Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Forced transition as a design and robustness variable for the lifting surfaces of solar stratospheric platforms

Melih Yıldız

melihy@erciyes.edu.tr

Erciyes University

Research Article

Keywords: laminar separation bubble, boundary-layer trip, low Reynolds number, high-altitude platform, transition prediction, aircraft sizing

Posted Date: November 5th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10907059/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Abstract

Forcing boundary-layer transition is the standard remedy for the laminar separation bubble that dominates the drag of low-Reynolds lifting surfaces, and on a solar stratospheric platform it is normally discussed for the wing. Here the forced-transition location is treated as an explicit design variable inside a coupled environment–sizing–mission closure, the trip is charged for its own excrescence drag, and the result is resolved by component. Section polars for five low-Reynolds airfoils are generated over chord Reynolds numbers $3 \times 10^4 \le Re_c \le 10^6$ with the transition location and the amplification factor n_{crit} as free axes; the trip height is sized from the roughness Reynolds number criterion. Three results follow. First, charging the trip for its own drag moves the profitability threshold from $Re_c \approx 3.5 \times 10^5$ to $\approx 2.0 \times 10^5$. Second, the platform straddles that threshold: the wing, at $Re_c = 1.97 \times 10^5$, loses from a trip at every height considered, while the empennage, at 1.08×10^5, gains throughout — tripping both surfaces costs 1.8% in gross mass, the empennage alone saves 4.6%, the wing alone costs 6.6%. Third, the benefit grows as the disturbance environment quiets, from 1.5% at $n_{\mathrm{crit}} = 7$ to 8.7% at $n_{\mathrm{crit}} = 13$, so transition control is worth most precisely where transition prediction is least verifiable; against wind-tunnel data, tripping reduces the sensitivity of predicted drag to the amplification factor eighteen-fold. The design rule is therefore component-wise, not chordwise: trip the surfaces below the bubble threshold and leave those above it clean.

1. Introduction

A solar-powered stratospheric platform flies where the air is thin enough that its wing chord Reynolds number falls close to 2×10^5. In that band the laminar boundary layer on the upper surface separates under the adverse pressure gradient, transitions in the detached shear layer, and reattaches as a turbulent layer. The resulting laminar separation bubble carries a pressure-drag penalty that can exceed the skin friction it replaces, and it makes the section polar sensitive to conditions that are not part of the geometry at all.

The standard remedy is to force transition ahead of the separation point with a boundary-layer trip. The remedy is old, and its section-level behaviour has been documented carefully: Gopalarathnam et al. [1] designed three airfoils with systematic variations of the transition ramp and tested them at $Re_c = 10^5$ and 3×10^5 with the trip placed at $x/c = 0.1,\ 0.2,\ 0.3$ and 0.4 on the upper surface. Their measurements already contain most of the section physics used here: the optimum location was the aft-most tested, forward trips raised the drag above the clean airfoil at the higher Reynolds number, and aft trips lost effectiveness at high lift coefficient “particularly when the laminar separation location is upstream of the trip”. Their overall conclusion was that, for the Reynolds numbers and shapes considered, airfoils designed with trips held no clear advantage over airfoils designed clean.

Two things are missing from that picture when the question is asked about a vehicle rather than a section.

The first is the device itself. A trip is a physical excrescence standing in the boundary layer, and its form drag is subtracted from whatever the earlier transition wins. A wind-tunnel measurement of a tripped airfoil contains that drag implicitly and correctly, but it cannot be separated from the result, and so it cannot be swept as a design variable. In a predictive model the two must be assembled separately, and at the Reynolds numbers of interest they are of the same order: the section gain at the design point of the platform considered here is $\Delta c_d \approx 1 \times 10^{-4}$, while a trip sized to be effective contributes between 3×10^{-4} and 1.6×10^{-3}. The sign of the net result therefore cannot be assumed.

The second is that an aircraft is not a section. The empennage of a stratospheric platform operates at a chord roughly half that of the wing and therefore at half its Reynolds number, which can place the two surfaces on opposite sides of the threshold where tripping stops paying. The literature on low-Reynolds turbulators is written almost entirely about the main wing; the tail is rarely mentioned. On the platform studied here the tail carries 18% of the reference area but about 22% of the parasite drag, and, as will be shown, essentially all of the available trip benefit.

Active flow control offers an alternative route to the same objective. Micro- blowing through a microporous wall reduces turbulent skin friction by up to 51% in the configuration studied by Najafi Khaboshan et al. [2], and plasma actuators and synthetic jets report gains of comparable magnitude. These methods carry power, ducting and certification costs that a solar platform operating on a closed diurnal energy budget cannot absorb, so passive tripping remains the only realistic option in this application. The question is not whether to control transition actively or passively, but where a passive device belongs and what it is worth once the vehicle is resized around it.

This paper answers that question by carrying the forced-transition location as an explicit axis of the section table, charging the trip for its own drag, and propagating the result through a coupled environment–sizing–mission closure to gross mass, cruise power and endurance. Section 2 describes the section model and its two new axes, including the physical meaning of the amplification factor and a validation against wind-tunnel measurements in both clean and tripped configurations. Section 3 develops the trip drag budget. Sections 4–7 give the section-level results, their propagation to the vehicle, the sensitivity to the disturbance environment, and the endurance budget. Section 8 states the limitations and Section 9 the conclusions.

The sizing closure used as a tool here, and the section-table method it rests on, are those of the companion papers [3, 4]; neither is repeated. The contribution of the present paper is the transition axis, the trip drag budget, the component-wise resolution, and the robustness result.

2. Section model with transition as a free axis

2.1. The polar table

Section drag is supplied by a tabulated polar, generated with NeuralFoil [5], a trained surrogate of XFOIL [6] evaluated in its largest configuration. The method is that of the companion paper [4]; what is new

here is that two quantities previously held fixed are carried as axes.

The table spans five low-Reynolds sections (FX 63–137, LA2573A, SG6043, E387, DAE-21) over $3 \times 10^4 \leq Re_c \leq 10^6$ at 25 logarithmically spaced stations, at section lift coefficients $c_l = 0.2,\ 0.4,\ 0.6,\ 0.8,\ 1.0,\ 1.2$, for eleven transition locations and four amplification factors, in two trip configurations. The full cross product is 63 000 rows, of which 2 275 are unreachable operating points and are carried as missing values so that the sizing loop can interpret them as “this section cannot fly this point”.

The lift-coefficient nodes deserve a remark, because the range matters more than it appears. A tail surface operates near $c_l = 0.2$, and at low Reynolds number on a cambered section the drag there is *higher* than at $c_l = 0.8$, not lower: for FX 63–137 at $Re_c = 1.07 \times 10^5$ the values are $c_d = 0.0291$ and 0.0261 respectively. A table whose lowest node is $c_l = 0.8$ therefore under-predicts tail drag systematically, and the error propagates to gross mass. The nodes used here extend to $c_l = 0.2$ for this reason.

2.2. Definition of the trip

The transition location x_{tr}/c is the station at which transition is forced. The value $x_{tr} = 1$ denotes free transition, that is, the e^N prediction with no device, and reproduces the baseline table of the companion work to within 0.01%. Two trip configurations are carried: the device on the upper surface only, and on both surfaces. Unless stated otherwise the results below use the upper-surface configuration, which is what a turbulator on a lifting surface normally means.

The wing and the empennage carry independent transition locations. This is not a formality: it is the degree of freedom in which the main result of this paper is expressed.

2.3. The amplification factor as a receptivity parameter

The e^N criterion collapses the receptivity of a boundary layer into a single amplification factor N. It matters for what follows that N is not a property of the atmosphere: it absorbs the free-stream turbulence, the acoustic field, and the roughness and waviness of the surface into one number, and it is normally calibrated for a class of configuration rather than computed.

The correlation of Mack [7] relates N to the free-stream turbulence intensity of a wind tunnel,

$$N = -8.43 - 2.4 \ln Tu, \quad 0.0007 \leq Tu \leq 0.0298, \quad (1)$$

and places the value $N = 9$ used throughout the companion work at $Tu \approx 0.07\%$. That is already a low-turbulence tunnel rather than a noisy one; the qualifier “average” in the usual tabulation refers to tunnels, not to flight. The correlation cannot be carried much further down, however: for $N \gtrsim 11$ the implied turbulence intensity falls below the calibrated range, where receptivity is governed by surface roughness, waviness and acoustics rather than by free-stream turbulence [8]. The flight end of the ladder

is therefore anchored not by the correlation but by the configuration-class values tabulated for XFOIL [6], which assign $N = 10–12$ to a clean wind tunnel, $11–13$ to a motorglider and $12–14$ to a sailplane.

