

# A Coupled Analytical-Simulation Approach to Redundant Small Spacecraft AOCS

Tomas Mrazek  
*Technical University of Munich*  
*Chair of Space Mobility and Propulsion*  
Munich, Germany  
tomas.mrazek@tum.de

Aman Arora  
*University of Luxembourg*  
*Space Robotics Research Group*  
Luxembourg, Luxembourg  
aman.arora@uni.lu

Matteo El Hariry  
*University of Luxembourg*  
*Space Robotics Research Group*  
Luxembourg, Luxembourg  
matteo.elhariry@uni.lu

Yasin Kale  
*Technical University of Munich*  
*Chair of Space Mobility and Propulsion*  
Munich, Germany  
yasin.kale@tum.de

Chiara Manfretti  
*Technical University of Munich*  
*Chair of Space Mobility and Propulsion*  
Munich, Germany  
chiara.manfretti@tum.de

**Abstract**—Missions including close-proximity operations and docking (CPOD) as well as in-orbit servicing are redefining the capabilities of small spacecraft. As a result, attitude and orbit control systems (AOCS) requirements have evolved to include not only attitude maintenance, but also 6-DoF translation and rotation. Cold-gas thruster arrays are key to these capabilities, but optimising their design for sustained as well as impulsive performance while maintaining redundancy cannot be addressed by existing impulsive-centric methods.

In this work, we present a coupled analytical and simulation design framework for defining thruster configurations while imposing deterministic redundancy. Established convex geometry methods for fault-tolerant design are reformulated to always ensure sufficient control authority, dictated by close-proximity operations and docking capability. A modular simulation stack allows for the inclusion of a Reinforcement Learning policy as a CPOD guidance module and of effects such as valve transients and small-scale nozzle efficiency. An optimised thruster layout is obtained using a genetic algorithm and verified for a rendezvous and servicing mission, representing an in-orbit servicing demonstration of Water Electrolysis Propulsion (WEP). The observed 15.7% reduction in thruster on-time is contextualised against typical performance drivers in cold-gas AOCS design.

This work establishes a baseline for AOCS hardware development within the Ice2Thrust EIC Pathfinder project. Ice2Thrust proposes WEP as a non-toxic, high-performance, and refillable alternative to conventional propulsion systems. By integrating redundancy and realistic system dynamics early in the design process, the present approach contributes to a system that is robust and prepared for mission extension applications.

**Index Terms**—close-proximity operations, in-orbit servicing, attitude and orbit control, fault-tolerant control, water electrolysis propulsion

## I. INTRODUCTION

Thrusters are the only viable attitude and orbit control method for a range of mission profiles. Of those, close-proximity operations and docking present perhaps the greatest challenge, requiring precise 6 Degrees of Freedom (DoF)

manoeuvring capability not offered by momentum exchange methods such as reaction wheels. Thruster arrays are equally indispensable for missions carrying vibration-sensitive payloads, and are favoured beyond LEO or where thrust imbalances exist, as they require no momentum desaturation and operate independently of external magnetic fields [6, 8].

Among spacecraft subsystems, Attitude and Orbit Control System (AOCS) exhibit the highest failure rates. Among AOCS actuator types, it is thrusters in particular that fail most often [19]. As a result, missions that rely on thruster-based attitude control incorporate redundancy in their designs. A common approach is to equip the spacecraft with two identical, duplicate actuator systems [2, 12]. More recently, resilient thruster configurations have been proposed for future spacecraft that incorporate redundancy without the need for a dedicated backup array [11]. Design methods for such redundant layouts have been documented by Crawford (1969) and Chen (2024) [5, 4].

In this paper, a design method is proposed for a optimising thruster array layout while retaining redundancy with respect to defined mission requirements. In addition to redundancy verification, a computationally efficient, quasi-static evaluation of candidate layouts enables the implementation of a genetic algorithm (GA) for further refinement. A high-fidelity Basilisk simulation is used for validation and establishes a baseline for further parameter space exploration [3]

This work contributes to the AOCS design of a WEP spacecraft within the EIC Pathfinder Ice2Thrust project. Demonstrating docking and refilling capability requires the AOCS to execute close-proximity manoeuvres in 6-DoF alongside pointing and thrust vector control. Using the proposed design method, a cold-gas thruster array incorporating a redundant layout is created and validated in a dynamic simulation. The correlation between the quasi-static candidate evaluation method and a full dynamic simulation is assessed. The sensitivity to total mission impulse of the thruster layout is

The authors gratefully acknowledge the financial support of the European Union (GA number 101161690) for the Ice2Thrust (S4I2T) project.

explored and compared to other typical performance levers and uncertainties in AOCS design.

## II. MISSION DEFINITION

A demonstration mission involving two 16U CubeSats is proposed in this work. An active and passive spacecraft are defined as the "chaser" and "target" respectively. Both spacecraft are equipped with the Ice2Thrust WEP system. An electrolyser decomposes water into hydrogen and oxygen using on-board solar arrays. A hot-gas thruster directly combusts the propellants for main propulsion, and a cold-gas thruster array employs either  $H_2$  or  $O_2$  interchangeably to maintain a target gas ratio in the tanks.

