

ADAPTIVE FLIGHT CONTROL LAWS APPLIED TO AIR REFUELING WITH RECEIVERS EQUIPPED WITH RECEPTACLE OUT OF THE PLANE OF SYMMETRY

Víctor Martínez Jiménez
Flight Dynamics & Control Laws Department
Airbus Defence and Space
Avda. de John Lennon, s/n, 28906 Getafe, Spain
Phone : +34 914433057
e-mail : victor.martinez@airbus.com

Francisco Asensio Nieto
Flight Dynamics & Control Laws Department
Airbus Defence and Space
Avda. de John Lennon, s/n, 28906 Getafe, Spain
Phone: +34 914431211
e-mail: francisco.asensio@airbus.com

Ignacio Romo Martínez
Flight Dynamics & Control Laws Department
Airbus Defence and Space
Avda. de John Lennon, s/n, 28906 Getafe, Spain
Phone : +34 914437672
e-mail : ignacio.romo@airbus.com

Alberto Sanz de Blas
Flight Dynamics & Control Laws Department
Airbus Defence and Space
Avda. de John Lennon, s/n, 28906 Getafe, Spain
Phone : +34 914431052
e-mail : alberto.sanz@airbus.com

Diego Lodares Gomez
Flight Dynamics & Control Laws Department
Airbus Defence and Space
Avda. de John Lennon, s/n, 28906 Getafe, Spain
Phone : +34 914437684
e-mail : diego.lodares@airbus.com

Abstract

This paper addresses the design and certification of the Airbus Automatic Air-to-Air Refueling (A3R) flight control laws for receivers whose receptacle is located outside the plane of symmetry, with the F-15 as the only such aircraft in active service. The strong receiver sidewash gradient close to the receptacle, combined with the long and flexible structure of the refueling boom, makes it difficult for classical PID controllers and prevents reliable contact. The proposed solution couples an open-loop *Direct Link* compensation— derived from CFD trim data — with a *Model Reference Adaptive Controller (MRAC)* that updates the feed-forward gain in real time by means of a fuzzy neural network whose weights are driven by a Lyapunov-based update law. Asymptotic stability of the tracking error is proved using Barbalat's lemma. The certification strategy combines a Lyapunov-based rationale, a linear analysis of the worst-case adaptive controller, an extensive non-linear clearance campaign and two dedicated flight-test campaigns. The system was formally certified in 2024 by the Spanish certification authority (INTA) after 360 automatic contacts with F-15SG receivers.

Keywords: boom, air-to-air refueling, adaptive control, MRAC, F-15

1. INTRODUCTION

This section introduces the air-to-air refueling boom system and the roadmap until its automation. The flying boom of the Airbus Multi Role Transport Tanker (MRTT) is a rigid body equipped with two ruddervators positioned in a V-shaped configuration that allow the boom operator to fly the boom in pitch and roll axes and there is a telescopic probe that is extended for contact.

The first air-to-air refueling boom system entered into service in the late 1940s, and the first Fly By Wire (FBW) boom system went into service in 1981. The next major step was the introduction by Airbus of the Boom Enhanced Vision System (BEVS) on the A330-MRTT, that allows the boom operator to be seated with the rest of the crew while keeping a full view of the operation through a low-latency vision system. The first A330-MRTT was certified in 2010 and entered into service with the Royal Australian Air Force (RAAF) in 2011 (Figure 1).



Figure 1 – Royal Australian Air Force A330-MRTT tanker

1.1 Air to Air Refueling manual operation

Manual operation of the boom is performed by means of a Flight Control Stick (FCS), which commands the boom in pitch and roll, and a Telescopic Control stick (TCS), which commands extension and retraction of the boom probe, as depicted in Figure 2:

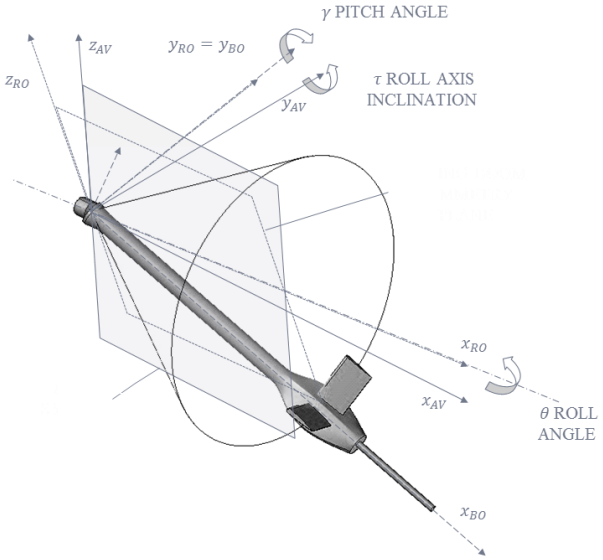


Figure 2 –Boom axes scheme

One of the most demanding control challenges of the boom is its structural flexibility: the natural frequencies of the first flexible modes lie close to those of the rigid-body control bandwidth. A weighted blending of sensors placed at different stations of the boom is used to attenuate the first flexible mode (see REF [1] for more details). Because the boom modal characteristics depend on a large number of operating conditions — flight condition, telescopic length, fuel inside the boom, receptacle type, pitch and roll angles —it has been observed in service some residual zones where first flexible mode is not completely suppressed. To improve this, the **Elastic Modes Adaptive Cancellation (EMAC)** algorithm was developed; it identifies the modal characteristics in real time and cancels them in the loop (REF [2], [3]). EMAC is also able to transition the flight control laws automatically whenever a mismatch between the commanded control mode and the physical state of the boom is detected (REF [4]).