A stratospheric platform wing belongs to the sailplane class in every respect that sets N: a single-piece composite surface, finished and polished, with a large chord and a low wing loading; and the strongly stratified air at 20 is quieter than any tunnel. We therefore take $N = 11–13$ as flight-representative, with $N = 11$ the conservative choice, and retain $N = 9$ only as the tunnel-calibrated reference against which the section model was validated. The lower end of the sweep, $N = 7$, is not hypothetical either: it represents a degraded surface — particulate accumulation, ultraviolet ageing, thermal cycling, repair marks — which a platform designed for months of continuous flight must be assumed to reach.

This is the first reason the amplification factor deserves to be carried as a free axis rather than fixed at a default. The second is that no single value is correct for the whole mission: N drifts downward as the surface degrades, so the vehicle traverses a band of N in service. A design whose closure depends strongly on N is therefore not merely uncertain — it is unstable over its own operating life.

2.4. Validation against measurement in both configurations

The section model is compared with the wind-tunnel measurements of Selig and McGranahan [9], who tested six low-Reynolds airfoils in the University of Illinois at Urbana–Champaign (UIUC) low-speed subsonic tunnel, in a clean configuration and with a zigzag trip of 0.11% chord thickness placed at $x/c = 0.02$ on the upper surface and $x/c = 0.05$ on the lower; the measurements are referred to as the UIUC data in the figures below. Two of those sections are carried in the present table: FX 63–137, the section used for the platform, and the E387 benchmark. The tabulated polars were extracted from the report and the model was run at the measured trip geometry rather than at a tabulated node, so the comparison is direct.

The amplification factor is carried at two values, $N = 7$ and $N = 9$. This is not a free choice: the flow-quality survey reported in the same document places the tunnel at a turbulence intensity of $0.10–0.17\%$ with the test rig installed, which through Eq. (1) corresponds to $N = 7.2–9.0$. The band between the two therefore represents the facility's own disturbance uncertainty.

Two results follow, and they are shown in Fig. 1 and Table 1.

On absolute level the model carries a mean absolute error of 13–20% over the design lift range $0.4 \le c_l \le 1.2$, and the deviation is strongly Reynolds-dependent: it reaches 47.6% at $Re_c = 10^5$ for the clean FX 63–137 and falls to 6.4% at 5×10^5. This is the familiar optimism of an integral bubble model at the low-Reynolds end, noted for the same class of method in Ref. [1] and consistent with the section-level validation reported in the companion work [3]. It is the reason every quantitative statement in this paper is expressed as a ratio at fixed Reynolds number rather than as an absolute drag level.

Table 1

Validation statistics over $0.4 \le c_l \le 1.2$ at $n_{crit} = 9$, and the shift in the predicted drag when n_{crit} is moved from 7 to 9 at fixed Reynolds number.

Section	Configuration	Mean $\lvert$error$\rvert$	Signed error range over Re
E387 (E)	clean	16.8%	$-15.8\% \ldots +7.6\%$
E387 (E)	trip F	12.7%	$-5.8\% \ldots +17.5\%$
FX 63–137 (C)	clean	17.4%	$-47.6\% \ldots -6.4\%$
FX 63–137 (C)	trip F	19.9%	$-35.6\% \ldots +19.8\%$
n_{crit} shift, clean		mean 3.9%, median 2.2%, **max 12.8%**	
n_{crit} shift, trip F		mean 0.1%, median 0.0%, **max 0.7%**	

The second result matters more for the argument developed here. Moving N from 7 to 9 shifts the predicted drag by up to 12.8% in the clean configuration but by at most 0.7% with the trip in place — a factor of eighteen. The claim that forcing transition removes the design's dependence on the transition model is therefore not an artefact of the model reasoning about itself; it is visible in the comparison with measurement.

The trip in these tests sits near the leading edge, so this comparison validates the forward-trip branch of the model directly. The aft trip location that emerges in Section 5 remains an extrapolation of a model whose response to tripping has been verified, and is identified as such in Section 8.

Use of large language models. A large language model (Claude, Anthropic) was used to automate the numerical analysis chain — scripting the polar-table generation, the trip-drag bookkeeping and the figure pipeline — under the author's direction. All methods, equations and results were specified and verified by the author, who takes full responsibility for the content of this work.

3. The drag of the trip itself

A section polar computed with forced transition answers what the flow does if transition occurs at x_{tr}; it says nothing about the device that makes it occur. The chain used here to supply that missing term is assembled from quantities the section solution already provides, so that no part of it is a free parameter.

The boundary-layer state at the trip station — momentum thickness θ, shape factor H and edge velocity u_e/V_∞ — is taken from the section solution, which reports them at 32 chordwise stations. Thwaites' correlation is inverted to obtain the local pressure-gradient parameter λ from the shape factor, and the Pohlhausen quartic profile with the corresponding Λ then gives the velocity profile and the wall shear [10]. At zero pressure gradient the pair reproduces the Blasius shape factor ($H = 2.554$ against 2.59) and the Thwaites shear function (0.2349 against 0.2205) to within 7%, which is the accuracy this estimate requires.

The trip height follows from the roughness Reynolds number criterion $Re_k = u_k k/\nu$ of Braslow and Knox [11]. Their critical value $Re_k = 600$ applies to three-dimensional distributed roughness in an attached laminar layer and is an upper bound in the present application, where the device does not have to destabilise a stable layer but only to perturb a shear layer already close to separation. A band $Re_k = 200$–600 is therefore carried, together with the height ratio $k/c \approx 0.002$ typical of low-Reynolds practice; for reference, the zigzag tape used in Ref. [9] has $k/c = 0.0011$ at a leading-edge station, where the boundary layer is far thinner than at the aft stations considered here. The excrescence drag, referred to chord and free-stream dynamic pressure, is taken as

$$\Delta c_{d,\mathrm{trip}} = C_{D,k}\,\phi\left(\frac{u_k}{V_\infty}\right)^2 \frac{k}{c}, \tag{2}$$

with $C_{D,k}$ the protuberance drag coefficient on frontal area and ϕ the spanwise coverage of a zigzag tape, whose serrated edge does not present a solid frontal strip. The product $C_{D,k}\phi$ is swept over 0.4–1.2 with 0.7 as the nominal value.

Table 2

Trip height and excrescence drag at the design point, for a trip at $x_{tr}/c = 0.60$. The raw section gain is the change in c_d from the polar alone; the net is that gain plus the device drag of Eq. (2) at $C_{D,k}\phi = 0.7$.

Surface	Criterion	k [mm]	k/δ	u_k/V_∞	$\Delta c_{d,\mathrm{trip}}$	Net
Wing	$Re_k = 200$	2.63	0.21	0.410	0.00029	$+0.00019$
$Re_c = 1.97 \times 10^5$	$Re_k = 400$	3.62	0.29	0.595	0.00085	$+0.00075$
$c_l = 1.36$	$Re_k = 600$	4.40	0.35	0.736	0.00157	$+0.00147$
raw gain -0.00010	$k/c = 0.002$	2.12	0.17	0.318	0.00014	$+0.00004$
Tail	$Re_k = 200$	2.15	0.29	0.501	0.00065	-0.00524
$Re_c = 1.08 \times 10^5$	$Re_k = 400$	3.01	0.40	0.717	0.00185	-0.00403
$c_l = 0.20$	$Re_k = 600$	3.71	0.50	0.872	0.00338	-0.00251
raw gain -0.00589	$k/c = 0.002$	1.17	0.16	0.244	0.00008	-0.00580

The outcome is asymmetric between the two lifting surfaces, and Table 2 gives it in full. On the wing the raw section gain is -1.0×10^{-4} while the device contributes between $+2.9 \times 10^{-4}$ and $+1.6 \times 10^{-3}$: the net is a loss for every trip height considered. The shape factor at the trip station is 4.56, above the laminar separation value, which explains the absence of a gain — at the design lift coefficient the flow has already separated upstream of $x/c = 0.6$, so the strip is placed inside the bubble rather than ahead of it. This is the mechanism identified experimentally in Ref. [1], where aft trips were observed to lose effectiveness once the laminar separation point moved upstream of the device. On the tail the raw gain is -5.9×10^{-3} against a device drag of $+0.7 \times 10^{-3}$ to $+3.4 \times 10^{-3}$: the net is a gain throughout, with the device consuming 31% of the section gain at the nominal setting.