The cold-gas thruster array is responsible for attitude- and fine orbit control, as well as for correcting any misalignment of the hot-gas thrust vector with respect to the Centre of Mass (CoM), substituting a full Thrust Vector Control (TVC) system. The specifications of both thrusters are given in Table I and physical properties of both spacecraft in Table II.

TABLE I  
THRUSTER PERFORMANCE FIGURES

	Thrust $F$ , N	Specific Impulse $I_{sp}$ , s
Hot-gas	2.2	370
Cold-gas	0.065	66 ( $O_2$ ) / 256 ( $H_2$ )

TABLE II  
REFERENCE SPACECRAFT PHYSICAL PROPERTIES

Mass (g)	34,831
Dimensions $[x, y, z]$ (m)	$[0.2263, 0.2263, 0.454]$
Inertia Tensor ( $kg^2 m^{-1}$ )	$\begin{pmatrix} 0.567 & 0.0 & 0.0 \\ 0.0 & 0.567 & 0.0 \\ 0.0 & 0.0 & 0.227 \end{pmatrix}$

From the chaser's perspective, a demonstration mission consists of the following phases:

*a) Detumbling:* Following deployment, the satellite experiences rotation about all its axes, referred to as "tip-off rates". These are assumed to be  $10^\circ s^{-1}$  along all axes for a typical deployment system [17].

*b) Commissioning:* Satellite systems are checked out. During this period, attitude and orbit are maintained within defined limits.

*c) Rendezvous:* An orbit transfer is executed with the hot-gas thruster to place the chaser in the same orbit as the target. A lack of hot-gas thrust vector control imposes additional requirements on the AOCS.

*d) Close-Proximity and Docking (CPOD):* Throughout the final approach, the spacecraft must follow a passively safe trajectory. This ensures that a collision is avoided even if anomalies, such as propulsion, navigation, or sensor failures occur. It is assumed that the chaser spacecraft is cooperative, which allows for both R-bar and V-bar approaches to be considered. An R-bar approach follows a radial vector emanating

from the centre of the earth, towards the target spacecraft. The R-bar approach is suitable for a quick approach and exploits the gravity gradient between the two spacecraft to achieve passive safety [7]. The V-bar approach follows the target's orbital velocity vector, allowing for time-unconstrained hold points, consistent illumination [7], and better Global Navigation Satellite System (GNSS) performance [22]. The V-bar hopping approach in particular benefits from high  $\Delta V$  efficiency [16] and is selected for this demonstration mission. Once the target and chaser are sufficiently close, the chaser enters the keep-out sphere through an approach corridor, defined as a cone with a half-angle of  $15^\circ$  - a typical value for close-proximity operations [7].

*e) Separation and disposal:* In this phase, the hot-gas thruster is used to move the chaser to a disposal orbit, again requiring thrust vector control to be achieved through cold-gas array firing.

To isolate it from broader mission planning, the AOCS design method presented in this work is validated against a simulation of the final CPOD approach only.

## III. DESIGN METHOD

### A. Layout Definition

A *working* system is defined by Crawford (1969) as one where every vector  $\vec{v}$  in the  $D$ -dimensional activity space is achievable by a linear combination of working activity vectors, with a constraint of only positive control inputs [5]. In addition, the Level of Redundancy (LR) of a system is defined as one less than the number of thruster failures that need to occur to render the system *non-working*. For the 6-dimensional task of CPOD and a desired LR of 1, a minimum condition for the number of thrusters for working system is given by [5]

$$2LR + D + 1 \leq N = 9. \quad (1)$$

For Ice2Thrust, a fixed value of  $N = 12$  was chosen to facilitate integration into a 16U CubeSat shape and leverage symmetry throughout the design. Possible thruster locations  $\vec{r}_n$  were defined as a discrete set of 40 points near the  $\pm Z$  face plates, shown in Figure 1. In addition, each thruster firing vector  $\vec{v}_n$  is defined by a face-tilt angle  $\theta$  and face-azimuth angle  $\phi$ . The range of  $\theta$  is limited to  $\pm 40^\circ$  to limit self-impingement with ejected propellant. Gas impingement on the target spacecraft is also a concern, but was not considered in this optimisation.

Given the requirements in Section III-B, thruster layout generation leverages symmetry to increase the likelihood of generating compliant configurations. Single seed thrusters are placed randomly along aforementioned discrete points, spawning a set of 2 or 4 K4 planar-symmetric thrusters that, when fired simultaneously, produce pure axial translation. This algorithm alone does not guarantee pure translation along each axis, but its simplicity allows for sufficiently fast iteration to obtain a large number of viable configurations. As each configuration is fully defined by its seed thrusters alone, the




Fig. 1. Possible thruster locations marked in red, hot-gas thrust vector marked in orange

solution space as well as number of parameters describing a configuration are reduced.

