Very Low Earth Orbit from a Thermal Perspective

Thermal Analysis of a 16U CubeSat

Vincent Bull

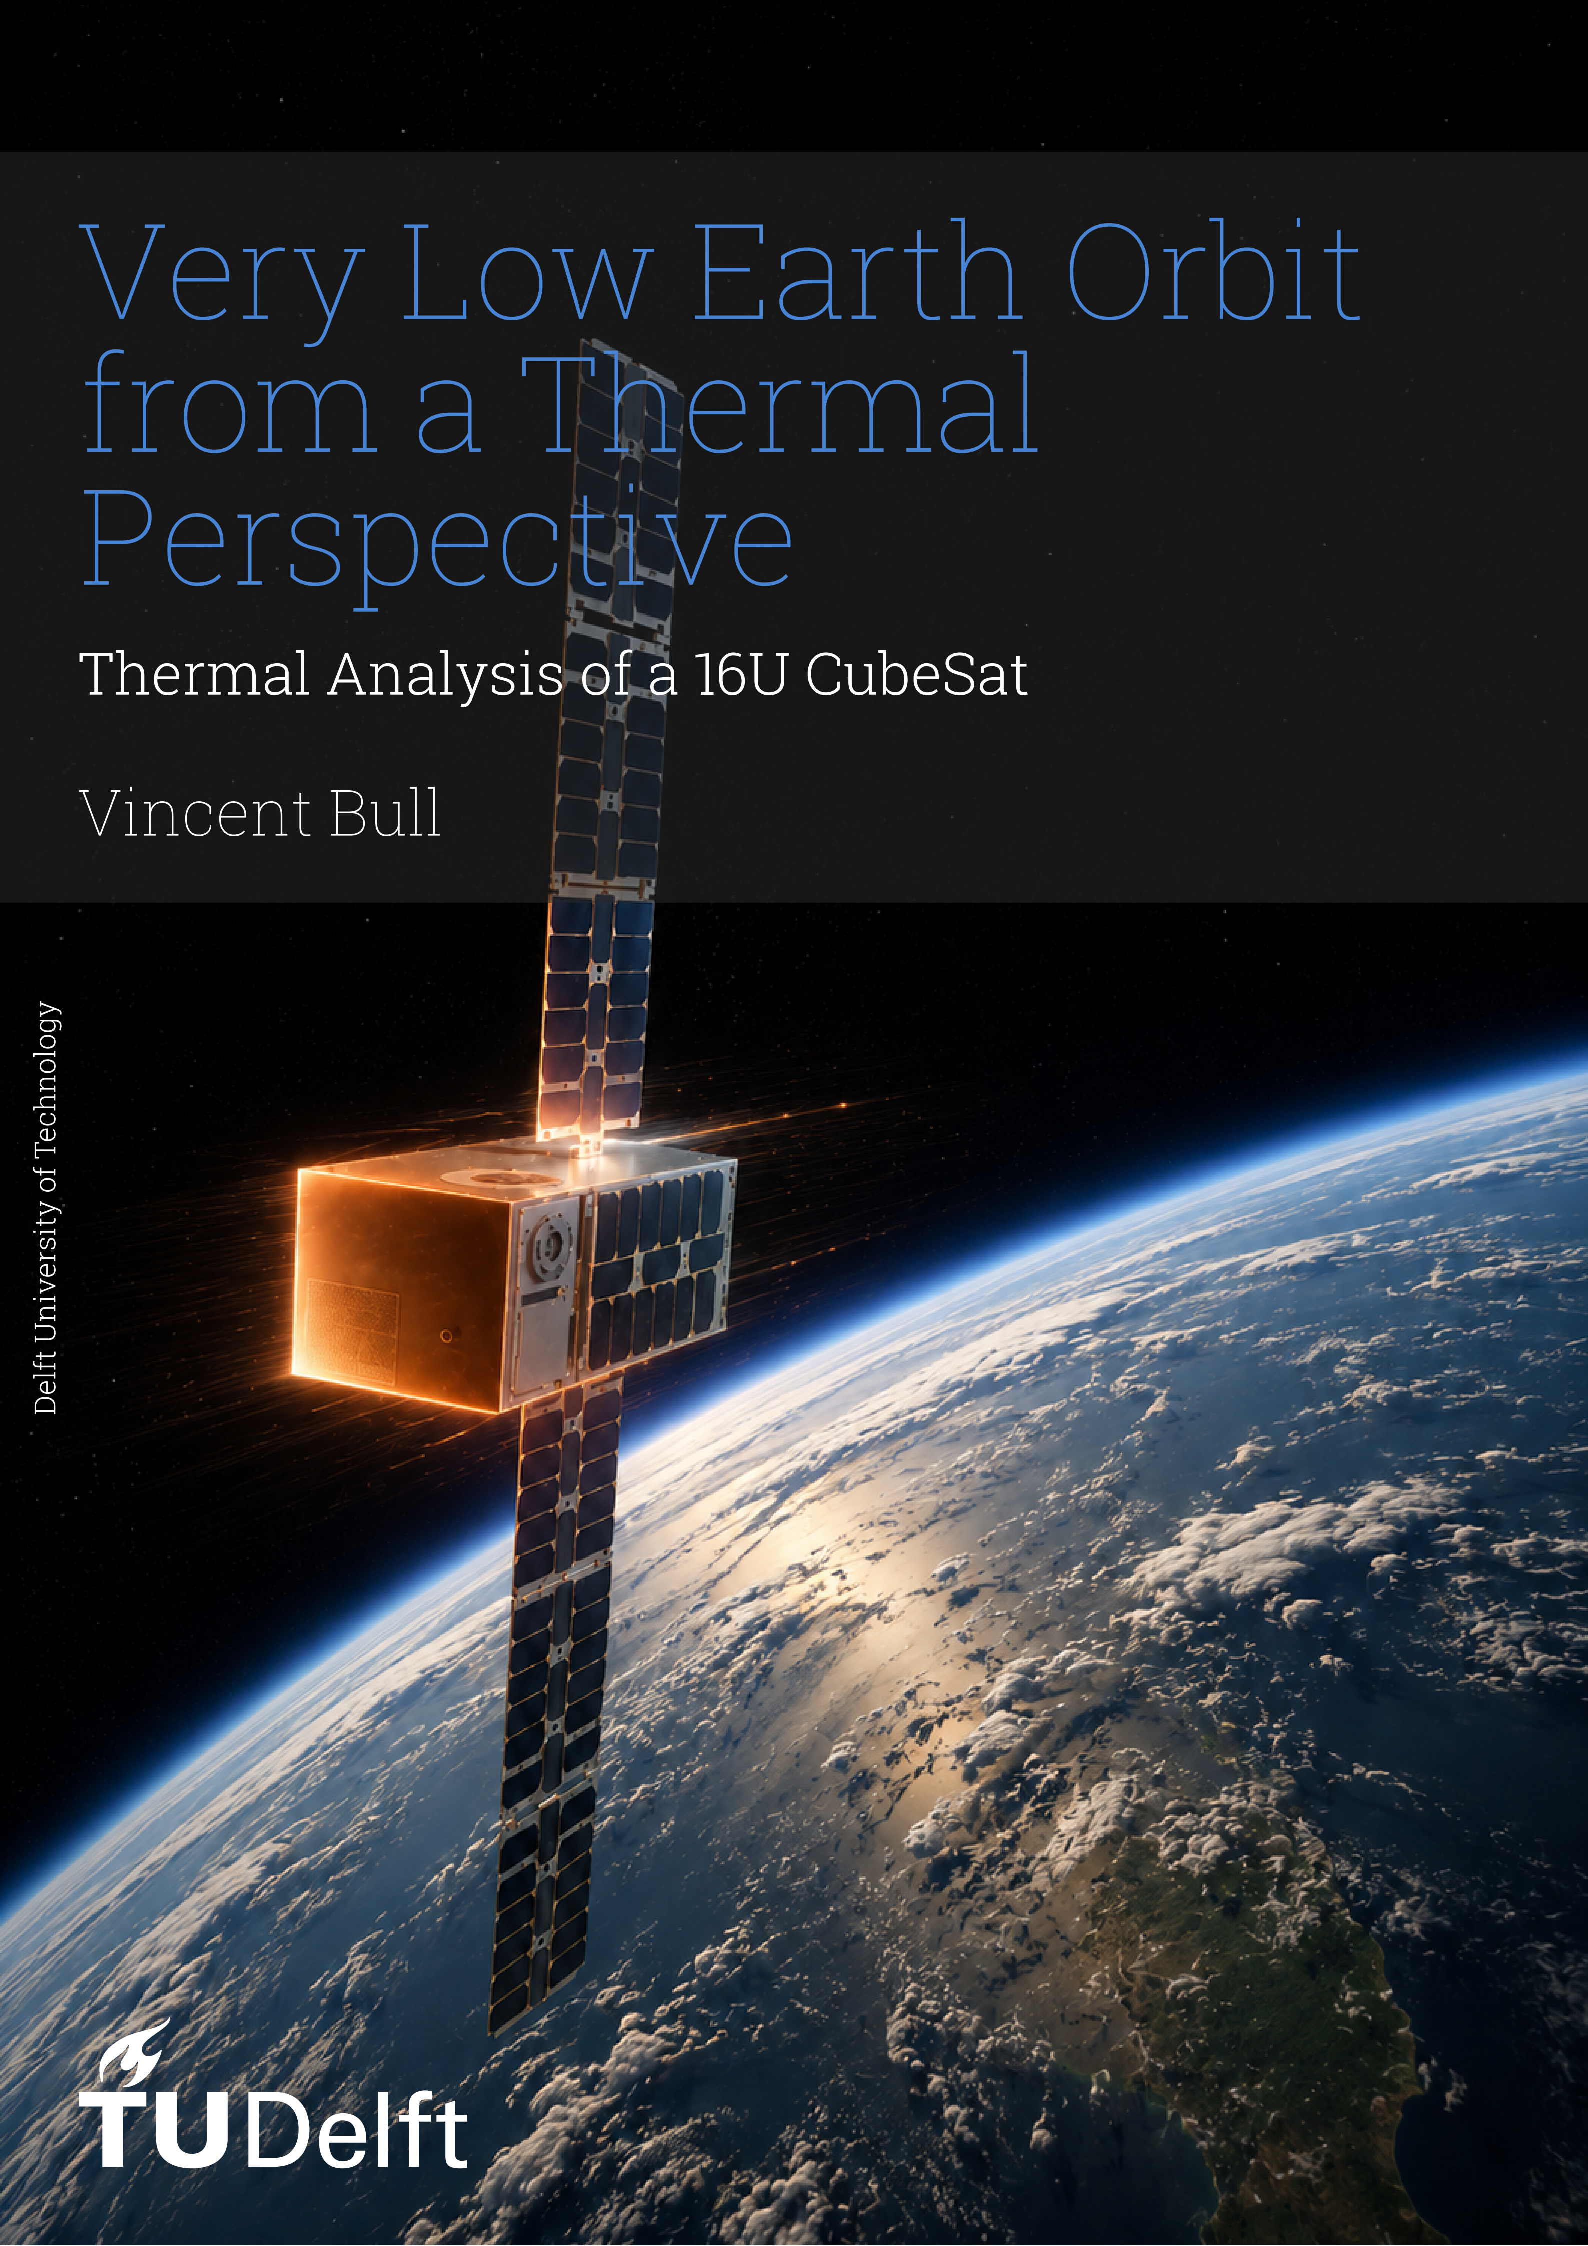

Delft University of Technology

TU Delft

This page is intentionally left blank.

Very Low Earth Orbit from a Thermal Perspective

Thermal Analysis of a 16U CubeSat

by

Vincent Bull

to obtain the degree of Master of Science

at the Delft University of Technology,

to be defended publicly on Friday, July 24, 2026 at 12:45.

Student number:	5326028	
Project duration:	October 20, 2025 – July 24, 2026	
Thesis committee:	Dr. ir. I. Uriol Balbin	TU Delft, supervisor
	Prof. C. Dransfeld	TU Delft, Chairperson
	Dr. S. Gehly	TU Delft, External Examiner
	Dr. P. Bambach	DLR, Supervisor

Deutsches Zentrum
für Luft- und Raumfahrt
German Aerospace Center

This page is intentionally left blank.