Charging for the device also moves the profitability threshold. On the raw section gain alone, a trip at $x/c = 0.60$ pays up to $Re_c \approx 3.5 \times 10^5$ at $c_l = 1.0$; once the device is charged for, the threshold

falls to $Re_c \approx 2.0 \times 10^5$. Figure 2 shows both curves together with the required trip height across the Reynolds range.

4. Section-level results

The threshold quoted above is specific to $c_l = 1.0$, and it is worth being explicit that it depends on the lift coefficient. A single threshold drawn across a vehicle would be misleading, because the two surfaces operate at very different lift coefficients as well as at different Reynolds numbers. Evaluated at each surface's own operating condition and with the device charged for, the thresholds are $Re_c = 0.99 \times 10^5$ for the wing at $c_l = 1.36$ and 2.03×10^5 for the tail at $c_l = 0.20$.

Figure 3 places the platform against those numbers. The wing operates at 1.99 times its own threshold and the tail at 0.53 times its own: the same aircraft sits on both sides of the same physical boundary. The wing would have to fly at roughly 29 before a trip began to pay, well above the 24 ceiling of the design. Figure 3(b) shows the consequence directly: swept over the transition location at each surface's own Reynolds number and lift coefficient, the wing curve never falls below unity while the tail curve reaches 0.795 near $x_{tr}/c = 0.6$.

The chordwise behaviour of the section ratio itself is shown in Fig. 4. The optimum location moves aft with Reynolds number, from $x/c \approx 0.30$ at $Re_c = 3 \times 10^4$ to 0.55–0.60 over 10^5–3×10^5, consistent with the aft optimum found experimentally over the narrower range of Ref. [1]. A conventional forward strip at $x/c \approx 0.05$ crosses from gain to loss at $Re_c \approx 1.1 \times 10^5$ on the raw section coefficient, and costs 36% in section drag at the design condition of the wing.

Figure 5 shows the same section swept over the amplification factor at free transition. Between $Re_c \approx 7 \times 10^4$ and 10^5 a change from $n_{\mathrm{crit}} = 9$ to 13 multiplies the predicted drag by a factor of 2.7 to 4.5, which is the quantitative form of the statement that a free-transition prediction in this band is only as good as the receptivity estimate behind it.

5. Propagation to the vehicle

The section results above are propagated through the coupled environment–sizing–mission closure of the companion work [3], in which gross mass is the smallest root of a mass balance whose battery term is set by the integral of the diurnal power deficit, and the planform is selected by sweeping span, aspect ratio and lift coefficient. The method is summarised in Fig. 6. The device drag of Eq. (2) enters the parasite drag build-up alongside the section terms,

$$C_{D0} = (c_{d,w} + \Delta\, c_{d,\mathrm{trip},w}) + \frac{S_t}{S}(c_{d,t} + \Delta\, c_{d,\mathrm{trip},t}) + c_{d,\mathrm{misc}}, \quad (3)$$

so that each surface is charged only for the device it actually carries.

The reference mission is a 5, 60 W payload held at 20 at 39° N on the summer solstice for seven days in solar mode. With both surfaces free the closure gives 173.8 gross mass at $b = 34 m$ and $AR = 32$, with $C_L = 1.357$, $Re_c = 1.97 \times 10^5$ at the wing and 1.08×10^5 at the tail, and 2033 W of cruise power.

Table 3
Configuration comparison at the reference mission, with the trip charged for its own drag at $C_{D,k}\phi = 0.7$. Starred rows repeat the same configuration with the device drag omitted, and are shown only to isolate its effect.

A	Configuration	m [kg]	P [W]	$\Delta m / m_0$
	both surfaces free	**173.8**	**2033**	**—**
B	both tripped at $x_{tr}/c = 0.60$	176.9	2007	$+1.8\%$
B*	as B, device drag omitted	163.8	—	-5.8%
C	**tail only tripped**	**165.8**	1903	-4.6%
C*	as C, device drag omitted	162.7	—	-6.4%
D	wing only tripped	185.3	2171	$+6.6\%$

Table 3 and Fig. 7 give the result. Tripping both surfaces, which is what a section-level recommendation of an aft turbulator would produce if applied uniformly, costs 1.8% in gross mass once the device is charged for; the same configuration appears to save 5.8% if it is not. Tripping the wing alone is the worst option at $+6.6\%$, because it pays for a device that returns nothing. Tripping the empennage alone saves 4.6%, and it does so across the whole of the assumed range of the device drag coefficient: -5.4%, -4.6% and -3.2% at $C_{D,k}\phi = 0.4$, 0.7 and 1.2 respectively.

6. Robustness to the transition model

The gross masses quoted so far are computed at $n_{\mathrm{crit}} = 9$, the tunnel-calibrated value. Section 2.3 argued that a flight vehicle sees $N = 11$–13, and that the value drifts over the life of the surface. Figure 8 and Table 4 give the closure over the whole band.

Table 4
Gross mass over the amplification-factor band, with the empennage trip charged for its own drag.

n_{crit}	A: both free [kg]	C: tail tripped [kg]	Gain
7	161.4	159.0	$-1.5\,\%$
9	173.8	165.8	$-4.6\,\%$
11	193.3	179.1	$-7.3\,\%$
13	222.2	203.0	$-8.7\,\%$
spread	37.7%	27.7%	

Two readings follow. The first is that the benefit of the trip grows as the air becomes quieter, from 1.5% at $n_{\mathrm{crit}} = 7$ to 8.7% at $n_{\mathrm{crit}} = 13$: transition control is worth most exactly where the transition prediction is least verifiable. The second is that the spread of the closure over the band narrows from 37.7% to 27.7%. That narrowing is real but modest at the vehicle level, and it is smaller than the section-level result of Section 2.4 would suggest in isolation, because the wing is left clean in configuration C and therefore continues to carry its own free-transition sensitivity.

This is the honest form of the robustness argument. At the section level the sensitivity to the amplification factor collapses by a factor of eighteen when a surface is tripped, and that collapse is confirmed against measurement. At the vehicle level only the tripped surface enjoys it, so the design as a whole becomes markedly, but not completely, less dependent on a quantity that is neither measured in flight nor constant over a mission.

7. Endurance budget

The mass differences above are the result of resizing. It is also useful to ask what the same choice is worth on a fixed airframe, because that is the form in which an operator meets it. Holding the baseline planform, mass, battery and solar array fixed and changing only the section aerodynamics and the device drag, the level-flight endurance on battery alone at 20 is 11.22 h in configuration A. Tripping the empennage adds 0.18 h; tripping both surfaces removes 0.08 h; tripping the wing alone removes 0.25 h.

The corresponding margin over the 24-hour battery deficit is more revealing, because it changes sign with season. On the design day the margin is $+0.53h$ in configuration A and $+0.74h$ in configuration C. At the equinox the same airframe is short of closure in every configuration, at $-2.08h$ and $-1.88h$

respectively: the empennage trip recovers about a fifth of an hour of a two-hour shortfall, which shifts the seasonal boundary slightly but does not move it. Figure 9 gives both quantities.

8. Limitations

Surrogate model accuracy. The section data come from a trained surrogate of an integral boundary-layer method whose bubble model is optimistic at low Reynolds number. Section 2.4 quantifies this: 13–20% mean absolute error over the design lift range, rising to 47.6% at $Re_c = 10^5$ for the clean FX 63–137. All quantitative statements in this paper are consequently expressed as ratios at fixed Reynolds number, where the systematic component of that bias largely cancels. Absolute masses should be read as the output of a consistent model, not as predictions of a built vehicle.

Scope of the tripped validation. The measurements available in both configurations for the sections used here place the trip near the leading edge, at $x/c = 0.02$ and 0.05. They therefore validate the forward-trip branch of the model directly, and the model's response to tripping in general, but not the aft location that the results recommend. The measurements of Ref. [1] extend to $x/c = 0.4$ on purpose-designed sections and support the aft trend qualitatively; a direct validation at $x/c \approx 0.6$ on one of the present sections does not exist and would be the most useful single addition to this work.