Each configuration can be described by a  $D \times N$  thruster configuration matrix  $\mathbf{A}$ , where  $D = 6$  corresponds to 6-DoF spacecraft control and  $N$  corresponds to the number of thrusters.

$$\mathbf{A} = \begin{pmatrix} \mathbf{V} \\ \mathbf{L} \end{pmatrix} \quad (2)$$

where  $\mathbf{V}$  is a matrix of thruster firing vectors  $\vec{v}_n$

$$\mathbf{V} = \begin{pmatrix} v_{x,1} & v_{x,2} & \cdots \\ v_{y,1} & v_{y,2} & \cdots \\ v_{z,1} & v_{z,2} & \cdots \end{pmatrix} \quad (3)$$

and  $\mathbf{L}$  is a matrix of thruster firing moments, obtained by thruster position vector  $\vec{r}_n$  and firing vectors  $\vec{v}_n$ .

$$\mathbf{L} = \begin{pmatrix} \vec{r}_1 \times \vec{v}_1 & \cdots & \vec{r}_N \times \vec{v}_N \end{pmatrix} \quad (4)$$

### B. System Requirements

With 6-DoF operation in mind, an additional requirement was defined to maximise performance in pure translation and rotation and mandate symmetry in generated thruster configurations. While this risks excluding optimal asymmetric configurations, it allows for significant reduction of the solution space. Therefore, the following requirement is only applied with no failures present: *For each axis, one subset of thrusters must exist that when fired at full thrust, produces net force and torque only along the commanded axis.*

All candidate configurations are subject to requirements for minimum authority. These arise from the need for thrust vector misalignment and sufficient control authority during CPOD. With no hot-gas thruster (HGT) TVC capability, the AOCS must correct for hot-gas thrust vector misalignment arising from installation uncertainty or CoM shifts during firing. Assuming a hot-gas thrust vector uncertainty of  $1^\circ$  about its X and Y axes, the minimum torque authority  $T_{min}$  is calculated from the worst-case moments produced by the hot-gas thruster. The minimum force authority  $F_{min}$  is defined by a required per-axis acceleration of  $0.5 \text{ mm s}^{-2}$ , giving a minimum control authority vector

$$\vec{W}_{min} = \begin{bmatrix} F_{min,x} \\ F_{min,y} \\ F_{min,z} \\ T_{min,x} \\ T_{min,y} \\ T_{min,z} \end{bmatrix} = \begin{bmatrix} 0.02 \text{ N} \\ 0.02 \text{ N} \\ 0.02 \text{ N} \\ 0.009 \text{ N m} \\ 0.009 \text{ N m} \\ 0.0001 \text{ N m} \end{bmatrix} \quad (5)$$

### C. Redundant Control Authority

Following a single thruster failure, the AOCS must remain *working*, as defined in Section III-A. This requires not only that 6-DoF control is retained, but that the minimum control authority vector  $\vec{W}_{min}$  remain reachable.

From the thruster configuration matrix, an effective thrust matrix  $\mathbf{A}_{eff} = \mathbf{A} \text{diag}(\vec{T})$  is defined using the maximum continuous thrust of each thruster. To verify sufficient control authority, we normalise the effective configuration matrix  $\mathbf{A}_{eff}$  by per-axis minimum control authority vector  $\vec{W}_{min}$

$$\mathbf{A}^* = \text{diag}(\vec{W}_{min})^{-1} \mathbf{A}_{eff}. \quad (6)$$

This gives a non-dimensional configuration matrix expressing the fractional satisfaction of per-axis requirements by each thruster.  $\mathbf{A}^*$  must be full-rank to ensure the entire space of torques and translations is covered with zero thruster failures.

For each set of  $D - 1$  linearly independent vectors of  $\mathbf{A}^*$ , defined as  $\tilde{\mathbf{A}}$ , a  $D - 1$  dimensional hyperplane can be formed containing all  $D - 1$  thrust vectors. This defines a direction when the maximum number of thrusters cannot contribute to propulsion as they lie orthogonal to hyperplane normal  $\vec{l}$ . The hyperplane normal for dimension  $j$  is defined by [4] as

$$\vec{l} = \left( \tilde{\mathbf{A}}^T \right)^{-1} \begin{pmatrix} 0 \\ \vdots \\ 0 \\ 1 \end{pmatrix} \quad (7)$$

where  $\vec{e}_j$  is column  $j$  in a D-order identity matrix.

Of the remaining  $n - (D - 1)$  vectors, contained in matrix  $\hat{\mathbf{A}}$ , at least  $LR + 1$  vectors must lie on either side of this hyperplane. We can define a vector  $\vec{d}$  representing the distance of each thruster to this hyperplane

$$\vec{d} = \hat{\mathbf{A}}^T \cdot \frac{\vec{l}}{\|\vec{l}\|}. \quad (8)$$

The sets of resulting positive and negative terms of  $\vec{d}$ ,  $\mathcal{S}^+$  and  $\mathcal{S}^-$ , give the fractional contributions of each thruster to the required per-axis thrust and torque. By ordering  $\mathcal{S}^+$  and  $\mathcal{S}^-$  in ascending order of magnitude and removing the "strongest" thrusters from each set, the  $LR$  is defined as the greatest number of thrusters removed for which




Fig. 2. Static analysis force vectors (orange) compared to force vectors obtained in dynamic Basilisk simulation (gray)

$$C(\text{LR}) = \sum_{j=1}^{N-\text{LR}} \mathcal{S}_j \geq 1.0 \quad (9)$$

for both the positive and negative sets. For the purposes of layout optimisation, the number of tolerable failures is returned for each hyperplane. The mean number of tolerable failures indicates the resilience of the configuration, while the minimum represents the configuration's LR, which must be greater or equal to 1.

