Multi-Objective PID Controller Optimization for Quanser Aero 2 Using NSGA-II: Simulation and Hardware Implementation

Ratiba FELLAG(1)*, Mahmoud BELHOCINE(1), Meziane HAMEL(2)
(1) Robotics and Industrial Automation Laboratory, Centre de Développement des Technologies Avancées (CDTA), Algiers, Algeria
(2) Laboratory of Energy and Mechanical Engineering (LEMI), M'Hamed Bougara University, Boumerdès, Algeria
*rfellag@cdta.dz

Abstract: This paper addresses the challenge of optimal Proportional-Integral-Derivative (PID) controller tuning for coupled Multi-Input Multi-Output (MIMO) mechatronic systems and validates the solution on a Quanser Aero 2 dual-rotor helicopter platform. The primary challenge in controlling such systems lies in satisfying conflicting performance objectives, namely, minimizing tracking error while simultaneously reducing control effort to ensure energy efficiency and respect physical actuator limitations. To address this, the controller tuning problem is formulated as a dual-objective optimization problem. The Non-dominated Sorting Genetic Algorithm II (NSGA-II) is applied to obtain a diverse set of Pareto-optimal solutions that characterize the inherent compromise between tracking accuracy and control effort. The methodology is validated through simulation and experimental validation, demonstrating significant performance gains over a conventional PD controller. A rigorous comparative analysis using multiple performance indices (ISE, ITAE, RMSE) provides a quantitative assessment of the different solutions on the Pareto front. The results confirm that the NSGA-II approach provides a systematic and effective method for designing high-performance, robust, and practical controllers for complex MIMO systems. The compromise solution from the Pareto front offers an excellent balance of responsiveness, stability, and energy efficiency in both simulation and real-world experiments.

Keywords: PID Control, Multi-Objective Optimization, NSGA-II, MIMO Systems, Quanser Aero 2, Hardware Implementation, Evolutionary Algorithms.

1. INTRODUCTION

The control of Multi-Input Multi-Output (MIMO) systems is a fundamental challenge in modern engineering, with critical applications in aerospace, robotics, and process control [1]. Unlike Single-Input Single-Output (SISO) systems, MIMO systems are characterized by complex cross-coupling between inputs and outputs, where a single control action can affect multiple system states simultaneously. This inherent coupling, combined with often conflicting performance requirements, makes controller design a non-trivial task. Industrial controllers must typically achieve a delicate balance: minimize tracking error, reduce settling time, limit overshoot to prevent system stress, and constrain control effort to conserve energy and avoid actuator saturation [2]. Traditional single-objective optimization techniques are often insufficient, as they require an a priori weighting of objectives, which can obscure the true nature of the performance trade-offs and lead to suboptimal designs [3].

Despite the development of advanced control theories, the Proportional-Integral-Derivative (PID) controller remains the workhorse of the industry, used in over 90% of control loops [4, 5]. Its enduring popularity is due to its simple structure, intuitive tuning parameters, and proven robustness. However, tuning PID controllers for optimal performance in complex MIMO systems remains a significant hurdle [6]. Classical tuning methods, such as Ziegler-Nichols, often produce aggressive gains, leading to excessive overshoot and control activity, which is unacceptable for systems with sensitive dynamics or strict actuator constraints [7, 8].

The Quanser Aero 2, a two-degree-of-freedom (2-DOF) laboratory helicopter, epitomizes these challenges. Its coupled pitch-yaw dynamics, inherent nonlinearities, and motor voltage saturation limits (±24V) make it an ideal testbed for advanced MIMO control strategies [9]. Effective control requires managing the aerodynamic coupling where the action of one rotor significantly influences the motion on the other axis [10].

To address these challenges, Multi-Objective Optimization (MOO) has emerged as a powerful paradigm. MOO algorithms seek not a single solution, but a set of Pareto-optimal solutions [11, 12]. Each point on the resulting Pareto front represents a non-dominated solution, where no single objective can be improved without degrading at least one other. This provides designers with a complete spectrum of optimal trade-offs, enabling an informed decision based on specific operational priorities [13].