The excrescence drag coefficient. The product $C_{D,k}\phi$ in Eq. (2) is carried as a band, 0.4– 1.2, on engineering grounds rather than from a measurement of the specific device. The configuration conclusion is unchanged across that band, and the sign of the result never changes within it, but a measured value would replace the weakest link in the chain. The trip geometry itself is better anchored: Ref. [9] documents a zigzag tape of $k/c = 0.0011$ at a leading-edge station.

Boundary-layer closure at the trip station. The Thwaites–Pohlhausen pair used to obtain the wall shear is at the edge of its validity where $H > 4$, which is the condition found at the wing trip station at the design lift coefficient. The wing conclusion is insensitive to this, because the raw section gain there is already of order zero and the sign of the net does not depend on the precision of the height estimate; the tail station, at $H = 4.33$, is in the same regime and the same caveat applies with the same weak consequence.

Single transition location per surface. Each surface carries one chordwise transition location, constant across the span. Spanwise variation — a clean root with a tripped tip, or a stall strip — would couple to the stall pattern and to the induced drag, and is outside the present model.

Two-dimensional section data. The closure is section-based, as in the companion work, and inherits that assumption. The maximum lift coefficient receives only a relative correction from the table, with the absolute level taken from an empirical model anchored on wind-tunnel data and the correction clipped to $(0.70,\ 1.35)$.

Surface degradation. The amplification factor is used here as a proxy for the disturbance environment including surface quality, but the model contains no explicit account of contamination, erosion or repair over a months-long mission. The direction of the effect is clear — degradation lowers the effective N, and Table 4 shows that this moves the design toward the lighter end of the band — but its rate is not modelled.

9. Conclusions

Forced transition has been carried as an explicit design variable in a coupled environment–sizing–mission closure for a solar stratospheric platform, with the trip charged for its own excrescence drag and the result resolved by component. Three conclusions follow.

First, the profitability of a trip cannot be established from the section polar alone. Charging the device for its own drag moves the threshold at $c_l = 1.0$ from $Re_c \approx 3.5 \times 10^5$ to $\approx 2.0 \times 10^5$, and the threshold is a function of lift coefficient as well as Reynolds number.

Second, the platform straddles that threshold, and the design rule that follows is component-wise rather than chordwise. The wing operates at twice its own threshold, where a trip is a net loss at every height considered, because at the design lift coefficient the laminar layer has already separated ahead of any useful trip station. The empennage operates at half of its own threshold, where a trip is a net gain throughout. Tripping both surfaces costs 1.8% of gross mass; tripping the empennage alone saves 4.6%; tripping the wing alone costs 6.6%. The rule is therefore: trip the surfaces that lie below the bubble threshold, and leave those above it clean.

Third, the value of transition control is best understood as the removal of a dependence rather than as a reduction of a mean. The benefit of tripping the empennage grows from 1.5% to 8.7% of gross mass as the amplification factor moves from 7 to 13, that is, as the environment becomes quieter and the free-transition prediction less verifiable. At the section level the sensitivity of the predicted drag to the amplification factor falls by a factor of eighteen when a surface is tripped, and this is confirmed against wind-tunnel measurement rather than asserted from the model. At the vehicle level the corresponding narrowing is from 37.7% to 27.7%, limited because the wing is deliberately left clean and retains its own sensitivity.

The most useful extension of this work would be a direct measurement at an aft trip location on one of the sections used here, which would replace the one extrapolation the present results still rest on.

Declarations

Competing interests.

The author has no relevant financial or non-financial interests to disclose.

Funding

The author declares that no funds, grants, or other support were received during the preparation of this manuscript.

Author Contribution

M.Y. did all the research, visualisation and validationn and wrote the main manuscript text. All authors reviewed the manuscript.

Data Availability

The section polar tables, the trip drag budget, and the scripts that generate every figure and quot-ed number are available from the author on reasonable request.

References

1. Gopalarathnam, A., Broughton, B.A., McGranahan, B.D., Selig, M.S.: Design of low Reynolds number airfoils with trips. J. Aircr. **40**(4), 768–775 (2003). 10.2514/2.3157
2. Najafi Khaboshan, H., Yousefi, E., Svorcan, J.: Analysis of the turbulent boundary layer and skin-friction drag reduction of a flat plate by using the micro-blowing technique. J. Appl. Mech. Tech. Phys. **63**(3), 425–436 (2022). 10.1134/S0021894422030075
3. Yıldız, M.: Design regions and operational limits of solar-powered stratospheric platforms from a coupled environment–sizing–mission model. under Rev., (2026)
4. Yıldız, M.: The minimum admissible chord: a Reynolds-endogenous wing sizing model and the design space of solar stratospheric platforms. under Rev., (2026)
5. Sharpe, P.D., Hansman, R.J.: NeuralFoil: an airfoil aerodynamics analysis tool using physics-informed machine learning, arXiv:2503.16323 ((2025). 10.48550/arXiv.2503.16323
6. Drela, M.: XFOIL: an analysis and design system for low Reynolds number airfoils, in: Low Reynolds Number Aerodynamics, Vol. 54 of Lecture Notes in Engineering, Springer, pp. 1–12. (1989). 10.1007/978-3-642-84010-4_1
7. Mack, L.M.: Transition and laminar instability, Tech. Rep. JPL Publication 77 – 15, Jet Propulsion Laboratory (1977)
8. van Ingen, J.L.: The e^N method for transition prediction: historical review of work at TU Delft, AIAA Paper 2008–3830, 38th AIAA Fluid Dynamics Conference (2008). 10.2514/6.2008-3830

9. Selig, M.S., McGranahan, B.D.: Wind tunnel aerodynamic tests of six airfoils for use on small wind turbines, (2004). Tech. Rep. NREL/SR-500-34515, National Renewable Energy Laboratory
10. Schlichting, H., Gersten, K.: Boundary-Layer Theory, 9th Edition, Springer, (2017). 10.1007/978-3-662-52919-5
11. Braslow, A.L., Knox, E.C.: Simplified method for determination of critical height of distributed roughness particles for boundary-layer transition at Mach numbers from 0 to 5, Tech. Rep. NACA TN 4363, National Advisory Committee for Aeronautics (1958)