1.2 Automatic Air to Air Refueling

Work on the automation of the refueling operation started in 2015. Two years later, the first fully automatic contact in history was achieved between an A310-MRTT (the Airbus flying test bed) and a Portuguese Air Force F-16A. The first automatic contact with a heavy receiver — an A310-MRTT acting as tanker and a Royal Australian Air Force (RAAF) A330-MRTT as receiver — followed in 2018. In 2019 the first commercial contract for A3R was signed with the RSAF, leading in 2022 to the certification of automatic operation with all F-16 variants and with the A330-MRTT as receivers. This is referred to in the rest of the paper as A3R Phase 1. The next step, addressed in this paper, was the certification of A3R for the F-15SG. The main programme milestones are summarized in Figure 3:

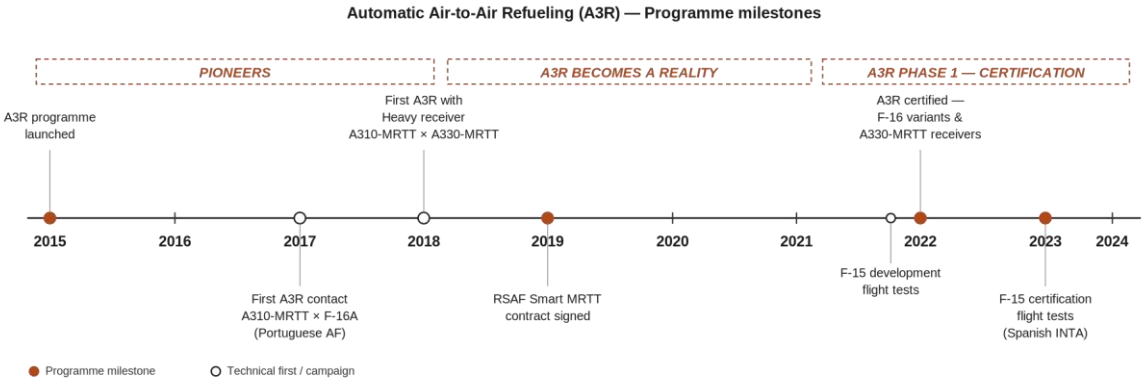


Figure 3 –A3R milestones

1.3 Unconventional receptacles

Since the standardization of the receptacle installation (Universal Aerial Refueling Receptacle Slipway Installation, UARRSI), the receptacle of most receivers is located in the plane of symmetry of the aircraft and its doors do not generate aerodynamic disturbances that affect the boom contact operation. A few receivers, however, retain an unconventional receptacle layout because structural or systems constraints preclude the standard installation. Two families can be identified:

- **Doors that open outwards** generate a complex aerodynamic flow field around the receptacle once they are deployed. The Boeing E-3 AWACS (Figure 4) and the B-52 Stratofortress are representative examples.



Figure 4– Receptacle doors unconventional opening

- **Receptacles located out of the plane of symmetry** — historically present on the F-84G Thunderjet and on early conversions of the B-29/B-50. Today the only receiver in active service with a wing-edge receptacle is the F-15, in all of its variants (Figure 5).



Figure 5 – Wing-edge receptacle layout

2. OUT OF PLANE OF SYMMETRY RECEPTACLES

2.1 Aerodynamic and Flexibility Issues

Early in-service experience with wing-edge receptacles showed that performing a successful contact was very difficult and that the boom-operator training requirements were significantly higher than for conventional receivers. Among the receivers that ever flew with this layout, only the F-15 remains in active service.

Receptacles located close to the leading edge of the wing sit inside a strong aerodynamic flow field. As the boom probe approaches the receptacle, this flow field induces a lateral force on the probe that grows rapidly with proximity, pushing the boom towards the canopy of the receiver. Because the boom is a long flexible structure (Section 1.1), the same lateral force also bends the probe, which in turn degrades the boom-tip position estimation produced by the Computer Vision (CV) subsystem. The effect is more severe at high dynamic pressure.

Figure 6 illustrates this phenomenon at high altitude and high speed for two values of the vertical distance in receiver Z axis (note that X-axis is positive towards the tail of the aircraft, Z-axis positive upwards and Y-axis positive towards the right wing) between the boom nozzle and the receptacle, ΔZ = 5.3 m and ΔZ = 1.0 m. The contour lines represent the local flow-field magnitude in m/s, the red arrows its direction in the X-Y plane of the receiver, the black line the front border of the F-15 fuselage and the yellow circle the position of the receptacle. The strong flow field gradient that is developed as the probe approaches the receptacle is the central control challenge addressed by this paper.

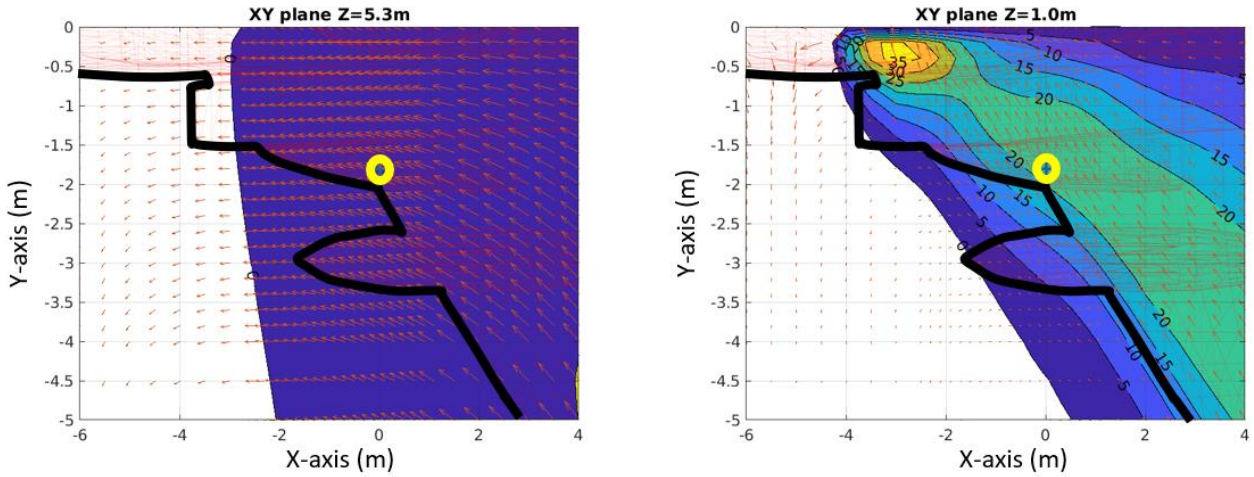


Figure 6 – Flow field in X-Y receiver plane evolution with respect to relative distance

2.2 Need for an open-loop compensation: the Direct Link

The PID-based outer-loop controller originally tuned for small and heavy receivers was retuned for tracking the F-15 receptacle. Although simulations showed acceptable tracking performance, the extension-to-contact phase revealed significant lateral misalignments as soon as the probe was extended towards the receptacle.