#### D. Efficiency Estimation

The efficiency of a configuration represented by  $\mathbf{A}_{\text{eff}}$  can be estimated using static analysis [10]. This method is particularly relevant for missions with a pre-determined trajectory or ones where the AOCS mainly counteracts a known external force [1, 15]. In a static analysis, a dynamic simulation approximated by a calculation of the thruster on-time  $t_{\text{on}}$  required to reach a discretised set of force and torque vectors.

For a mission with balanced impulse demand, the set of force and torque vectors can be defined as a set of unit vectors. Should mission impulse demand be known, this set of unit vectors can be appropriately scaled or replaced with the known set of mission vectors. In this work, a scaled set of unit vectors shown in Figure 2 was used, informed by the accumulated forces and torques of a single dynamic simulation described in Section III-F. As vectors  $\vec{f}_i$  represent translational and rotational impulse rather than force and torque, sum  $\sum u_i$  corresponds to a thruster on-time and can be linked to the propellant use of each configuration.

For the  $i^{\text{th}}$  unit vector  $\vec{f}_i$ , a least-squares solution of control inputs  $u$  to the thruster allocation problem is found

$$u_i = \mathbf{A}_{\text{eff}}^T (\mathbf{A}_{\text{eff}} \mathbf{A}_{\text{eff}}^T)^{-1} \vec{f}_i. \quad (10)$$

The choice of thrust allocation method affects the resulting control authority. The least-squares solution is also known as the minimum power controller, and leverages the Moore-Penrose solution for computational efficiency. However, the maximum attainable forces are lower compared to alternatives such as the minimum flow rate controller [21]. Nonetheless, this method was chosen as it is also applied in the `forceTorqueThrForceMapping` module in Basilisk, which is later used for dynamic validation. In future work, both the static and dynamic analysis could apply a more propellant-efficient allocation scheme.

#### E. Optimisation Method

The definition of an optimal thruster layout is a non-linear optimisation problem. A solution is required that satisfies a set of  $N$  control demands  $\mathcal{W}$  using a set of 12 thrusters. For every required torque  $f_i \in \mathcal{W}$ , the achievable error and cumulative input effort  $u_i$  are minimised in Equation 11.

$$\min_{A, \{u_i\}} \sum_{i=1}^N \left[ \|\mathbf{A}u_i - f_i\|^2 + \alpha \sum_{j=1}^{12} u_i^{(j)} \right] \quad (11)$$

This is a non-convex optimisation problem, making finding a global optimum non-trivial. An implementation of constraints that ensure redundancy and respect consistency of the  $\mathbf{A}$  matrix with physical constraints of a CubeSat would further increase complexity. Instead, this problem is approached with a GA, leveraging methods laid out in Sections III-B-III-D to assign each configuration a performance score.

The GA is implemented using the Distributed Evolutionary Algorithms in Python (DEAP) framework [9] with parameters listed in Table III. Starting with an initial population generated stochastically, the face-tilt and face-azimuth angles  $\phi$  and  $\theta$  of the seed thrusters are varied. Crossover combines the seed