Figures

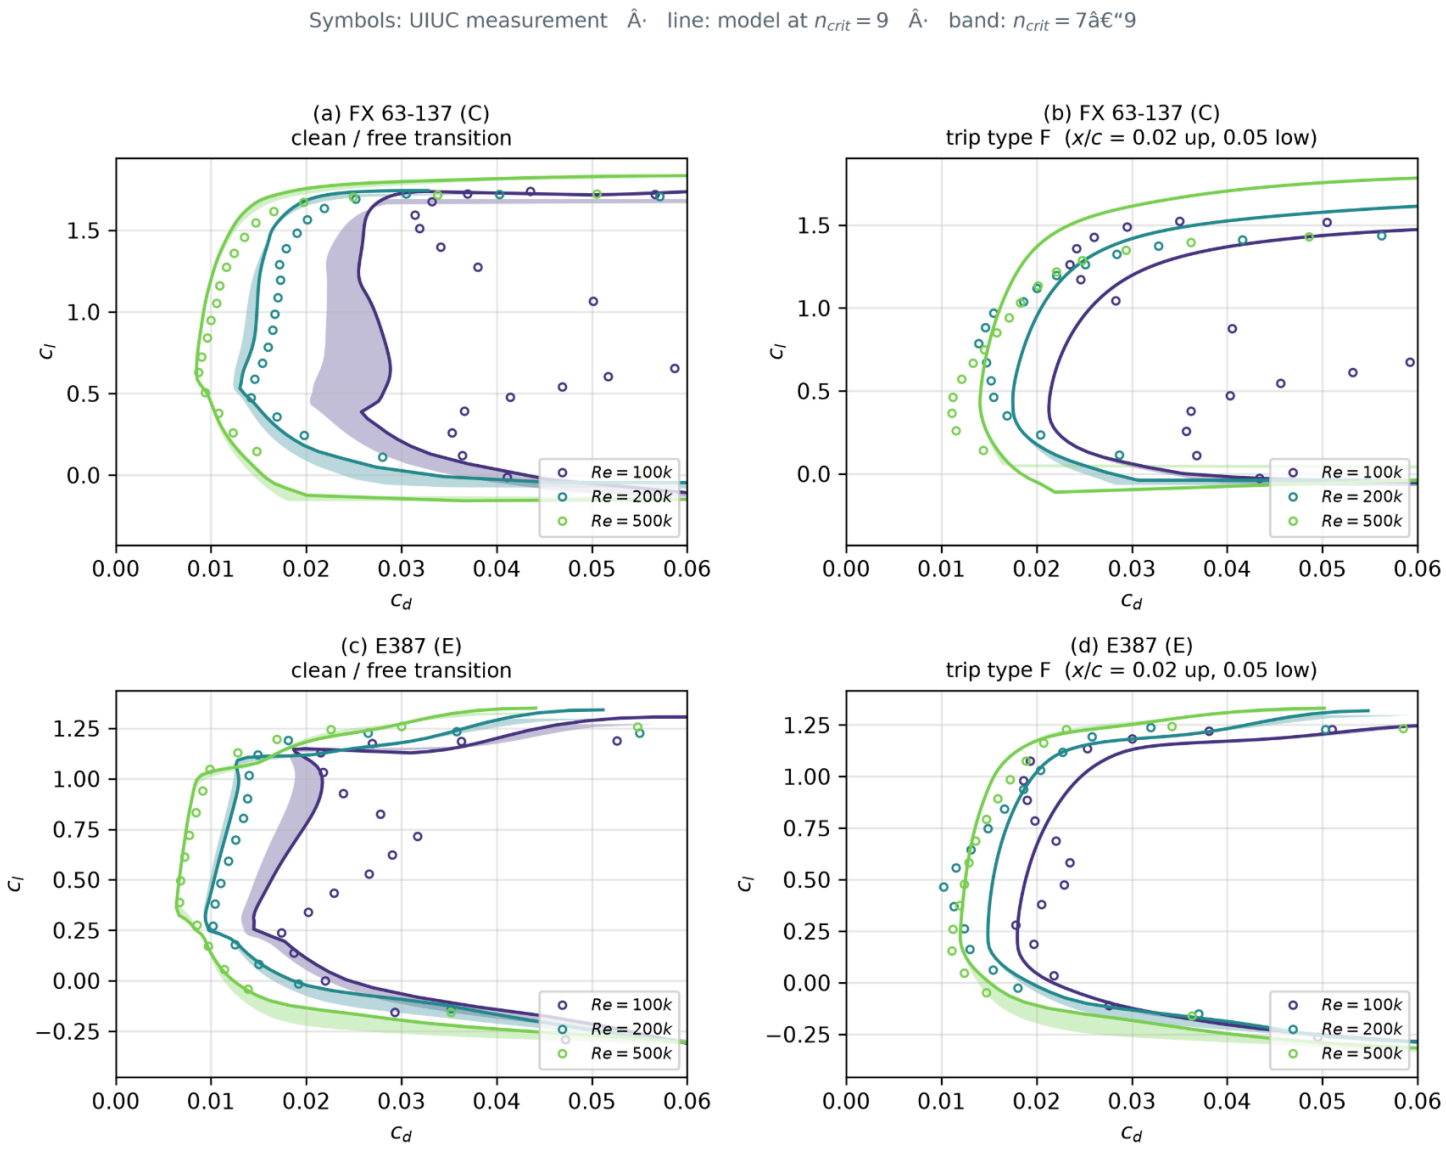


Figure 1

Comparison with the UIUC measurements of Ref. [9] in both configurations. Symbols are measurement, the line is the model at n_crit=9, and the shaded band spans n_crit=7–9, the tunnel's own disturbance

uncertainty. The band is visible in the clean panels (a, c) and all but vanishes in the tripped panels (b, d).

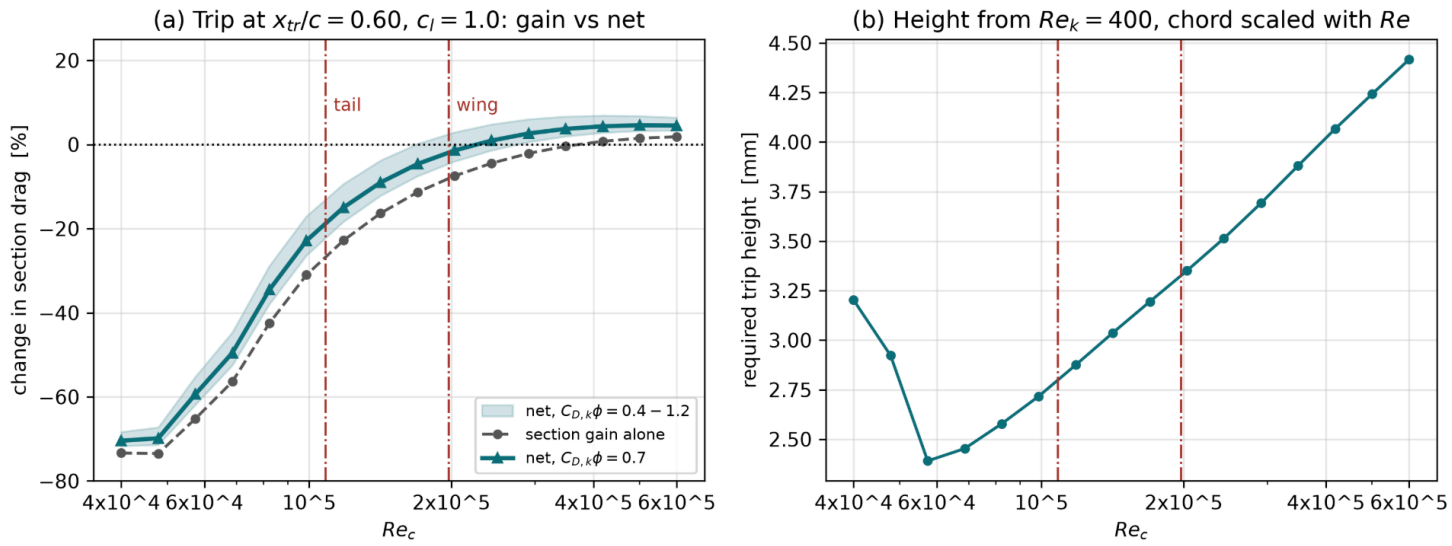


Figure 2

(a) Change in section drag from a trip at x_tr/c=0.60 at c_l=1.0: the section gain alone, and the net once the device is charged for, with the band spanning C_(D,k) ϕ=0.4–1.2. (b) Trip height required by Re_k=400. The operating points of the wing and the tail are marked.