Among MOO techniques, evolutionary algorithms (EAs) are particularly well-suited due to their global search capabilities and ability to handle complex, non-convex search spaces. The Non-dominated Sorting Genetic Algorithm II (NSGA-II) is a state-of-the-art EA that has become a benchmark in the field [14]. It efficiently finds a well-distributed set of Pareto-optimal solutions by using elitism and a parameter-less diversity preservation mechanism.

This paper presents an extended, in-depth investigation into the application of NSGA-II for the optimal tuning of decentralized PID controllers for the Quanser Aero 2. Building upon preliminary conference work [15], this journal article provides: 1) an expanded analysis including a new simulation scenario; 2) a new section on hardware implementation and validation; 3) a rigorous quantitative comparative analysis; and 4) a detailed discussion of the results, including Pareto front quality and robustness. The goal is to provide a complete framework, from theoretical formulation to practical validation, for applying multi-objective optimization to a real-world mechatronic system.

2. SYSTEM DESCRIPTION AND MODELING

The Quanser Aero 2 is a laboratory-scale helicopter designed for research in control systems and mechatronics [9]. The system, shown in Fig. 1, features a cross-shaped frame mounted on a pivot, allowing for two degrees of freedom: rotation around the horizontal axis (pitch, θ) and rotation around the vertical axis (yaw, ψ). Actuation is provided by two DC motors driving propellers [16]. The front motor primarily controls the pitch angle, while the rear motor controls the yaw angle. High-resolution optical encoders measure the pitch and yaw angles for feedback control.

The linearized state-space model around the equilibrium is given by:

$$\begin{cases} \dot{x} = Ax + Bu \\ y = Cx \end{cases} \quad (1)$$

Where $x = [\theta, \psi, \dot{\theta}, \dot{\psi}]^T$ is the state vector, $u = [V_p, V_y]^T$ is the input vector (voltages to the

Fig. 1 Quanser Aero 2 platform [9] .

pitch and yaw motors), and $y = [\theta, \psi]^T$ is the output vector.

$$A = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ -K_{sp}/J_p & 0 & -D_p/J_p & 0 \\ 0 & 0 & 0 & -D_y/J_y \end{bmatrix} \quad (2)$$

$$B = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ D_t * K_{pp}/J_p & D_t * K_{py}/J_p \\ D_t * K_{yp}/J_y & D_t * K_{yy}/J_y \end{bmatrix} \quad (3)$$

$$C = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \quad (4)$$

The model parameters are listed in Table 1 [17].

3. MULTI-OBJECTIVE OPTIMIZATION FORMULATION

The following section defines the PID controller tuning as a multi-objective optimization problem by outlining the decision variables, objective functions, and constraints. It also introduces the NSGA-II algorithm selected to obtain the solution.

Table 1. Dynamic parameters of the Aero 2 [10].

Symbol	Description	Value
J_p	Pitch axis inertia	0.0232 $Kg.m^2$
J_y	Yaw axis inertia	0.0238 $Kg.m^2$
D_p	Pitch axis damping	0.0020 $N.m/V$
D_y	Yaw axis damping	0.0019 $N.m/V$
K_{pp}	Pitch thrust gain from front rotor	0.0032 N/V
K_{py}	Pitch thrust gain from rear rotor	0.0014 N/V
K_{yy}	Yaw thrust gain from rear rotor	0.0061 N/V
K_{yp}	Yaw thrust gain from front rotor	-0.0032 N/V
K_{sp}	Pitch stiffness	0.0074$N.m/V$
D_t	Distance btw pivot and center of rotor	0.1674m

A. Problem Formulation

A decentralized control structure with two independent PID controllers is used. The parameter vector to be optimized k, contains the six gains:

$$k = [k_{p\theta}, k_{i\theta}, k_{d\theta}, k_{p\psi}, k_{i\psi}, k_{d\psi}] \quad (5)$$

The problem is to find the vector k that minimizes two objective functions [15], and that respects the mentioned constraints for pratical implementation :

Objective 1 - Tracking Performance (f_1): This function measures the system's ability to follow a reference trajectory. It is defined as the Integrated Squared Error (ISE) with a temporal weighting term to penalize persistent errors more heavily.

$$f_1 = \int_0^T w(t)\left[e_p^2(t) + e_y^2(t)\right]dt \quad (6)$$

where w(t) = 1 + 0.1t provides an increasing penalty for steady-state errors, and $e_p(t)$, $e_Y(t)$ are the pitch and yaw tracking errors, respectively.

