to be the initial angle of the HSV upon entry to the flight path [3]. However, the same technique can be readily incorporated between outer-loop updates in a full codesign procedure with more realistic design parameters. Following Parker et al. [21], the drag and lift coefficient functions C_D and C_L are modeled as low-order polynomials in α given by

$$C_D(\alpha) = C^0_D + C^1_D\alpha + C^2_D\alpha^2, \quad C_L(\alpha) = C^0_L + C^1_L\alpha.$$

The values of the parameters used in our simulations are given in Table I.

Optimal values of W_{fs} are plotted against their corresponding values of γ_0 as the dotted black line in Figure 2. The coloured lines give the values of each μ_i and ν_i. Note that in this example, both terminal constraints were violated for different values of γ_0, as well as the dynamic pressure constraint. The corresponding trajectories $x^\star$ to the optimal controls $u^\star$ have been displayed in Figure 3. The dashed lines represent the infeasible trajectories corresponding to values of γ_0 with nonzero values of W^α_{fs}.

Figure 4 depicts the optimal values of W^α_{fs} against the corresponding values of γ_0, as well as the parameter updates via Algorithm 1. The arrows depict the Newton updates. In this example, γ_0 was initialized at $-6.5°$ (-0.1134464 radians) in black, and $6.5°$ (0.1134464 radians) in red. Feasible trajectories were found after two iterations of the algorithm in the black example, and three in the red example (the third update obstructs the second). Figure 5 reinitializes this experiment at $-6.9°$ (-0.1204277 radians) in black, and

C_D^0	1.0131×10^{-2}	C_L^0	-1.8714×10^{-2}
C_D^1	$4.5315 \times 10^{-2}\text{rad}^{-1}$	C_L^1	$4.6773 \times 10^{0}\text{rad}^{-1}$
C_D^2	$5.8224 \times 10^{0}\text{rad}^{-2}$	c	7km
m	201.59kg	ρ_0	1.23kg/m^3
S	4.5m	T	30s
$\underline{v}$	0m/s	$\overline{v}$	∞ m/s
$\underline{\gamma}$	$\frac{\pi}{4}$rad	$\overline{\gamma}$	$\frac{\pi}{4}$rad
$\underline{y}$	32.5km	$\overline{y}$	21.5km
$\underline{\alpha}$	-0.436332rad	$\overline{\alpha}$	0.436332rad
$\underline{q}$	10kPa	$\overline{q}$	100kPa
$\underline{y}_T$	24km	$\overline{y}_T$	26km
T_r	4031.8N/m	v_0	1500m/s
γ_0	[-1.2217rad,1.2217rad]	y_0	30km

TABLE I

PARAMETER VALUES USED IN THE TRIMMED HSV MODEL.

6.9° (0.1204277 radians) in red. Contrasting Figures 4 and 5 with Figure 2, we note that the parameter updates land near points where constraints are no longer violated. This is due to the behaviour of the cost function, and the fact that Newton's method returns parameters that would have zero cost if the cost functional was linear.

IV. CONCLUSIONS AND FUTURE WORKS

A. Conclusions

Nested codesign is a challenging optimization problem whose difficulties are highlighted when applied to optimal design of bespoke hypersonic vehicles, tailored to specific mission objectives. In particular, the constrained dynamics are highly nonlinear, and depend on lift and drag functions that need to be estimated via expensive CFD simulations. This can lead to design initialization in regions where the dynamics problem is infeasible. We have proposed a method that mitigates this issue by allowing for detection of not only infeasibility itself, but which constraints have been violated by the design parameter settings. The nonlinear program that was used in this detection was differentiated to obtain a gradient that can be used in a parameter update to move towards a feasible design.

B. Future Works

While the current method targets poor initialization, it can immediately be inserted between parameter updates to ensure feasibility throughout the codesign procedure. Similarly, mitigation of the non-global minima found due to the problems inherent nonconvexity requires attention. Formal convergence analysis involving conditions under which one can obtain convergence guarantees, as well as estimates on the error in the cost induced by the discretization are required in order to make the method more robust. Adaptation of reachability techniques based on the Hamilton-Jacobi-Bellman equations [18] to the current setting may also provide additional benefits over current techniques.

V. ACKNOWLEDGMENTS

This work is in collaboration with BAE systems funded through an Australian Research Council Linkage Project grant LP190100104, and AFOSR grant FA2386-22-1-4074.

REFERENCES

[1] J. T. Allison and D. R. Herber. Special section on multidisciplinary design optimization: Multidisciplinary design optimization of dynamic engineering systems. *AIAA Journal*, 52(4):691–710, 2014.

[2] J. T. Allison and S. Nazari. Combined plant and controller design using decomposition-based design optimization and the minimum principle. volume Volume 1: 36th Design Automation Conference, Parts A and B, pages 765–774, 08 2010.