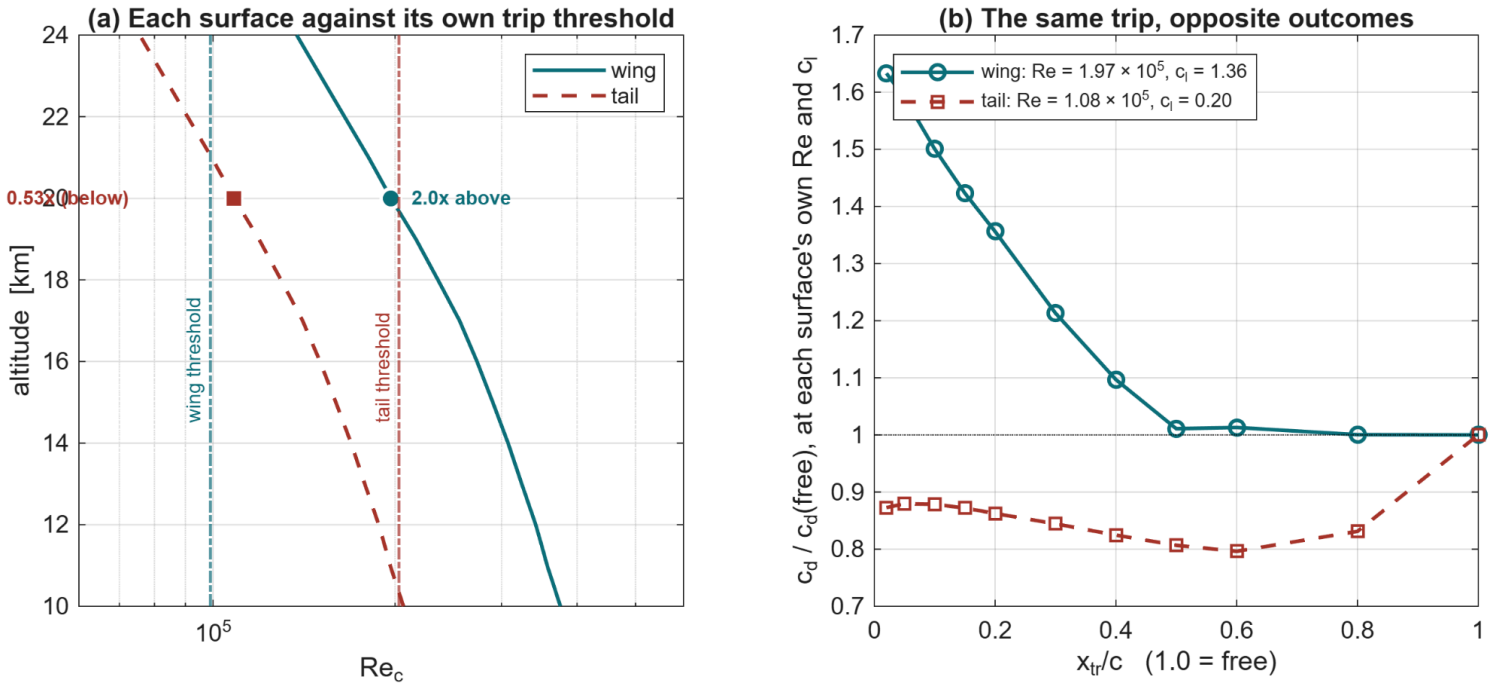


Figure 3

(a) Reynolds number of the wing and the tail against altitude for the baseline design, each with its own trip-profitability threshold. (b) The same trip evaluated at each surface's own operating condition: it is a loss on the wing at every location and a gain on the tail.

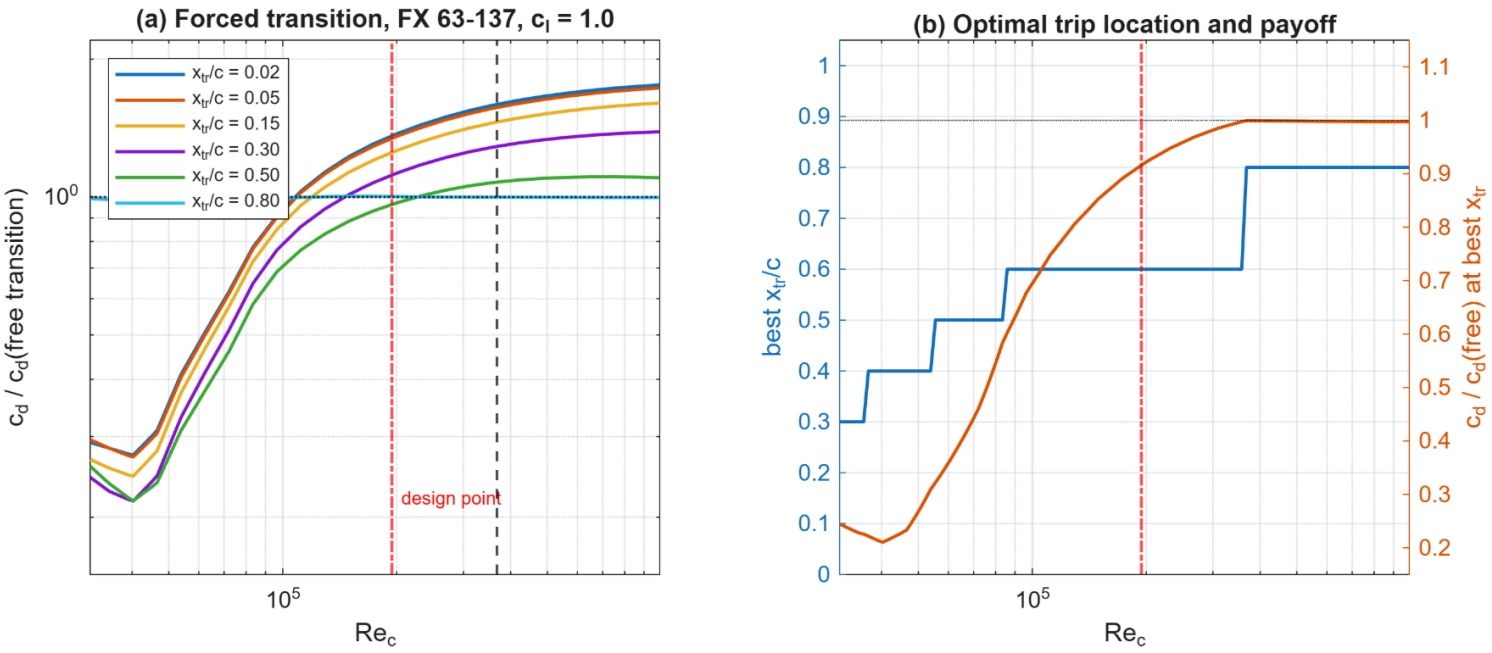


Figure 4

(a) Section drag relative to free transition for several forced transition locations, FX 63-137 at c_l=1.0. (b) The best location at each Reynolds number and the gain obtained there.