Objective 2 - Control Effort (f_2): \)): This objective quantifies the energy consumption and aggressiveness of the control signals. It is formulated as a combined measure of control signal variance and peak values to promote smooth and energy-efficient control actions.

$$f_2 = Var(u_p) + Var(u_y) + \alpha_{peak}\left(\left|u_p\right|_{\max} + \left|u_y\right|_{\max}\right) \quad (7)$$

Where $Var(u)$ is the variance of the control signal, $(|u|_{\max})$ is the peak absolute value, and α_{peak} is a weighting factor to balance the components.

Constraints: Box constraints are applied to the PID gains to ensure they remain within a practical and physically realizable range. The search space is defined as follows:

- Proportional gains:
  - $k_{p\theta} \in [10,180]$, $k_{p\psi} \in [5,80]$.
- Integral gains:
  - $k_{i\theta} \in [0.1,15]$, $k_{i\psi} \in [0.1,12]$.
- Derivative gains:
  - $k_{d\theta} \in [5,120]$, $k_{d\psi} \in [3,70]$.

B. The NSGA-II Algorithm

The Non-dominated Sorting Genetic Algorithm II (NSGA-II) is an elitist multi-objective evolutionary algorithm well-reputed for its efficiency and robust performance in finding a well-distributed set of Pareto-optimal solutions [14]. It operates on a population of candidate solutions (sets of PID gains) and iteratively evolves them toward the optimal trade-off front. The general workflow is depicted in Fig. 2.

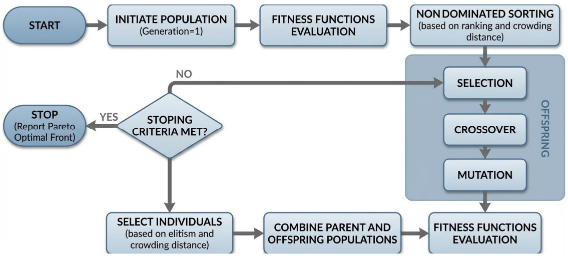

Fig. 2. NSGA II Algorithm flowchart [18]

The algorithm begins by initializing a random population of individuals within the predefined gain constraints. Each individual is then evaluated by running a closed-loop system simulation to calculate the objective function values, **(f_1)** and **(f_2)**. The core of NSGA-II lies in its selection mechanism. First, the population is sorted into different non-domination fronts. Solutions that are not dominated by any other solution are assigned to the first front. This process is repeated to form subsequent fronts.

To maintain diversity and prevent premature convergence to a small region of the search space, a crowding distance metric is calculated for each solution within a front [19]. This metric estimates the density of solutions, favoring individuals in less-crowded regions. An offspring population is then generated using genetic operators: binary tournament selection (which considers both rank and crowding distance), simulated binary crossover, and polynomial mutation [14]. Finally, an elitist strategy combines the parent and offspring populations, and the best individuals are selected for the next generation, ensuring that high-quality solutions are preserved. This iterative process continues until a termination criterion, such as a maximum number of generations, is met.

4. SIMULATION RESULTS

The NSGA-II was configured with a population of 100 individuals and run for 200 generations. Crossover and mutation probabilities were 0.90 and 0.12, respectively, with distribution indices of 20 for both operators. The reference signals were phase-shifted square waves for pitch and yaw, respectively, with a disturbance of magnitude 0.4 rad and a duration of 2s introduced at 10s to test rejection capabilities. The performance of the NSGA-II-tuned controllers was compared with that of a baseline PD-Ref controller from the constructor [9].