The fundamental reason is bandwidth: the lateral aerodynamic force grows so rapidly when decreasing ΔTL (nozzle-receptacle distance in telescopic sense) that a feedback controller of PID type cannot respond fast enough to compensate it without destabilising the loop. The solution adopted is to add a feed-forward open-loop term, denoted δ_{OL}, that injects the asymmetric ruddervator deflection required to balance the boom aligned with the receptacle as a function of the relative receiver-receptacle position. This open-loop path has no feedback delay and provides the necessary compensation bandwidth. The differential ruddervator values used to build the open-loop schedule depend on the flight condition, the receptacle position inside the attitude envelope, and the longitudinal distance measured in X-boom axis (see Fig 2) between the receptacle and the boom.

The resulting baseline outer-loop control law is the legacy PID controller tuned for the F-15 augmented with the open-loop Direct Link during the extension-to-contact phase. The Direct-Link schedule as a function of the longitudinal distance measured in X-boom axis between the boom-tip and the receptacle for one representative receptacle position is shown in Figure 7, where each subplot corresponds to a different flight level and Mach/dynamic-pressure combination:

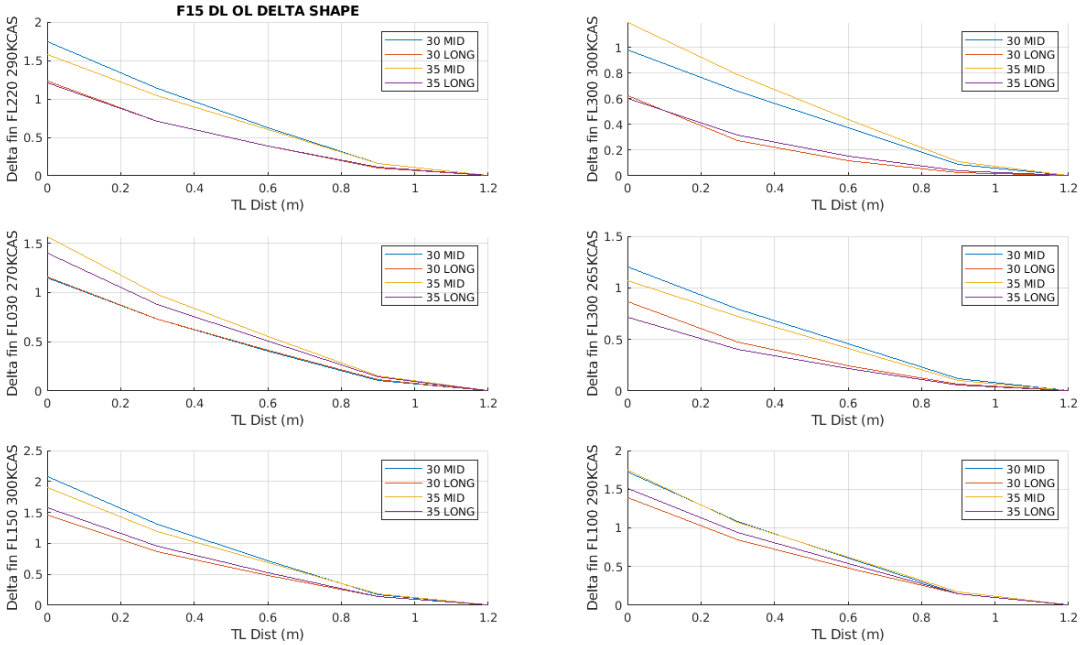


Figure 7 – Direct link required in the whole flight envelope for a specific condition

2.3 Model vs reality

Like any feed-forward compensation, the success of the Direct Link relies on the fidelity of the model that has been used to generate it. In the present case the uncertainty is significant: the detailed geometry of the F-15 was not available, and the CFD aerodynamic dataset had to be built from a number of hypotheses about the airframe shape. The only way to evaluate the accuracy of these hypotheses is through flight test, and the only way to recover the nominal tracking performance in the presence of the resulting modelling errors is to wrap the open-loop schedule in a real-time adaptive law. The design of the adaptive law to handle this uncertainty is described in section 3.

3. APPROACH

3.1 Flight Control Laws Structure

As discussed in section 2, the baseline flight control law consists of a classical PID feedback term acting on the boom vs receptacle alignment error (as for SMALL and HEAVY receiver type FCL), augmented with a feed-forward Direct Link (DL) that injects an open-loop ruddervator asymmetric deflection computed from F15 upwash aerodynamic trim data (δ_{OL}). The direct link compensates the F-15 sidewash gradient that the feedback path alone cannot achieve. The resulting structure is shown in Figure 8:

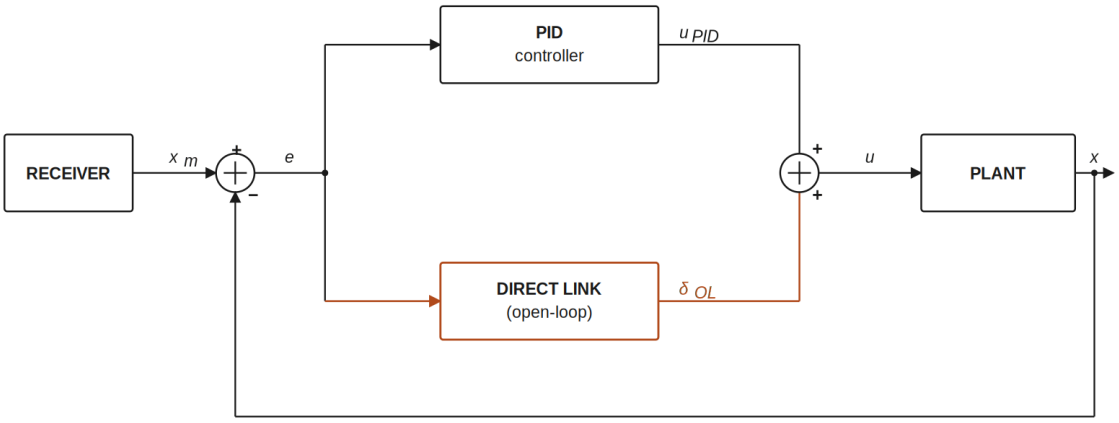


Figure 8 – Baseline flight control law: PID feedback augmented with an open-loop Direct Link