Preface

This thesis marks the conclusion of my efforts toward obtaining a Master of Science in Aerospace Engineering at Delft University of Technology. The project was conducted between October 2025 and July 2026, as a collaboration with the German Aerospace Center (DLR), which provided the mission context and simulation environment for a 16U CubeSat study in Very Low Earth Orbit. The project was as demanding as it was rewarding, and in hindsight the months went by faster than I expected.

I would like to thank Ines Uriol Balbin and Patrick Bambach for their guidance and support throughout this project, and everyone else at DLR RSC3 who contributed their time, expertise, and feedback along the way.

I am grateful to my family, my mother, father, and siblings, for their steady support throughout my studies. Thank you to my friends for the encouragement and the necessary distractions along the way. Finally, I want to thank my partner, whose support carried me through the more stressful stretches of this project.

Vincent Bull
Delft, July 2026

This page is intentionally left blank.

Abstract

Very Low Earth Orbit (VLEO) improves the geometric performance of optical payloads and increases communication link margins relative to conventional Low Earth Orbit (LEO), but exposes spacecraft to aerothermal heating, aerodynamic disturbances, and large environmental uncertainty. These effects are particularly challenging for CubeSat-class platforms, where mass, volume, and available thermal-control hardware are limited. This thesis assesses the thermal feasibility of a predominantly commercial-off-the-shelf 16U CubeSat platform operating down to a minimum perigee altitude of 130 km.

A simulation-based thermal-analysis framework is developed to construct conservative VLEO thermal cases and propagate them into spacecraft temperatures. Atmospheric, aerothermal, aerodynamic, and Earth-radiation inputs are defined for bounding hot and cold conditions. Analytical free-molecular aerothermal and aerodynamic models are benchmarked against DSMC reference data and used for rapid altitude- and attitude-dependent load evaluation. Earth albedo and infrared radiation are represented with a CERES-derived zonal-harmonics model after comparison with constant and ECSS-based alternatives. The resulting cases are analysed with a reduced-order LTspice thermal network and a spatially resolved COMSOL finite-element model. Morris screening, deterministic design sweeps, and surrogate-based uncertainty propagation are then used to size and verify the passive thermal configuration.

The baseline configuration is non-compliant in the hot case even without internal dissipation, identifying insufficient bus heat rejection as the limiting mechanism. A passive design combining a ram-facing heat shield, high-emissivity side panels, and thermally controlled external interfaces permits a deterministic payload duty cycle of 20 % at 130 km. The propagated maximum structure temperature is centred at 39.7 °C with a standard deviation of 0.85 °C and remains more than 7 °C below the 50 °C structure-to-payload interface requirement. No requirement exceedance is observed within the propagated uncertainty set. The results indicate that passive thermal control of a 16U COTS-based CubeSat is feasible for short-duration operation down to 130 km perigee, provided that interface conductances and subsystem-level thermal behaviour are verified by further modelling and test correlation.

This page is intentionally left blank.

Contents

Nomenclature

Abbreviations

Abbreviation	Definition
ADCS	Attitude Determination and Control System
AO	Atomic Oxygen
BDF	Backward Differentiation Formula
BMSP	Body-Mounted Solar Panel
CDF	Cumulative Distribution Function
CERES	Clouds and the Earth's Radiant Energy System
COMSOL	COMSOL Multiphysics
COTS	Commercial-Off-The-Shelf
DLR	German Aerospace Center
DOY	Day of Year
DSMC	Direct Simulation Monte Carlo
ECSS	European Cooperation for Space Standardization
FMH	Free-Molecular Heating
HWM	Horizontal Wind Model
ISO	International Organization for Standardization
LEO	Low Earth Orbit
LTAN	Local Time of the Ascending Node
LTspice	Linear Technology SPICE circuit simulator
MSFC	Marshall Space Flight Center
MSIS	Mass Spectrometer and Incoherent Scatter
NRLMSISE-00	Naval Research Laboratory Mass Spectrometer and Incoherent Scatter Extended atmospheric model, 2000 version
OBC	On-Board Computer
OLR	Outgoing Longwave Radiation
PCB	Printed Circuit Board
PWL	Piecewise Linear
QoI	Quantity of Interest
RMSE	Root-Mean-Square Error
RSP	Radiative Side Panel
SP	Solar Panel
SSO	Sun-Synchronous Orbit
SUP	Support Structure
SYS	Subsystem node
TCC	Thermal Contact Conductance
TCS	Thermal Control System
TOA	Top of Atmosphere
VLEO	Very Low Earth Orbit