After 200 generations, the algorithm converged to the well-distributed Pareto front shown in Fig. 3. The front clearly illustrates the trade-off between tracking error **(f_1)** and control effort **(f_2).** Three representative solutions were selected for analysis: PID-f1 (the best tracking), PID-f2 (the lowest effort), and PID-Comp (a balanced compromise).

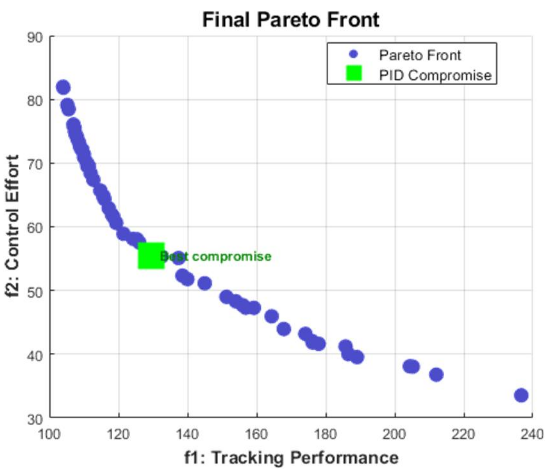

Fig. 3. Simulated pitch response with disturbance

The simulated square wave tracking responses for pitch and yaw tracking are shown in Figs. 4 and 5, respectively. The PID-f1 controller demonstrates the fastest response and quickest disturbance rejection, closely following the reference. In contrast, PID-f2 is noticeably slower and more sluggish. The PID-Comp controller provides a response that is nearly as fast as PID-f1 but with less overshoot.