3.2 Adaptive controller

The open-loop direct-link compensation introduced in Section 3.1 relies on an aerodynamic dataset obtained from CFD simulations. Since the geometry of the receiver was not available in detail, the predicted flow-field gradient is subject to some uncertainty. To deal with it, regardless of the mismatch between model and reality, a real-time adaptive law is added on top of the baseline controller. The chosen architecture is a **direct Model Reference Adaptive Control (MRAC)** that adjusts the open-loop compensation so as to cancel the unknown sidewash disturbance acting on the boom plant. The overall scheme is shown in Figure 9:

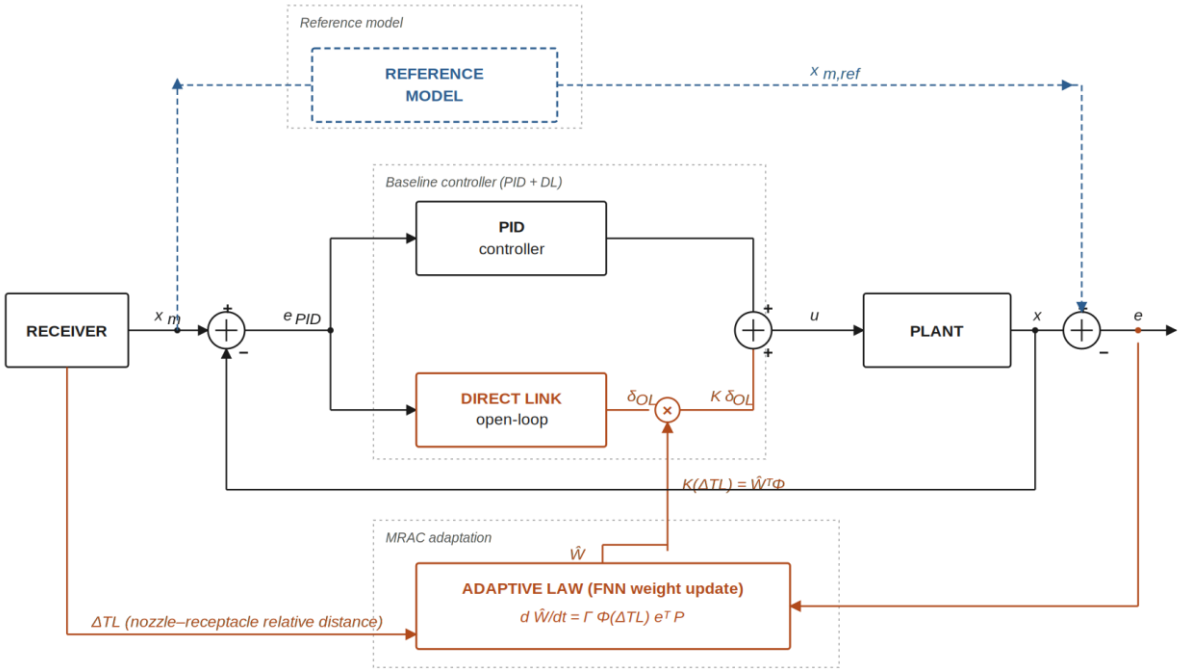


Figure 9 – MRAC architecture

3.2.1 Plant model and disturbance representation

Let $\mathbf{x}(t) \in \mathbb{R}^n$ denote the state of the boom plant and $u(t) \in \mathbb{R}$ the commanded asymmetric ruddervator deflection. The dynamics in the vicinity of the contact point are written in standard state-space form:

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}u, \quad \mathbf{y} = \mathbf{C}\mathbf{x} + \mathbf{D}u \tag{1}$$

The aerodynamic flow field generated by the receiver acts on the boom structure as an unknown additive input matched with the control channel. Denoting this disturbance by $f_{sw}(\mathbf{x}, t)$, the perturbed dynamics become:

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}[u + f_{sw}(\mathbf{x}, t)] \tag{2}$$

The goal of the adaptive law is to obtain a control input $u_{ad}(t)$ that cancels the unknown disturbance, so that the closed-loop response matches that of a disturbance-free reference plant:

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}[u + f_{sw} + u_{ad}] \tag{3}$$

3.2.2 Disturbance parametrization

Flight test evidence shows that the asymmetric ruddervator deflection required to balance the boom aligned with the receptacle in the presence of F-15 sidewash varies smoothly with the relative distance between the boom nozzle and the receptacle, ΔTL. The required compensation can therefore be expressed as an unknown scalar function of ΔTL. By virtue of the universal approximation theorem, this function can be expressed as:

$$K(\Delta TL) = \Sigma_{j=1}^{N} W_j \phi_j(\Delta TL) + \sigma \tag{4}$$

where W_j and ϕ_j denote, respectively, the weight and the activation function associated with the j-th node of a single-layer neural network, and σ is a uniformly bounded approximation error whose magnitude can be made arbitrarily small by a suitable choice of N. Since f_{sw} is smooth in ΔTL, a single hidden layer is sufficient.