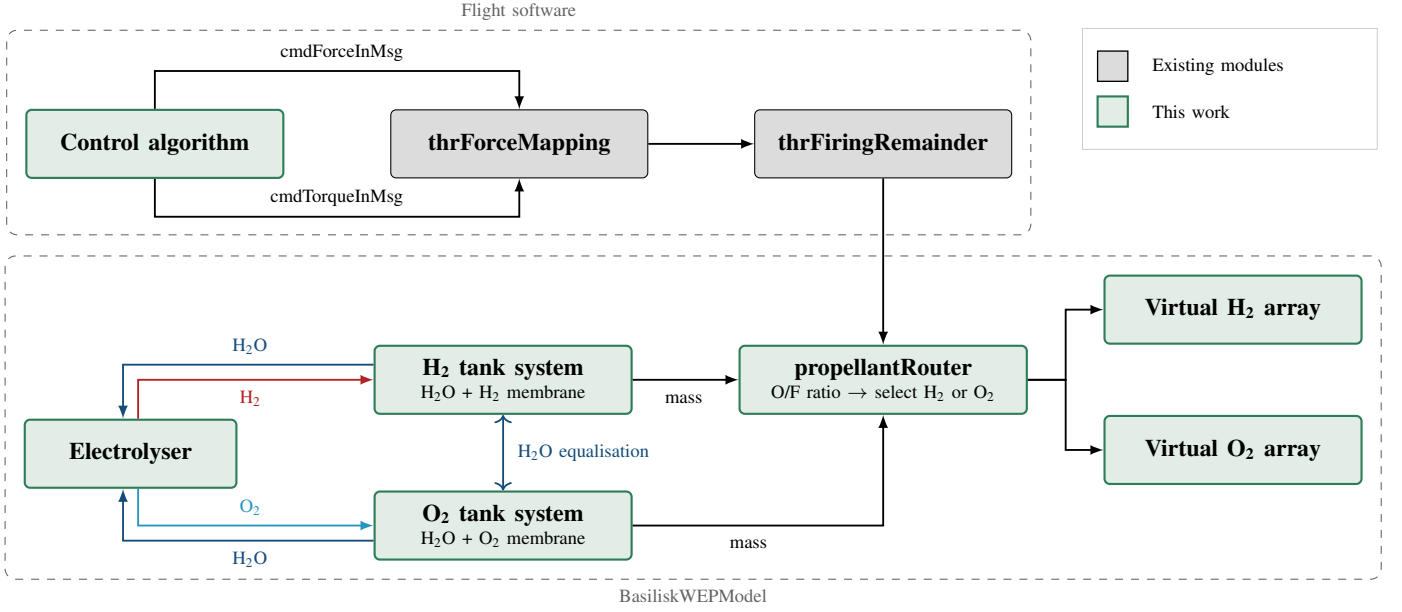


Fig. 3. Simulation architecture of the AOCS only, showing the interaction between the flight software and propulsion system model components

TABLE III  
GENETIC ALGORITHM PARAMETERS

Parameter	Value
Population Size	1,133
Generations	30
Crossover Probability	0.85
Mutation Probability	0.3

thrusters of two configurations to spawn a new 12-thruster configuration. The optimisation objective was chosen to be  $t_{on,stat} = \sum u_i$ .

#### F. Dynamic Simulation

To verify that improved performance predicted by the optimisation translates to dynamic conditions, a simulation of the final close-proximity operations is carried out in Basilisk. This step allows for evaluating the relative impact of optimising the thruster layout compared to performance parameters such as the Minimum Impulse Bit (MIB) or  $I_{sp}$ , further informing a suitable focus for follow-on research. Finally, by establishing a dynamic simulation environment, changes to the thruster layout can be tested in conjunction with representative control algorithms, discussed in Section III-G.

The initial condition for the dynamic simulation represents the chaser at the start of its final approach, immediately outside of the keep-out sphere (KOS). A final approach is only permitted within a  $\pm 15^\circ$  approach corridor and permissible relative initial orientation  $|\theta_{x,y,z}| < 6^\circ$ . The initial condition space was sampled with 15 points, and a simulation was performed for each starting position to evaluate a configuration. The chaser was then commanded to move to its final position for capture and docking, within tolerances informed by docking adapter

specifications. The CPOD controller gains were selected to achieve a docking time of 600 seconds.

To model the Ice2Thrust system in Basilisk, a library of custom modules, *BasiliskWEPModel*, was developed. Models implemented for the AOCS are shown in Figure 3. Selectable-propellant cold-gas thrusters are represented by a set of virtual arrays sharing the same location. In addition, a bi-propellant hot-gas thruster model was implemented to support orbit change burns foreseen in future work. To reduce the number of independent variables, the impact of propellant selection on efficiency was not considered in this work.

#### G. CPOD Policy

The current CPOD implementation uses two decoupled, analytically-derived closed-loop controllers: a Clohessy–Wiltshire-compensated proportional-derivative (CW+PD) controller for translational guidance, and a Modified Rodrigues Parameter proportional-derivative (MRP-PD) controller for attitude alignment. Both run simultaneously at the flight-software rate of 10 Hz, with orbital mechanics propagated by the Basilisk astrodynamics framework [3] at 100 Hz.

Each control step, the *SimpleNav* Basilisk module provides chaser and target navigation states in the Earth-centred inertial (ECI) frame. The Hill (LVLH) frame is constructed from the target ECI state as

$$\hat{\mathbf{x}}_H = \frac{\mathbf{r}_{tgt}}{|\mathbf{r}_{tgt}|}, \quad \hat{\mathbf{z}}_H = \frac{\mathbf{r}_{tgt} \times \mathbf{v}_{tgt}}{|\mathbf{r}_{tgt} \times \mathbf{v}_{tgt}|}, \quad \hat{\mathbf{y}}_H = \hat{\mathbf{z}}_H \times \hat{\mathbf{x}}_H, \quad (12)$$

and the Hill-to-inertial DCM is  $\mathbf{C}_{HN} = [\hat{\mathbf{x}}_H, \hat{\mathbf{y}}_H, \hat{\mathbf{z}}_H]^\top$ . The relative attitude error  $\sigma_{BR}$  and relative angular rate  $\omega_{BR}^B$  are computed inside Basilisk by the *locationPointing* and *attTrackingError* modules, which define the docking

reference as the chaser's  $+x_B$  axis aligned with the target's  $+x_B$  axis (port-to-port).