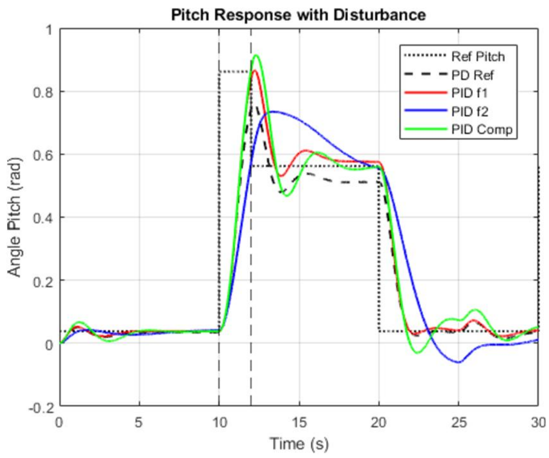

Fig. 4. Simulated pitch response with disturbance

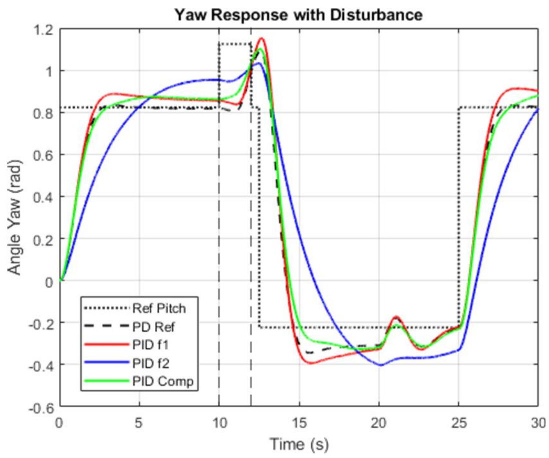

Fig. 5. Simulated yaw response with disturbance

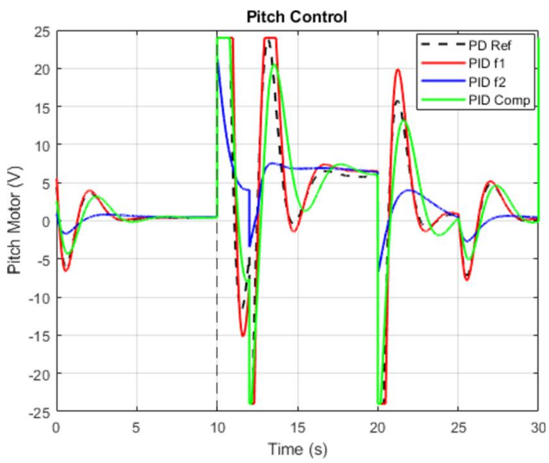

Fig. 6. Simulated controllers' comparison for pitch tracking

This behavior is further explained by the control signals shown in Figs. 6 and 7. The PID-f1 controller utilizes aggressive, high-amplitude control voltages to achieve its fast tracking, which could lead to actuator saturation and wear in a real system. The PID-f2 controller, by design, utilizes minimal voltage, resulting in a slower response. The PID-Comp controller's signals are more moderate, balancing aggressive action with energy conservation.

Quantitative analysis in Tables 2 and 3 confirms these observations. For pitch tracking (Table 2), PID-f1 has a slightly higher ISE (0.8706) than the PD-Ref (0.8184) but a lower ITAE (28.47 vs 29.45), indicating faster settling. PID-f2 has the worst performance across all metrics. For yaw tracking (Table 3), the trends are similar, with PID-Comp (ISE 3.4230) performing comparably to the baseline PD-Ref (ISE 3.3055) but with the added benefit of integral action for steady-state error elimination. These results highlight the effectiveness of the multi-objective approach in generating a spectrum of controllers, with PID-Comp emerging as a strong candidate for practical implementation.

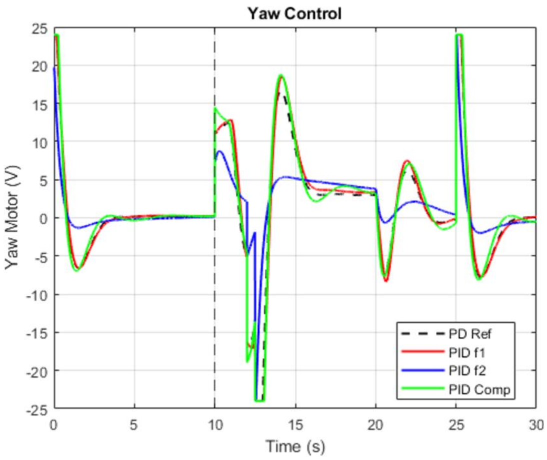

Fig. 7. Simulated controllers' comparison for yaw tracking

Table 2. Pitch simulation performance comparison

Controller	ISE	ITAE	RMSE
PD-ref	0.8184	29.45	0.1652
PID-f1	0.8706	28.47	0.1704
PID-f2	1.2159	51.54	0.2013
PID-comp	0.9205	34.56	0.1752

Table 3. Yaw simulation performance comparison

Controller	ISE	ITAE	RMSE
PD-ref	3.3055	75.37	0.3320
PID-f1	3.5718	85.95	0.3451
PID-f2	5.1542	132.4	0.4145
PID-comp	3.4230	79.83	0.3378

5. Experimental Validation

To validate the simulation results and assess real-world performance, the three selected PID controllers and the baseline PD-Ref were implemented on the Quanser Aero 2 hardware. The experimental setup is pictured in Fig. 8, and the obtained hardware implementation responses to the same square wave reference and disturbance used in simulations are illustrated in Figs. 9 and 10.

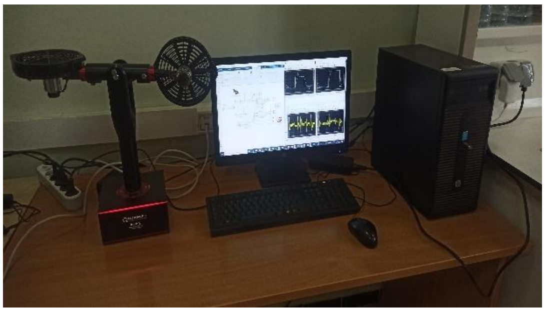

Fig. 8. Experimental setup at lab.