A **fuzzy neural network (FNN)** parametrization is adopted. The activation functions ϕ_j are linear basic functions defined on a one-dimensional uniform grid $\{\Delta TL_j\}$ of the input space $\Delta TL \in [\Delta TL_{min}, \Delta TL_{max}]$, and satisfy the partition-of-unity conditions:

$$\Sigma_{j=1}^{N} \phi_j(\Delta TL) = 1 \tag{5}$$

$$\phi_j(\Delta TL_j) = 1, \quad \forall\, j \in \{1, \ldots, N\} \tag{6}$$

$$0 \le \phi_j(\Delta TL) \le 1, \quad \forall\, j \in \{1, \ldots, N\} \tag{7}$$

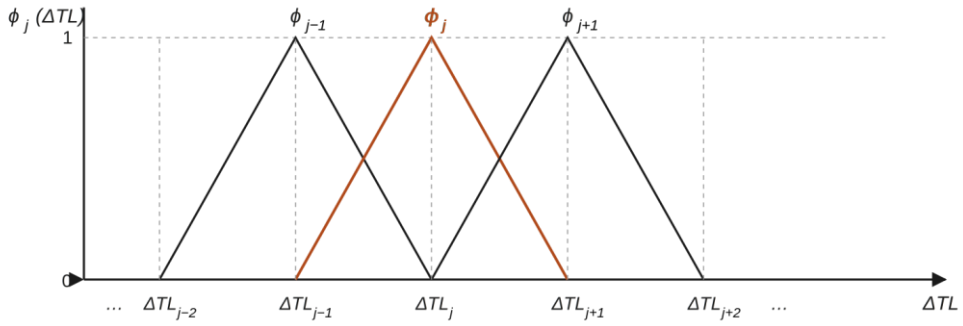


Figure 10 – Schematic representation of the 1D linear activation functions

This choice has a clear engineering interpretation: each weight W_j represents the compensation factor applied at the corresponding breakpoint, which simplifies monitoring, certification and post-flight

inspection. Collecting the activation functions and weights in the regressor vectors $\mathbf{\Phi} = [\phi_1, \ldots, \phi_N]^T$ and $\mathbf{W} = [W_1, \ldots, W_N]^T$, the disturbance and the adaptive compensation read:

$$f_{sw} = \boldsymbol{W}^{*T}\mathbf{\Phi}, \quad u_{ad} = -\widehat{\boldsymbol{W}}^{T}\mathbf{\Phi} \tag{8}$$

where $\boldsymbol{W}^*$ denotes the (unknown) ideal weight vector and $\widehat{\boldsymbol{W}}$ its on-line estimate. Substituting (8) into (3) yields the perturbed closed-loop dynamics:

$$\dot{\mathbf{x}} = \mathbf{A}\mathbf{x} + \mathbf{B}[u + \widetilde{\boldsymbol{W}}^{T}\mathbf{\Phi} + \sigma] \tag{9}$$

with the parameter estimation error defined as $\widetilde{\boldsymbol{W}} = \mathbf{W}^* - \widehat{\boldsymbol{W}}$. The approximation error σ is hereafter neglected for the purpose of the stability proof and treated as an ultimate bound in the implementation.

3.2.3 Reference model and tracking error

The reference model captures the desired closed-loop behaviour in the absence of sidewash perturbations. Let $\mathbf{x}_m \in \mathbb{R}^n$ be the reference state, governed by:

$$\dot{\mathbf{x}}_m = \mathbf{A}\mathbf{x}_m + \mathbf{B}u_m \tag{10}$$

where u_m is the reference command derived from the receptacle motion. The vector $\mathbf{x}$ gathers the two states most representative of the boom roll dynamics — the rotation angle and the angular rate about the X_{RO} axis (cf. Figure 2):

$$\mathbf{x} = [\theta,\ \dot{\theta}]^T \tag{11}$$

The tracking error and its dynamics are then:

$$\mathbf{e} = \mathbf{x} - \mathbf{x}_m \tag{12}$$

$$\dot{\mathbf{e}} = \mathbf{A}\mathbf{e} + \mathbf{B}\widetilde{\boldsymbol{W}}^{T}\mathbf{\Phi} \tag{13}$$

Note that ideally, $\mathbf{x}_m = \mathbf{x}_{receptacle}$ (see Figure 9)

3.2.4 Lyapunov-based adaptation law

The update law for $\widehat{\boldsymbol{W}}$ is derived following the direct Lyapunov approach (see REF [5]). Consider the candidate Lyapunov function, which is positive-definite:

$$V(\mathbf{e}, \widetilde{\boldsymbol{W}}) = \mathbf{e}^T\mathbf{P}\mathbf{e} + |b|\ \widetilde{\boldsymbol{W}}^{T}\mathbf{\Gamma}^{-1}\widetilde{\boldsymbol{W}} > 0 \tag{14}$$

where $\mathbf{\Gamma} = \boldsymbol{\Gamma}^{T} > 0$ is the symmetric positive-definite adaptation-rate matrix, b the control power, and $\mathbf{P} = \boldsymbol{P}^{T} > 0$ that solves the following Lyapunov equation:

$$\boldsymbol{A}^{T}\mathbf{P} + \mathbf{P}\mathbf{A} = -\mathbf{Q} \tag{15}$$

Where $\mathbf{Q}$ is any matrix that satisfies $\mathbf{Q} = \boldsymbol{Q}^{T} > 0 \in \mathbb{R}^2$ x $\mathbb{R}^2$

Differentiating equation (14) and after a few calculations (see REF [5]):

$$\dot{V}\left(\boldsymbol{e}, \widetilde{\boldsymbol{W}}\right) = -\boldsymbol{e}^{T}\boldsymbol{Q}\boldsymbol{e} + 2\widetilde{\boldsymbol{W}}^{T}|b|(\mathbf{\Phi}\ \boldsymbol{e}^{T}\overline{\boldsymbol{P}}\ sgn(b) + \ \mathbf{\Gamma}^{-1}\ \widetilde{\boldsymbol{W}}) \tag{16}$$

Where $\boldsymbol{\Theta}$ is the activation function matrix. The first term in (16) is negative semi-definite by construction. Therefore, if the second term is null, it is ensured that (16) is negative semi-definite. Therefore, from the second term it is extracted the adaptation law. Note that $\dot{\widetilde{\boldsymbol{W}}} = \dot{\widehat{\boldsymbol{W}}}$ and sgn(b) is known and negative, therefore:

$$\dot{\widehat{\boldsymbol{W}}} = \mathbf{\Gamma}\ \mathbf{\Phi}\ \boldsymbol{e}^{T}\overline{\boldsymbol{P}} \tag{17}$$