Relative motion in the Hill frame obeys the Clohessy–Wiltshire equations

$$\ddot{x} = 3n^2x + 2n\dot{y} + F_x/m, \quad (13)$$

$$\ddot{y} = -2n\dot{x} + F_y/m, \quad (14)$$

$$\ddot{z} = -n^2z + F_z/m, \quad (15)$$

where  $(x, y, z)$  is the chaser position relative to the target in the Hill frame (radial, along-track, cross-track) and  $(F_x, F_y, F_z)$  is the commanded specific force in that frame. The controller cancels orbital coupling terms analytically and applies PD feedback:

$$F_x = m[-(K_p + 3n^2)x - K_d\dot{x} - 2n\dot{y}], \quad (16)$$

$$F_y = m[-K_p y - K_d\dot{y} + 2n\dot{x}], \quad (17)$$

$$F_z = m[-(K_p - n^2)z - K_d\dot{z}]. \quad (18)$$

The gains  $K_p = 3 \times 10^{-4} \text{ s}^{-2}$  and  $K_d = 0.035 \text{ s}^{-1}$  place the closed-loop poles at  $-\omega_n \pm \omega_n \sqrt{\zeta^2 - 1}$  with natural frequency  $\omega_n = \sqrt{K_p} \approx 0.017 \text{ rad s}^{-1}$  and critical damping ( $\zeta = 1$ ), giving a linear-regime settling time  $t_s \approx 4/\omega_n \approx 235 \text{ s}$ . The commanded force in the Hill frame is rotated to the chaser body frame via

$$\mathbf{F}_B = \text{sat}_{F_{\max}}(\mathbf{C}_{BH} \mathbf{F}_H), \quad \mathbf{C}_{BH} = \mathbf{C}_{BN}^{\text{chas}} \mathbf{C}_{HN}^\top, \quad (19)$$

where  $\mathbf{C}_{BN}^{\text{chas}}$  is the chaser body-to-inertial DCM derived from its MRP state, and each component is saturated at  $F_{\max} = 0.065 \text{ N}$  (the per-axis thrust limit of the cold-gas array).

Attitude is regulated by a PD control law on the MRP error directly:

$$\tau_B = \text{sat}_{\tau_{\max}}(-K_{p,a} \sigma_{BR} - K_{d,a} \omega_{BR}^B), \quad (20)$$

where  $\sigma_{BR}$  is the MRP attitude error of the chaser body with respect to the target body, and  $\omega_{BR}^B$  is the relative angular rate expressed in the chaser body frame. The gains  $K_{p,a} = 4 \times 10^{-3} \text{ Nm/MRP}$  and  $K_{d,a} = 4 \times 10^{-2} \text{ N m s rad}^{-1}$  a natural frequency  $\omega_{n,a} = \sqrt{K_{p,a}/I_{zz}} \approx 0.133 \text{ rad s}^{-1}$ , a damping ratio  $\zeta_a = K_{d,a}/(2\sqrt{K_{p,a} I_{zz}}) \approx 0.67$ , and a settling time  $t_{s,a} \approx 30 \text{ s}$ . The saturation limit  $\tau_{\max} = 5 \times 10^{-3} \text{ N m}$  matches the physical torque authority of the cold-gas array.

Evaluation over 20 episodes with initial distances uniformly sampled from 8 m to 10 m shows that the combined controller achieves a 100 % success rate on the *pose* stage (capture radius 0.25 m,  $|\sigma_{BR}| < 0.30 \text{ MRP}$ ,  $|\omega_{BR}| < 0.10 \text{ rad s}^{-1}$ , docking speed  $< 0.10 \text{ m s}^{-1}$ ). The attitude error settles from  $|\sigma_{BR}^0| \approx 0.54 \text{ MRP}$  to near zero within approximately 50 s, decoupled from translational dynamics. Episode trajectories form a tight bundle with negligible inter-episode variance, confirming the deterministic nature of the controller.

Although the CW+PD / MRP-PD controller solves the CPOD maneuver reliably, several assumptions limit its applicability to real mission scenarios. The CW equations assume a circular reference orbit and linearise relative motion; both approximations degrade for highly eccentric orbits or large

separations. The decoupled translational and attitude loops provide no coordinated guidance: the chaser may approach the docking port along a trajectory that is geometrically valid in the Hill frame yet dynamically unfavourable (e.g. an off-axis approach that wastes fuel correcting cross-track drift).

#### IV. RESULTS

Using the layout definition algorithm proposed in Section III-A, configurations compliant with requirements in Sections III-B-III-C were generated. A population of 1133 configurations was obtained out of 50 million iterations, with a runtime of 5 h 9 min on a 12-core AMD Zen 4 desktop. Configurations in this section are compared against a baseline configuration developed prior to this work, shown in Figure 4. This configuration is also compliant with requirements in Sections III-B-III-C.




Fig. 4. Existing baseline thruster layout

For each configuration, the static analysis method in Section III-D was applied along with force vectors in Figure 2 to find  $t_{on}$ . The results were compared to the total thruster on-time in a dynamic simulation in Figure 5. A Spearman coefficient of 0.632 was noted between the static and dynamic analysis thruster on-times. This result suggests the static method is suitable for predicting mission performance for a close-proximity operations final approach.