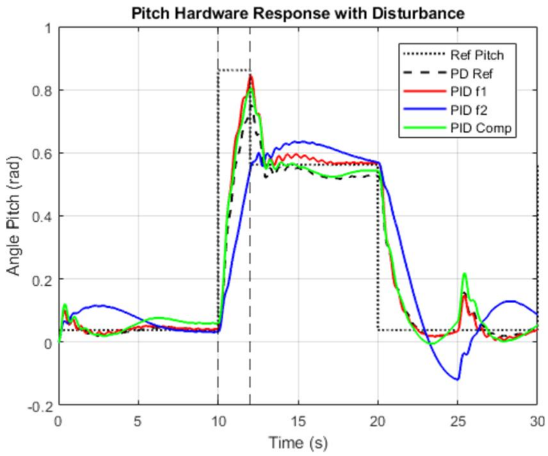

Fig. 9. Experimental pitch response with disturbance

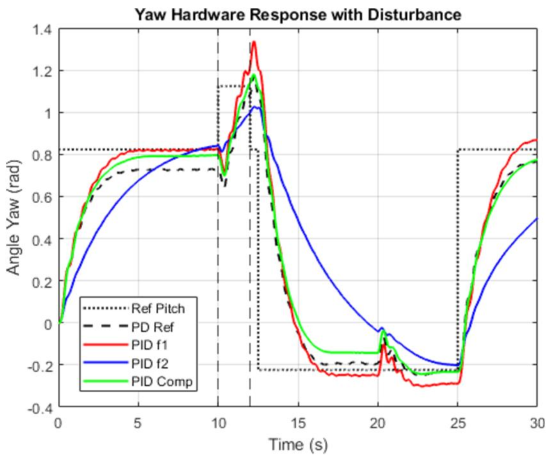

Fig. 10. Experimental yaw response with disturbance

The experimental results largely corroborate the simulation results. The PID-f1 controller again provides the most responsive tracking, while PID-f2 is the most conservative. The PID-Comp controller maintains its excellent balance of speed and stability. The experimental control signals (Figs. 11 and 12) are noisier than in simulation, which is expected due to sensor noise, unmodeled dynamics, and environmental factors. Notably, the aggressive nature of PID-f1 is even more apparent in the hardware implementation, with control signals frequently chattering and approaching the saturation limits.