Substituting (17) back into (16):

$$\dot{V}\left(\boldsymbol{e},\widetilde{\boldsymbol{W}}\right) = -\boldsymbol{e}^{\boldsymbol{T}}\boldsymbol{Q}\boldsymbol{e} \;\leq\; 0 \qquad (18)$$

Hence $\dot{V}$ is non-increasing along solutions, which implies that $\mathbf{e}(t)$ and $\widetilde{\boldsymbol{W}}(t)$ are uniformly bounded for all $t \geq 0$. From (13) the time derivative $\dot{\mathbf{e}}$ is also bounded, so $\dot{V}$ is bounded and V is uniformly continuous. Barbalat's lemma then yields:

$$\lim_{t\to\infty} V(t) = 0 \;\Rightarrow\; \lim_{t\to\infty} \mathbf{e}(t) = 0 \qquad (19)$$

That is, the tracking error converges asymptotically to the origin while all closed-loop signals remain bounded.

3.2.5 Learning monitors

The stability proof above assumes that the only source of state deviation with respect to the reference model is the receiver sidewash. In practice, several phenomena can produce a non-zero tracking error that is *not* attributable to the disturbance f_{sw}, and that should therefore not drive the adaptation law. The following monitors freeze the update law (17) when their associated condition is met:

- **Boom in contact with the receiver structure:** if the Elastic Modes Adaptive Cancellation algorithm (see REF [4]) detects that the boom is leaning on the receiver surface, the boom dynamics are no longer those of a free-flying body and the error signal is not reliable.
- **Nozzle far from the receptacle:** adaptation is enabled only when the direct-link compensation is active, i.e. when ΔTL is below a prescribed threshold; outside this region, no compensation is required and the regressor $\boldsymbol{\Phi}$ carries no useful information.
- **High atmospheric turbulence:** the boom motion induced by turbulence is uncorrelated with the sidewash disturbance, and incorporating it into the update would degrade the parameter estimate.
- **Probe in extension to contact:** learning is restricted to the phase in which the relative distance between the nozzle and the receptacle is being reduced, ensuring monotonic exploration of the input space of the FNN.

In addition, the weight estimate $\hat{\mathbf{W}}$ is projected onto a pre-defined compact set [$\mathbf{W}_{min}$, $\mathbf{W}_{max}$] which guarantees that, even under a fault of the learning mechanism, the adaptive compensation remains within bounds ensuring always a safe behavior of the boom. This is a key factor for the certification argument.

4. CERTIFICATION

4.1 Strategy

Certifying an adaptive controller is challenging because its behavior is not strictly deterministic: the closed-loop response depends on the on-line state of the learnt parameters, which in turn depends on the receiver's operating history. The strategy used to ensure that the system behavior is always safe, smooth and predictable rests on six pillars:

- **Rationale (Lyapunov theory).** The adaptation law has been derived from a Lyapunov candidate function. The system is asymptotically stable and the tracking error converges to zero, as proved in Section 3.2.4.
- **System protections:** The automatic refueling system is also equipped with several runtime protections (misalignment monitors, stability monitors) that interrupt the contact and ensure a safe separation between the boom and the receiver at any moment if a safe contact is not ensured. The boom operator monitors the operation at all times and can take manual control by simply grabbing the flight control stick or by commanding an emergency retraction.
- **Linear analysis**: Assuming, conservatively, that the system is always learning, the adaptive

gain is bounded between known values (by virtue of the projection introduced in Section 3.2.5) and can therefore be treated as an additional proportional gain. The resulting linear closed-loop is used to compute classical stability margins (Section 4.3).

- **Non-linear analysis**: Through an extensive clearance sweeping all conditions affecting to the adaptive controller is possible to demonstrate that the system behavior is always safe, smooth and predictable. See Section 4.2
- **Impact avoidance**: Real-time and predictive minimum clearance between the boom and the receiver geometry is computed using an internal receiver mesh. In case a potential unsafe condition arises, the system triggers an immediate automatic separation at a high rate to clear the boom from the receiver.
- **Flight test**: Two dedicated flight-test campaigns explored the corners of the refueling flight envelope and the full boom spatial envelope (Section 5)

4.2 Non-linear Clearance

The main goal of the non-linear clearance is to demonstrate, by means of non-linear simulations, that no unsafe situation arises in any plausible operational scenario. The receiver motion is generated from the following sources:

1. **Flight-test data**: Recorded receiver trajectories from all in-house test campaigns and from customer reports.
2. **Adapted flight-test data.** The recorded trajectories are modified in the boom coordinates (roll, elevation, telescopic length) so as to sweep the whole boom spatial envelope and flight envelope.
3. **Adverse atmospheric conditions:** Turbulence is introduced over a certain receiver position, again sweeping the boom spatial envelope and flight envelope.
4. **Multi-Strategy Adaptive Global Optimization:** Clearance based on adaptive optimization is considered to find out the worst situations in terms of relative distance between any part of the boom with respect any part of the receiver aircraft. For more information see REF [6] & [7]
5. **Adaptive gain at extreme values:** Because the adaptive gain is bounded between known minimum and maximum values, the cases in (i)–(iii) above are run again with the adaptive gain frozen at each of the two extremes. This isolates the worst behavior of the system even under a hypothetical learning fault.