[3] K. An, Z. yun Guo, X. ping Xu, and W. Huang. A framework of trajectory design and optimization for the hypersonic gliding vehicle. *Aerospace Science and Technology*, 106:106110, 2020.

[4] J. Andersson, J. Gillis, G. Horn, J. Rawlings, and M. Diehl. CasADi – A software framework for nonlinear optimization and optimal control. *Mathematical Programming Computation*, 11(1):1–36, 2019.

[5] J. A. Andersson and J. B. Rawlings. Sensitivity analysis for nonlinear programming in CasADi. *IFAC-PapersOnLine*, 51(20):331–336, 2018. 6th IFAC Conference on Nonlinear Model Predictive Control 2018.

[6] K. E. Atkinson. *An Introduction to Numerical Analysis*. John Wiley & Sons, New York, second edition, 1989.

[7] J. F. Bonnans and A. Shapiro. *Perturbation Analysis of Optimization Problems*. Springer New York, New York, NY, 2000.

[8] L. W. Cook, A. A. Mishra, J. P. Jarrett, K. E. Willcox, and G. Iaccarino. Optimization under turbulence model uncertainty for aerospace design. *Physics of Fluids*, 31(10):105111, 2019.

[9] Z. T. Dydek, A. M. Annaswamy, and E. Lavretsky. Adaptive control and the nasa x-15-3 flight revisited. *IEEE Control Systems Magazine*, 30(3):32–48, 2010.

[10] M. Garcia-Sanz. Control co-design: An engineering game changer. *Advanced Control for Applications*, 1(1):e18, 2019.

[11] P. Han, J. Shan, and X. Meng. Re-entry trajectory optimization using an hp-adaptive radau pseudospectral method. *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, 227(10):1623–1636, 2013.

[12] D. R. Herber and J. T. Allison. Nested and Simultaneous Solution Strategies for General Combined Plant and Control Design Problems. *Journal of Mechanical Design*, 141(1), 10 2018. 011402.

[13] K. W. Lee, W. Moase, A. Ooi, C. Manzie, and E. C. Kerrigan. Optimization framework for codesign of controlled aerodynamic systems. *AIAA Journal*, 54(10):3149–3159, 2016.

[14] K. W. Lee, W. H. Moase, S. Z. Khong, A. Ooi, and C. Manzie. Aerodynamic shape optimization via global extremum seeking. *IEEE Transactions on Control Systems Technology*, 23(6):2336–2343, 2015.

[15] K. W. Lee, W. H. Moase, C. Manzie, N. Hutchins, A. Ooi, J. Vethecan, and P. Riseborough. Is there a need for fully converged CFD solutions? global extremum seeking applied to aerodynamic shape optimisation. In *2013 Australian Control Conference*, pages 385–390, 2013.

[16] K. W. Lee, W. H. Moase, A. Ooi, and C. Manzie. Adjoint assisted geometry design of a feedback controlled missile. *IFAC Proceedings Volumes*, 47(3):451–456, 2014. 19th IFAC World Congress.

[17] V. Liu, C. Manzie, and P. M. Dower. Constraint handling of an airbreathing hypersonic vehicle via predictive reference management. In *2022 American Control Conference (ACC)*, pages 3991–3996, 2022.

[18] V. Liu, C. Manzie, and P. M. Dower. Reachability of linear time-invariant systems via ellipsoidal approximations. *IFAC-PapersOnLine*, 56(1):126–131, 2023. 12th IFAC Symposium on Nonlinear Control Systems NOLCOS 2022.

[19] K. Mackle, I. Jahn, and R. Gollan. Understanding the impacts of aerodynamic uncertainty on optimal trajectories for hypersonic vehicles. In *AIAA AVIATION 2021 FORUM https://doi.org/10.2514/6.2021-2507*.

[20] A. March, K. Willcox, and Q. Wang. Gradient-based multifidelity optimisation for aircraft design using bayesian model calibration. *The Aeronautical Journal*, 115(1174):729–738, 2011.

[21] J. T. Parker, A. Serrani, S. Yurkovich, M. A. Bolender, and D. B. Doman. Control-oriented modeling of an air-breathing hypersonic vehicle. *Journal of Guidance, Control, and Dynamics*, 30(3):856–869, 2007.

[22] A. Skoraczynski, C. Manzie, and P. M. Dower. Incentivizing local controllability in optimal trajectory planning. In *2023 American Control Conference (ACC)*, In print, 2023.

[23] A. K. Sundarrajan and D. R. Herber. Towards a fair comparison between the nested and simultaneous control co-design methods using an active suspension case study. In *2021 American Control Conference (ACC)*, pages 358–365, 2021.

[24] K. S. Won and T. Ray. A framework for design optimization using surrogates. *Engineering Optimization*, 37(7):685–703, 2005.