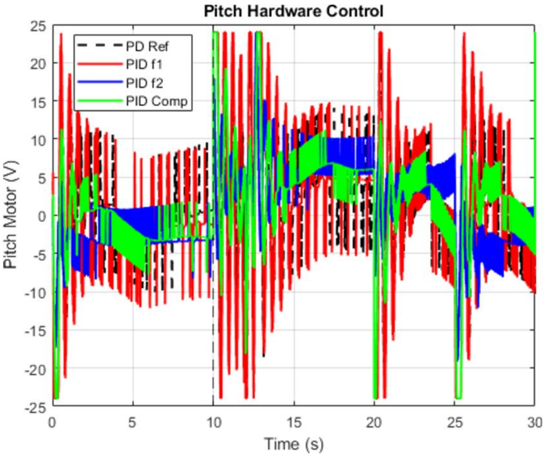

Fig. 11. Experimental comparison of controllers for pitch tracking

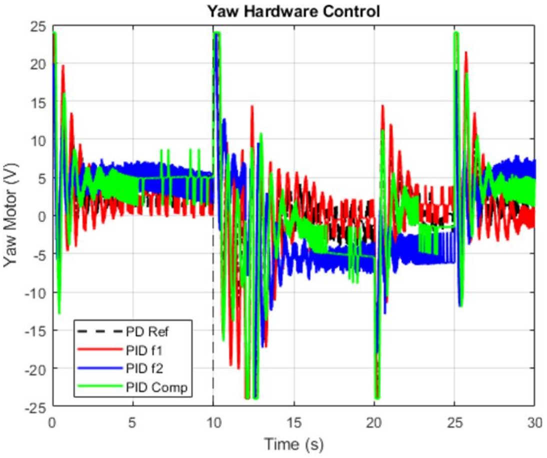

Fig. 12. Experimental comparison of controllers for yaw tracking

The quantitative performance metrics from the hardware tests are presented in Tables 4 and 5. For pitch control (Table 4), PID-f1 achieves the lowest ISE (0.5414) and ITAE (24.20), outperforming the PD-Ref and confirming its superior tracking ability in practice. The PID-Comp controller (ISE 0.5554) performs nearly as well, offering a robust alternative with significantly less control effort. For yaw control (Table 5), PID-f1 again shows strong performance (ITAE 81.04), closely followed by the PD-Ref (ITAE 84.09). The PID-Comp controller (ITAE 90.07) remains a viable choice, especially considering its smoother control action.

The consistency between simulation and experimental data validates the linearized model and the effectiveness of the NSGA-II optimization process for a real-world application.

Table 4. Experimental performance comparison for pitch

Controller	ISE	ITAE	RMSE
PD-ref	0.5799	26.59	0.1391
PID-f1	0.5414	24.20	0.1344
PID-f2	1.0874	49.71	0.1904
PID-comp	0.5554	29.01	0.1361

Table 5. Experimental performance comparison for yaw

Controller	ISE	ITAE	RMSE
PD-ref	3.1618	84.09	0.3247
PID-f1	3.4238	81.04	0.3378
PID-f2	6.4867	166.8	0.4650
PID-comp	3.3381	90.07	0.3336

6. Discussion

The results validate the multi-objective optimization approach for this mechatronic system, demonstrating strong correlation between simulation and experiments. The performance hierarchy (PID-f1 > PID-Comp > PID-f2) remained consistent across both domains. For pitch control, PID-Comp achieved ISE values of 0.9205 (simulation) and 0.5554 (experiment), confirming the linearized model's fidelity and the robustness of optimized gains during hardware transfer.

Experimental control signals (Figs. 11, 12) exhibited high-frequency noise absent in simulation (Figs. 6, 7), attributable to sensor noise, unmodeled nonlinearities (friction, aerodynamic drag), and actuator dynamics. The aggressive PID-f1 controller amplified this noise due to high derivative gains, highlighting practical limitations of tracking-optimized controllers.

Disturbance rejection tests confirmed closed-loop stability across all optimized controllers. The integral action ensured zero steady-state error, essential for precision applications [20]. PID-Comp demonstrated excellent disturbance rejection without excessive control effort, proving suitable for real-world conditions with unexpected perturbations.

7. Conclusion

This paper presents a comprehensive framework for tuning decentralized PID controllers for the coupled MIMO Quanser Aero 2 system using NSGA-II multi-objective optimization. The approach generates a Pareto front that explicitly maps the trade-off between tracking precision and control effort. Simulation and hardware validation confirmed that the selected controllers, PID-f1 (high performance), PID-f2 (low effort), and PID-Comp (balanced), behaved as predicted, with PID-Comp delivering near-optimal tracking with practical actuation levels.

This validated case study bridges the gap between simulation and practice for evolutionary optimization applied to real-world MIMO systems. The methodology is broadly applicable to industrial PID-controlled platforms, offering reduced tuning time and enhanced performance. Future work will incorporate robustness as a third objective, benchmark alternative optimization algorithms, and extend the framework to more complex systems such as fully actuated UAVs.

References

[1] J. Maciejowski, "Multivariable Feedback Design Addison," ed: Wesley Publishing Company-1989, 1989.

[2] R. K. Mudi and N. R. Pal, "A robust self-tuning scheme for PI-and PD-type fuzzy controllers," IEEE Transactions on fuzzy systems, vol. 7, no. 1, pp. 2-16, 2002.

[3] C. M. Fonseca and P. J. Fleming, "An overview of evolutionary algorithms in multiobjective optimization," Evolutionary computation, vol. 3, no. 1, pp. 1-16, 1995.

[4] K. J. Astrom, "PID controllers: theory, design, and tuning," The international society of measurement and control, 1995.