4.3 Linear Clearance

In parallel with the non-linear clearance, a linear analysis is carried out considering the system is always learning. The simplified A3R controller can be expressed as:

$$\theta^{dmd} = K_I \int (\theta^{rcv} - \theta^{noz}) + K_p(\theta^{rcv} - \theta^{noz}) + K_d \left(\dot{\theta}^{rcv} - \dot{\theta}^{noz}\right) + \delta_{OL} \quad (20)$$

Introducing the adaptive correction γ that scales the Direct Link, the loop equation becomes

$$\theta^{dmd} = K_I \int (\theta^{rcv} - \theta^{noz}) + K_p(\theta^{rcv} - \theta^{noz}) + K_d \left(\dot{\theta}^{rcv} - \dot{\theta}^{noz}\right) + \delta_{OL}[1 + \gamma \int \left(\dot{\theta}^{rcv} - \dot{\theta}^{noz}\right)] \quad (21)$$

There are monitors to disable learning when certain conditions are not fulfilled, but for linear analysis it can be considered that the system is always learning and Eq (24) would be valid. The last term can be considered as a proportional gain and Eq (24) would be:

$$\theta^{dmd} = K_I \int (\theta^{rcv} - \theta^{noz}) + (K_p + \gamma)(\theta^{rcv} - \theta^{noz}) + K_d \left(\dot{\theta}^{rcv} - \dot{\theta}^{noz}\right) + \delta_{OL} \quad (22)$$

The classical gain and phase margins computed on (24) meet the certification requirements. Both the non-linear and the linear clearance therefore confirm that the system is safe with good performance across the full operating envelope.

5. FLIGHT TESTS

Two flight-test campaigns were carried out: a development campaign in 2022 and a certification campaign in 2023. The model-based direct-link design mentioned in section 3 carried a substantial risk due to the limited fidelity of the F-15 aerodynamic dataset, which made flight testing an indispensable validation step. The two campaigns are described in the next subsections together with the most relevant results.

5.1 Development campaign

The development campaign took place in Singapore in March 2022. Six flights were used to explore the boom control behavior inside the F-15 refueling flight envelope. A total of 216 automatic contacts were performed, including the first automatic air-to-air refueling contact in history between a tanker and an F-15 fighter (Figure 11).

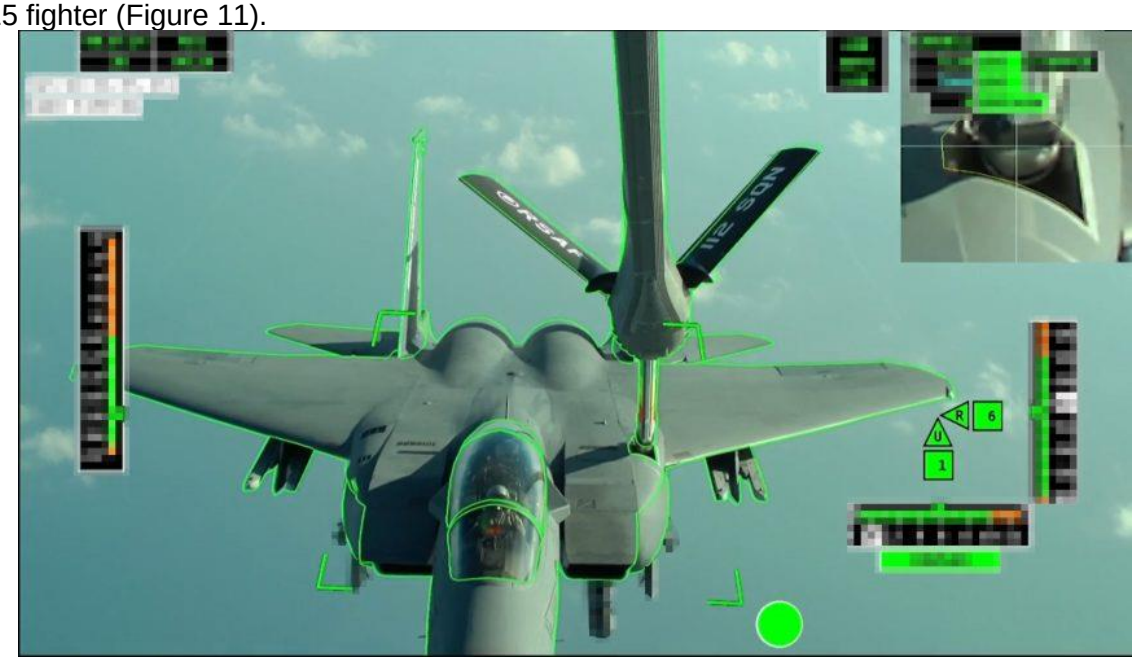


Figure 11 – First automatic air-to-air refueling contact with a F-15 fighter

The Direct-Link schedule was refined off-line between flights using the adaptive controller in a *post-flight identification* mode since it was not active in flight: the recorded data were replayed against the simulation, the sidewash dataset was tuned to match the flight trajectories, and the resulting Direct-Link values were updated for the next flight.

The conclusion of the development campaign was that the system was mature enough, pending a few minor issues, to enter a formal certification campaign in the following year.

5.2 Certification campaign

The certification campaign was carried out in Singapore in August 2023 with minor refinements over the configuration flown in 2022 and the adaptive algorithm active. It was structured in two stages:

- **Three pre-certification flights** to confirm that the issues identified during the 2022 development campaign had been fully resolved and that the system was ready to enter formal certification.
- **Three certification flights** to demonstrate compliance with the certification requirements with the Spanish certification authority, *Instituto Nacional de Técnica Aeroespacial* (INTA), on board.

The pre-certification flights confirmed the closure of all open development issues. The three certification flights then yielded 144 successful automatic contacts, covering the corners and the center of the F-15 refueling flight envelope and sweeping the full boom spatial envelope in elevation, roll and telescopic

length. Formal certification of the system was granted on the basis of these results.

Post-flight analysis showed tracking 1- σ error of less than 0.2 degrees in roll and 0.12 degrees in elevation across the receiver flight envelope, in line with tracking performance with F-16 and A330MRTT as receivers. The overall contact success rate, defined as the fraction of A3R activations that ended in a successful contact, was 64%, which is explained by the nature of the tests since the receiver is guided towards the limits of the boom spatial envelope and the contacts are very often interrupted automatically by the system if the spatial envelope limits are exceeded. Operational flights carried out after the certification campaign showed success rate over 80%, in line with those obtained for F-16 and A330-MRTT receiver types. Impact map analysis confirms that all boom nozzle contacts fell within the receptacle ramp; the few contacts near the ramp edge are within predicted tolerance and are considered acceptable. This demonstrates that this new direct-link adaptive strategy proposed for the F-15 as receiver presents enough accuracy to operate A3R.