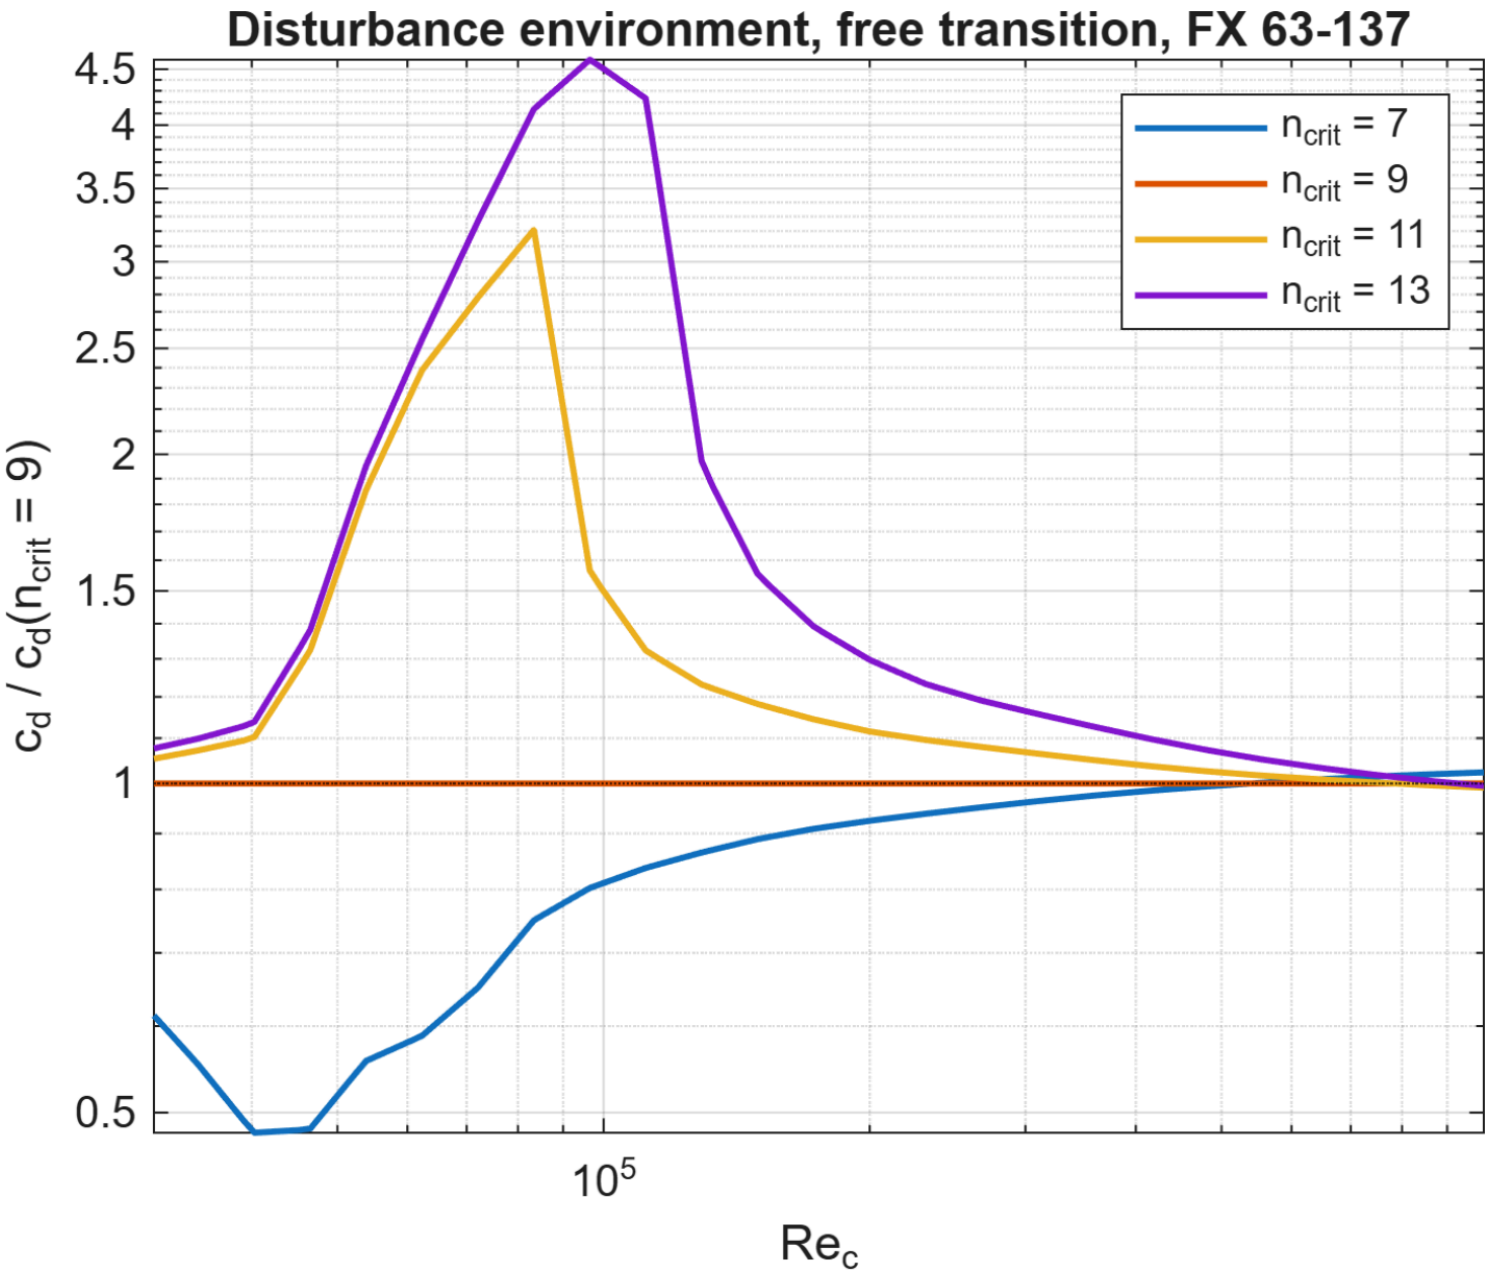


Figure 5

Section drag relative to n_crit=9 at free transition, FX 63-137 at c_l=1.0.

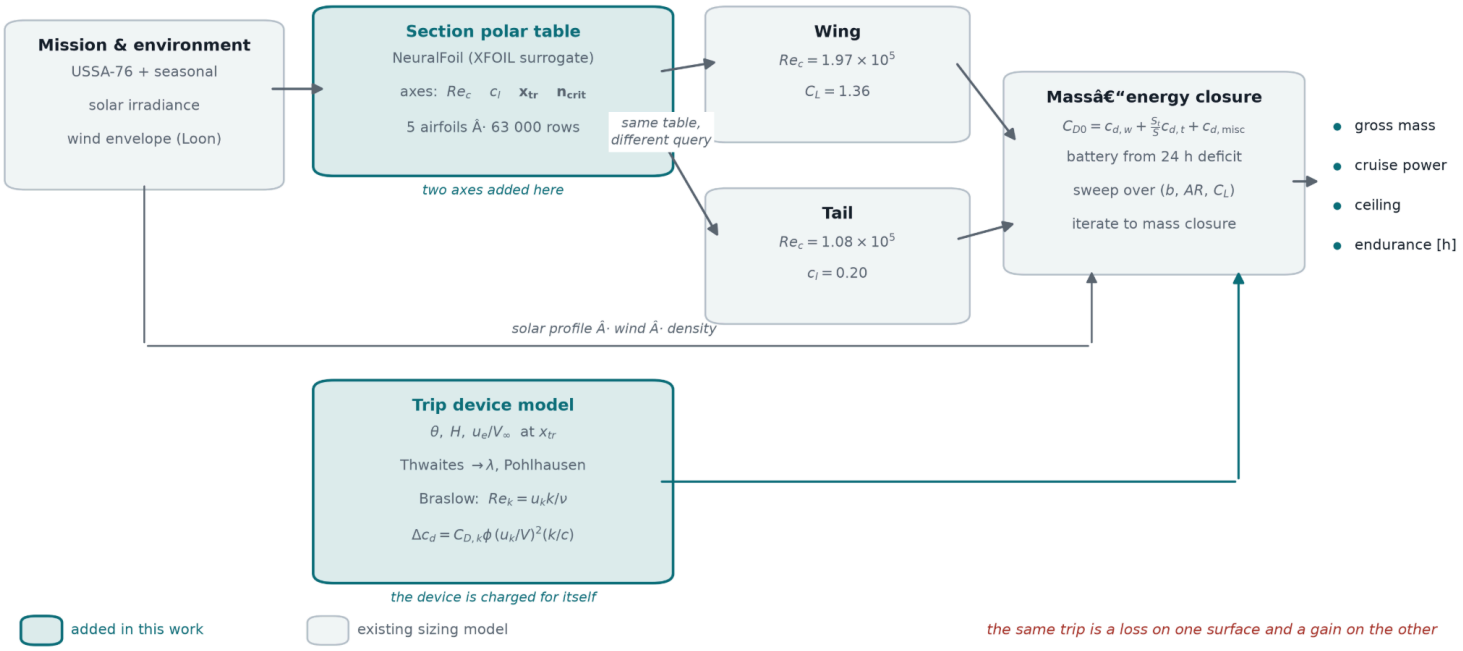


Figure 6

Method chain. The two axes added to the section table and the trip device model are the additions of this paper; the environment and sizing closure are those of the companion work.