The runtime of a single static analysis, including requirements and redundancy verification, was measured at 0.15 seconds. In comparison, a full dynamic simulation required 39 seconds on the same hardware. The static method enables large scale exploration of the configuration parameter space that would otherwise be computationally infeasible.

The correlation was negatively impacted by a set of outliers apparent in Figure 5 - configurations that performed well statically but not dynamically. The common feature of these configurations was asymmetry introduced by thrusters with non-zero azimuth and cant angles on a single principal axis. These configurations could statically reach all the required forces and torques, but produced roll instability in conjunction with the control and allocation algorithms in a dynamic simulation. Small stochastic attitude oscillations were observed across all dynamic simulations, further contributing to data spread. A






Fig. 5. Thruster on-time for static and dynamic configuration simulation (random sample of 100 configurations plotted for readability)

better understanding of the CPOD controller behaviour as well as additional effort in choosing suitable static analysis force vectors are expected to improve correlation in follow-on work.




Fig. 6. Thruster layout obtained through optimisation with GA

Compared to the baseline, the initial population of configurations produced a best normalised on-time of 0.852 when compared in a dynamic simulation. Using DEAP, the best normalised on-time was reduced to 0.843. Varying the GA parameters, a more significant improvement could not be achieved. Instead, further improvements may be possible by allowing the GA to vary thruster positions as well as angles, or by allowing for asymmetric configurations in the thruster layout. Nonetheless, the convergence between both stochastic and GA methods suggests the final configuration, displayed in Figure 6 may be approaching a local optimum within the explored parameter space and geometric restrictions.

## V. DISCUSSION

A stochastic population of requirements-compliant thruster configurations was generated and refined through genetic optimisation. This yielded a 15.7% reduction in thruster on-time relative to an existing baseline for a CPOD task. To contextualise this result and motivate further research, the layout optimisation gains are benchmarked against other losses typical in a cold-gas AOCS.

For low-thrust, small-scale nozzles, manufacturing effects such as surface roughness can lead to a degradation of nozzle efficiency on the order of 20% [20], proportionally increasing propellant demand.

In a purely attitude-hold mission, a reduction in the MIB translates to an equal reduction in total mission impulse. However, this relationship is no longer linear when sustained AOCS firing is required, such as in orbit control or close-proximity operations. Reducing thruster MIB is an active research area, with transient dynamics in particular requiring further research [18, 14].

Finally, optimal thruster allocation is becoming a more significant driver for resilient AOCS design. Unlike systems that achieve redundancy through duplication, resilient design produces highly coupled firing axes. In such cases, linear programming solutions have been proposed with up to 50% propellant savings compared to heuristic thruster allocation [13].

The layout optimisation method presented in this work represents one component of the broader AOCS design process, and alone does not constitute an optimal solution for a given mission. Rather, its strength lies in providing system designers with a selection of high-performing, requirements-compliant configurations at the preliminary design stage, from which a final selection can be made once practical constraints such as mechanical integration are taken into account.

Although the results presented are mission- and configuration specific, they indicate performance improvements comparable in magnitude to other established AOCS design considerations, highlighting the relevance of layout design to AOCS development. This is of particular relevance for emerging propulsion systems, where such gains can be realised at the preliminary design stage prior to any hardware investment and at little computational cost.

## REFERENCES

- [1] Alberto Anselmi et al. “Control Propellant Minimization for the Next Generation Gravity Mission”. In: *Modeling and Optimization in Space Engineering*. Ed. by Giorgio Fasano and János D. Pintér. Vol. 144. Cham: Springer International Publishing, 2019, pp. 1–32. ISBN: 978-3-030-10500-6 978-3-030-10501-3. DOI: 10.1007/978-3-030-10501-3\_1. (Visited on 2026-03-03).
- [2] M. Armano et al. “LISA Pathfinder microneutron cold gas thrusters: In-flight characterization”. In: *Physical review d* 99.12 (June 2019), p. 122003. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.99.122003. (Visited on 2025-07-01).