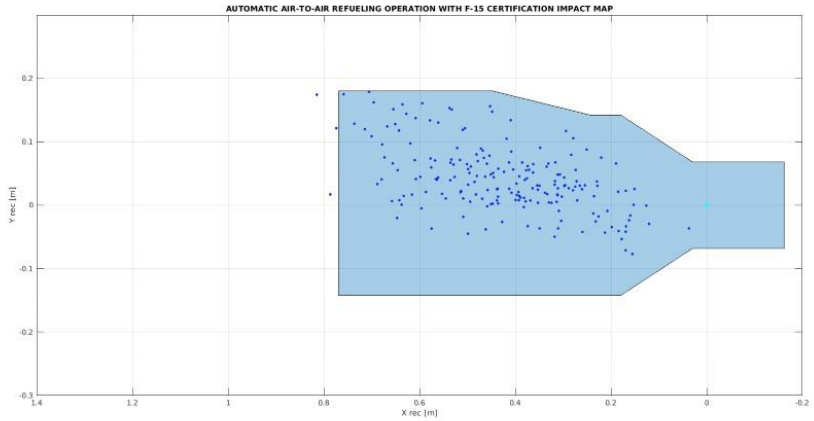


Figure 12 – Certification flights. Impact map analysis

6. CONCLUSIONS

Air-to-air refueling is one of the most demanding manual tasks in flight, and the difficulty is amplified when the receptacle is located out of the plane of symmetry of the receiver. In a fly-by-wire boom system such as the one fitted to the A330-MRTT, the operator does not feel the aerodynamic force on the boom coming from the receiver flow field, the only available cue is the visual lateral drift of the boom away from the receptacle. The operator therefore has to learn to apply a corrective lateral input whenever the probe is being extended towards the receptacle, in order to compensate the strong local sidewash effect.

This paper has shown how the automatic system reproduces that learnt behavior. The expected ruddervator compensation is computed from CFD-based receiver upwash aerodynamic data and injected as an open-loop Direct Link, while a Model Reference Adaptive Controller (MRAC) running on top of the Direct Link absorbs the residual mismatch between the model and reality. The adaptation law has been derived from a Lyapunov candidate and proves asymptotic convergence of the tracking error, while a fuzzy neural-network parameterization gives every weight a transparent physical interpretation. The certification strategy combining Lyapunov rationale, linear analysis, non-linear clearance and dedicated flight-test campaigns has been described, and the system with the new direct-link plus adaptive controller was certified after 360 automatic contacts performed across two campaigns.

7. FUTURE

The current state of automatic air-to-air refueling (REF [8]) is automatic but not fully autonomous: the system tracks the receptacle and extends the boom whenever the conditions are met, and it reacts automatically to any detected failure, but it still requires the boom operator to monitor the operation and to take manual control in a small number of off-nominal situations. Ongoing developments aim at full autonomy in three complementary directions:

- **Autonomous Assets Air-to-Air Refueling (A4R)** — a demonstrator in which an A310-MRTT tanker manages and guides a formation of unmanned receivers; see REF [9]. The end goal is a leader-on-air capable of controlling a heterogeneous formation of manned and unmanned receivers through the standard refueling positions.
- **Decision-making algorithms** — a higher-level layer that selects boom actions on the basis of a Maslow-like hierarchy of needs, emulating the levels of prioritization of an experienced operator (REF [10] & [11]). The aim is a boom that responds quickly to off-nominal situations with the smoothness and predictability of an experienced boom operator.
- **Active hose for hose-and-drogue refueling** — a controller that stabilizes the hose, reducing both the bounce induced by atmospheric turbulence and the lateral drift produced by the receiver flow field as the receiver probe approaches the drogue.

References

[1] Carlos A., Daniel P., Melchor B. *Methods and systems for reducing the phenomenon of structural coupling in the control system of an in-flight refuelling boom*, Patent US-8332081-B2. 2008.

[2] Francisco Jose A., Rodney R., Alberto S. *Method for Adaptively cancelling in real time elastic modes in discrete-time signals*. Patent US-10514318-B2. 2019.

[3] Rodney R., Alberto S. *Elastic Modes Adaptive Cancellation Algorithm for the A330 MRTT Flying Boom Control Laws.* ICAS, Seoul, Korea. 2016

[4] Rodney R., Pablo P., Francisco Jose A. Patent US-12091186: *Method for detecting a physical status of a flexible structure* 2024.

[5] Nhan T. Nguyen. *Model-Reference Adaptive Control*: A Primer. 2018

[6] Rodney R., Antonio A., Gregorio V., Nidal O., Francisco A.: *Flight Control Laws Clearance based on Coverage-granted Multi-Strategy Adaptive Global Optimization.* EUCASS Conference. 2022

[7] Rodney R., Miguel S., Francisco A.: *Flight Control Laws Carefree Handling Clearance of a Highly Manoeuvrable Aircraft using Multi-Strategy Adaptive Global Optimization.* EUCASS Conference. 2017

[8] Víctor M., Ignacio R., Alberto S., Francisco A.: *Advanced Refueling Boom System (ARBS) Automatic Air to Air refueling Flight Control Laws.* AIAA Conference. 2024

[9] Airbus newsroom: *Aerial refueling without human intervention.* https://www.airbus.com/en/newsroom/stories/2023-11-aerial-refuelling-without-human-intervention

[10] Ignacio R., Francisco A.: *Decision-making modelling: application to autonomous Air to Air refueling.* EUCASS Conference. 2019

[11] Ignacio R., Víctor M., Francisco A.: *Decision-making modelling application to Air to Air refueling approach phase.* AIAA Conference. 2024

Copyright Statement
The authors confirm that they, and/or their company or organization, hold copyright on all of the original material included in this paper. The authors also confirm that they have obtained permission, from the copyright holder of any third party material included in this paper, to publish it as part of their paper. The authors confirm that they give permission, or have obtained permission from the copyright holder of this paper, for the publication and distribution of this paper as part of the ICAS proceedings or as individual off-prints from the proceedings.