[5] J. D. Rojas, O. Arrieta, and R. Vilanova, Industrial PID controller tuning. Springer, 2021.

[6] H. V. H. Ayala and L. dos Santos Coelho, "Tuning of PID controller based on a multiobjective genetic algorithm applied to a robotic manipulator," Expert Systems with Applications, vol. 39, no. 10, pp. 8968-8974, 2012.

[7] J. G. Ziegler and N. B. Nichols, "Optimum settings for automatic controllers," Transactions of the American society of mechanical engineers, vol. 64, no. 8, pp. 759-765, 1942.

[8] H. A. Mintsa, G. E. Eny, N. Senouveau, and R. M. A. Nzué, "Optimal tuning PID controller gains from ziegler-nichols approach for an electrohydraulic servo system," Journal of Engineering Research and Reports, vol. 25, no. 11, pp. 158-166, 2023.

[9] Quanser, Quanser Aero 2 laboratory guide. 2022.

[10] R. Fellag and M. Belhocine, "Comparative analysis of PID, fuzzy PID, and ANFIS controllers for 2-DOF helicopter trajectory tracking: simulation and hardware implementation," Archive of Mechanical Engineering, vol. vol. 71, no. No 3, pp. 323-349, 24.09.2024 2024, doi: 10.24425/ame.2024.151331.

[11] K. Deb, "Multi-objective optimisation using evolutionary algorithms: an introduction," in Multi-objective evolutionary optimisation for product design and manufacturing: Springer, 2011, pp. 3-34.

[12] C. A. C. Coello, G. B. Lamont, and D. A. V. Veldhuizen, Evolutionary algorithms for solving multi-objective problems. Springer, 2007.

[13] J. L. J. Pereira, G. A. Oliver, M. B. Francisco, S. S. Cunha Jr, and G. F. Gomes, "A review of multi-objective optimization: methods and algorithms in mechanical engineering problems," Archives of Computational Methods in Engineering, vol. 29, no. 4, pp. 2285-2308, 2022.

[14] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan, "A fast and elitist multiobjective genetic algorithm: NSGA-II," IEEE transactions on evolutionary computation, vol. 6, no. 2, pp. 182-197, 2002.

[15] R. Fellag, M. Belhocine , and M. Hamel, "NSGA-II Multi-Objective PID Optimization for Quanser Aero 2 " presented at the International Conference on Computational Engineering, Artificial Intelligence and Smart systems IC2EAIS2 2025, Djanet, Algeria, 29-31 October 2025, 2025.

[16] R. Fellag and M. Belhocine, "2-DOF Helicopter Control Via State Feedback and Full/Reduced-Order Observers," in 2024 2nd International Conference on Electrical Engineering and Automatic Control (ICEEAC), 2024: IEEE, pp. 1-6.

[17] R. Fellag, M. Belhocine, F. Demim, and M. Hamel, "Robust Trajectory Tracking of a 2-DOF Helicopter Using Passivity-Based Sliding Mode Control," presented at the National conference on computational engineering, artificial intelligence and smart systems nc2eais2 2024 10-12 December 2024; , Tamanrasset, Algeria;, 2024. [Online]. Available: http://dlibrary.univ-boumerdes.dz:8080/handle/123456789/15097.

[18] N. L. Lukić, M. Božin-Dakić, J. A. Grahovac, J. M. Dodić, and A. I. Jokić, "Multi-objective optimization of microfiltration of baker's yeast using genetic algorithm," Acta Periodica Technologica, no. 48, pp. 211-220, 2017.

[19] H. Ma, Y. Zhang, S. Sun, T. Liu, and Y. Shan, "A comprehensive survey on NSGA-II for multi-objective optimization and applications," Artificial Intelligence Review, vol. 56, no. 12, pp. 15217-15270, 2023.

[20] G. Liu and S. Daley, "Optimal-tuning PID control for industrial systems," Control Engineering Practice, vol. 9, no. 11, pp. 1185-1194, 2001.