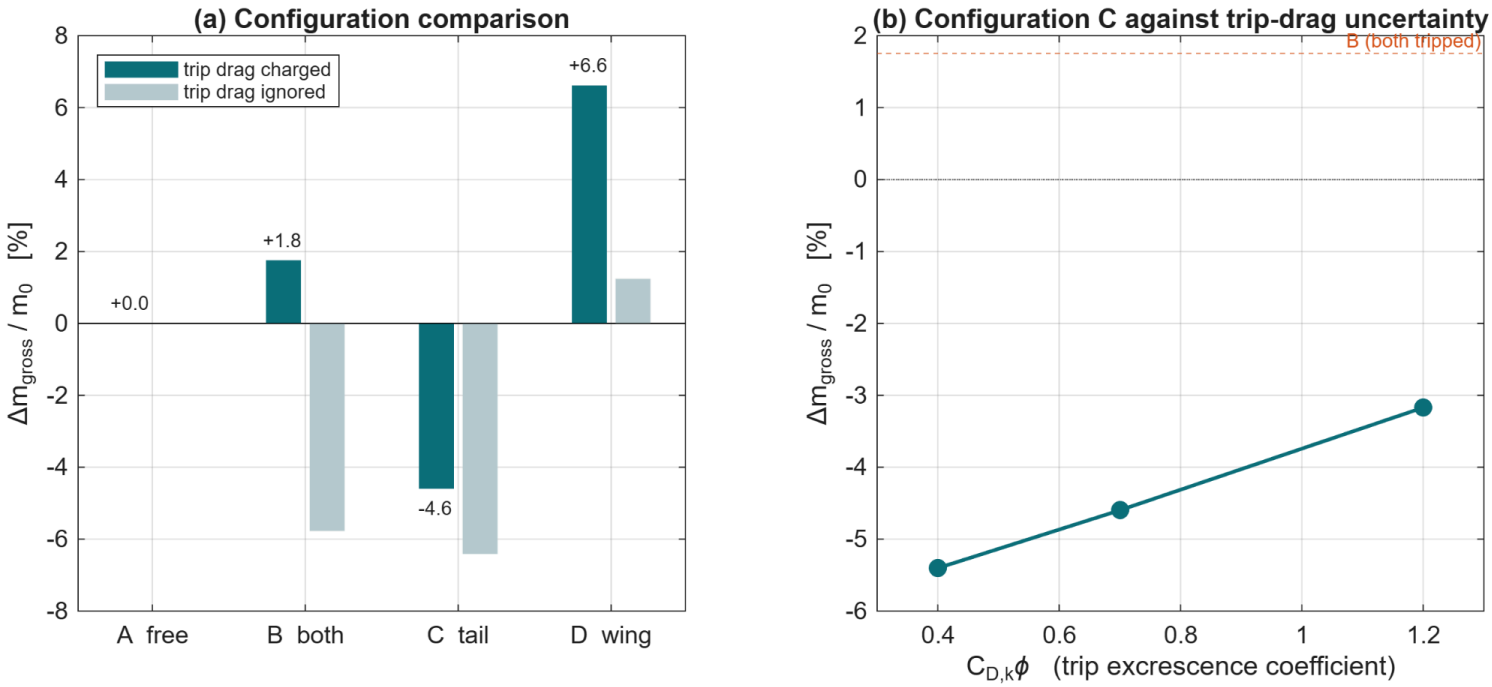


Figure 7

(a) Gross-mass change for the four configurations, with and without the device charged for. (b) Configuration C against the assumed excrescence drag coefficient; the dashed line is configuration B.

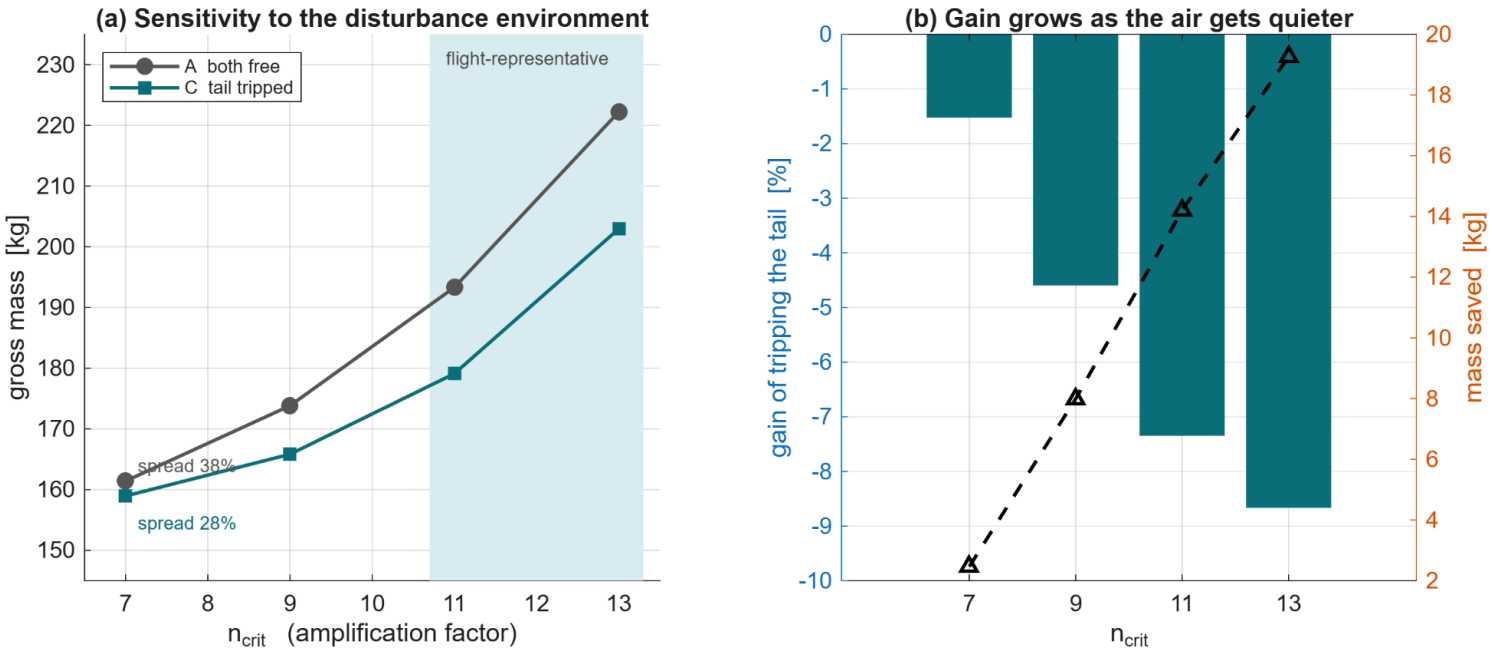


Figure 8

(a) Gross mass against the amplification factor for the free and tail-tripped configurations; the shaded region is the flight-representative band. (b) The gain from tripping the empennage, and the mass saved, against the amplification factor.

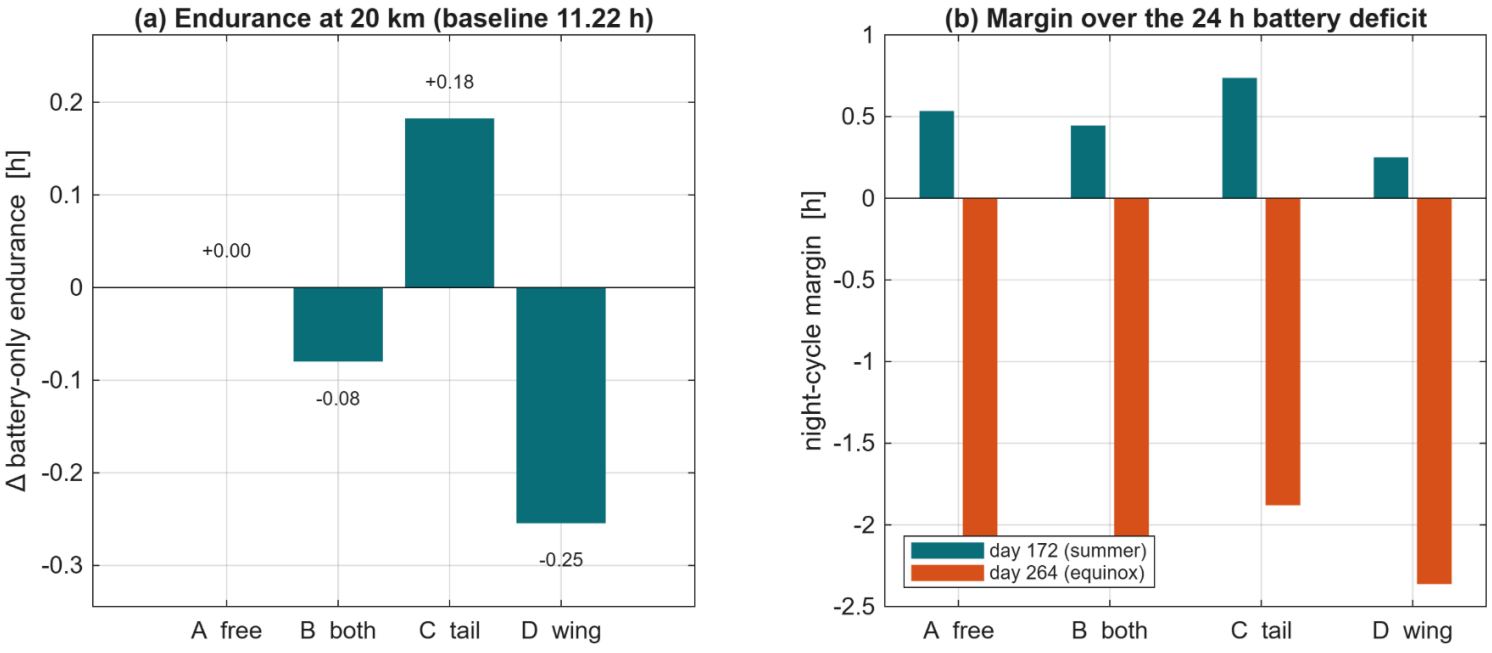


Figure 9

(a) Change in battery-only endurance at 20 on a fixed airframe. (b) Margin over the 24-hour battery deficit at the design day and at the equinox.