- [3] Basilisk. *Basilisk*. 2026. URL: <https://hanspeterschaub.info/basilisk/> (visited on 2024-04-16).
- [4] Mingwei Chen et al. "A Design Method for Redundant Microthruster Layout With Allocation Error". In: *International journal of aerospace engineering* 2024.1 (Jan. 2024). Ed. by Chuang Liu, p. 7627257. ISSN: 1687-5966, 1687-5974. DOI: 10.1155/2024/7627257. (Visited on 2025-05-30).
- [5] B. Crawford. "Configuration design and efficient operation of redundant multi-jet systems". In: *AIAA Guidance, Control, and Flight Mechanics Conference*. Princeton, NJ, U.S.A.: American Institute of Aeronautics and Astronautics, Aug. 1969. DOI: 10.2514/6.1969-845. (Visited on 2025-12-03).
- [6] R Delanoë et al. "Overview of the Euclid Reaction Control System and Micro Propulsion Feed Assembly". In: *Space Propulsion 2016*. Rome, IT, May 2016.
- [7] Wigbert Fehse. *Automated Rendezvous and Docking of Spacecraft*. Cambridge Aerospace Series. Cambridge University Press, 2003.
- [8] Thibault L. B. Flinois et al. "MICROTHRUSTER-BASED CONTROL FOR PRECISION POINTING OF NEXT GENERATION SPACE TELESCOPES". In: *Proceedings of the 44th Annual American Astronautical Society Guidance, Navigation, and Control Conference*, 2022. Ed. by Matt Sandnas and David B. Spencer. Vol. 179. Cham: Springer International Publishing, 2024, pp. 1415–1431. ISBN: 978-3-031-51927-7 978-3-031-51928-4. DOI: 10.1007/978-3-031-51928-4\_77. (Visited on 2025-11-12).
- [9] Félix-Antoine Fortin et al. "DEAP: evolutionary algorithms made easy". In: *Journal of machine learning research* 13 (July 2012), pp. 2171–2175.
- [10] Mahdi Ghobadi, Maziar Shafaei, and Mahdi Jafari Nadoushan. "Reliability Approach to Optimal Thruster Configuration Design for Spacecraft Attitude Control Subsystem". In: *Journal of aerospace technology and management* 12 (June 2020), e2320. ISSN: 2175-9146. DOI: 10.5028/jatm.v12.1112. (Visited on 2025-12-03).
- [11] Kaname Kawatsu et al. "On-Orbit Demonstration Plan for an Autonomous Health Monitoring Technology Towards Realization of Resilient Spacecraft Propulsion System". In: *IAF Space Propulsion Symposium*. Sydney, Australia: International Astronautical Federation (IAF), 2025, pp. 631–637. ISBN: 979-8-3313-2938-9. DOI: 10.52202/083090-0070. (Visited on 2026-02-23).
- [12] Jonas Marie et al. "In-orbit experience of the Gaia and LISA Pathfinder cold gas micro-propulsion systems". In: *2018 SpaceOps Conference*. Marseille, France: American Institute of Aeronautics and Astronautics, May 2018. ISBN: 978-1-62410-562-3. DOI: 10.2514/6.2018-2716. (Visited on 2026-03-18).
- [13] Matt Nehrenz and Matt Sorgenfrei. "A Comparison of Thruster Implementation Strategies for a Deep Space Nanosatellite". In: *AIAA Guidance, Navigation, and Control Conference*. Kissimmee, Florida: American Institute of Aeronautics and Astronautics, Jan. 2015. ISBN: 978-1-62410-339-1. DOI: 10.2514/6.2015-0866. (Visited on 2026-02-20).
- [14] Aysu Özden. "PERFORMANCE CHARACTERISTICS OF A COLD GAS THRUSTER WITH REAL GAS EFFECTS". MA thesis. MIDDLE EAST TECHNICAL UNIVERSITY, Sept. 2019.
- [15] Francesco Pacentra. "Thruster layout optimization for a space mission with demanding control requirements". MA thesis. Torino, IT: Politecnico di Torino, 2019.
- [16] Takahiro Sasaki, Yu Nakajima, and Toru Yamamoto. "Proximity Approaches and Design Strategies for Non-Cooperative Rendezvous: V-bar Hopping vs. Spiral Approach". In: *Transactions of the japan society for aeronautical and space sciences* 64.3 (2021), pp. 136–146. ISSN: 0549-3811, 2189-4205. DOI: 10.2322/tjsass.64.136. (Visited on 2026-04-09).
- [17] Till Siebert and Michael Tolstoj. *EXOpod NOVA User Manual*. Mar. 2024.
- [18] R. C. Stechman and John G. Campbell. *Water Electrolysis Satellite Propulsion System*: tech. rep. Fort Belvoir, VA: Defense Technical Information Center, Jan. 1973. DOI: 10.21236/AD0755384. (Visited on 2024-09-11).
- [19] Mak Tafazoli. "A study of on-orbit spacecraft failures". In: *Acta astronautica* 64.2-3 (Jan. 2009), pp. 195–205. ISSN: 00945765. DOI: 10.1016/j.actaastro.2008.07.019. (Visited on 2026-02-26).
- [20] Federico La Torre. "Gas Flow in Miniaturized Nozzles for Micro-Thrusters". In: *Technische Universiteit Delft* (2011).
- [21] Peter J. Wiktor. "Minimum control authority plot - A tool for designing thruster systems". In: *Journal of guidance, control, and dynamics* 17.5 (Sept. 1994), pp. 998–1006. ISSN: 0731-5090, 1533-3884. DOI: 10.2514/3.21301. (Visited on 2026-01-27).
- [22] Koji Yamanaka et al. "Guidance and navigation system design of R-bar approach for rendezvous and docking". In: *17th AIAA International Communications Satellite Systems Conference and Exhibit*. Yokohama, Japan: American Institute of Aeronautics and Astronautics, Feb. 1998. DOI: 10.2514/6.1998-1299. (Visited on 2025-12-27).