Symbols

Symbol	Definition	Unit
a	Albedo coefficient	–
a_a	Average Earth albedo coefficient	–
A_i	Area of surface or node i	m^2
A_{ij}	Effective contact area between nodes i and j	m^2
A_{ref}	Aerodynamic reference area	m^2
A_p	Planetary geomagnetic activity index	–
B_{ij}	Gebhart factor from surface i to surface j	–
C_D	Drag coefficient	–
C_L	Lift coefficient	–
C_i	Thermal capacitance of control volume or node i	JK^{-1}
C_{mp}	Most probable molecular speed	$\mathrm{m\,s^{-1}}$
c_p	Specific heat capacity	$\mathrm{J\,kg^{-1}\,K^{-1}}$
EE_i	Elementary effect of input parameter x_i	dependent on QoI and parameter scaling
F_D	Drag force	N
F_L	Lift force	N
F_{ij}	View factor from surface i to surface j	–
$F_{E,i}$	Earth view factor of node i	–
$F_{\mathrm{SC}\rightarrow E}$	View factor from spacecraft surface to Earth	–
$F_{10.7}$	Daily 10.7 cm solar radio flux	sfu
$F_{10.7a}$	81-day averaged 10.7 cm solar radio flux	sfu
f_{c}	Bolted-interface contact factor	–
G_{ij}	Equivalent conductive conductance between nodes i and j	$\mathrm{W\,K^{-1}}$
h	Spacecraft altitude	m
h_{a}	Apogee altitude	m
h_{p}	Perigee altitude	m
h_c	Effective contact conductance	$\mathrm{W\,m^{-2}\,K^{-1}}$
h_{HS}	Heat-shield-to-structure contact conductance	$\mathrm{W\,m^{-2}\,K^{-1}}$
h_{MSP}	Body-mounted-solar-panel-to-structure contact conductance	$\mathrm{W\,m^{-2}\,K^{-1}}$
I_{yy}	Principal moment of inertia about the body-fixed y-axis	$\mathrm{kg\,m^{-2}}$
I_{zz}	Principal moment of inertia about the body-fixed z-axis	$\mathrm{kg\,m^{-2}}$
K_n	Knudsen number	–
k	Thermal conductivity; number of input parameters where stated	$\mathrm{W\,m^{-1}\,K^{-1}}$ or –
k_B	Boltzmann constant	$\mathrm{J\,K^{-1}}$
k_{aero}	Aerothermal heat-flux scale factor	–
k_α	Aerodynamic stiffness parameters in pitch	m^3
k_β	Aerodynamic stiffness parameters in yaw	m^3
L	Maximum degree of zonal-harmonics expansion	–
l_{ref}	Reference length used for Knudsen-number calculation	m
M	Mean molar mass of atmospheric gas	$\mathrm{kg\,mol^{-1}}$

Symbol	Definition	Unit
M_y	Aerodynamic pitch moment about the body-fixed y-axis	N m
M_z	Aerodynamic yaw moment about the body-fixed z-axis	N m
m_i	Mass of node or control volume i	kg
n	Number density	m^{-3}
P_k	Dissipated power of subsystem k	W
p	Static pressure	Pa
p_f	Probability of exceeding the interface requirement	–
$Q(\mathbf{x})$	Quantity of interest as a function of uncertain inputs	K
Q_{gen}	Internal heat dissipation	W
$Q_{out,rad}$	Radiative heat rejection to deep space	W
$\dot{Q}$	Heat flow rate	W
$\dot{Q} * i$	Total heat rate into surface or node i	W
$\dot{Q} * E, i$	Combined Earth-radiation heat input to node i	W
$\dot{Q} * \text{aero}, i$	Aerothermal heat input to node i	W
$\dot{Q} * \text{env}, i$	Environmental heat input to node i	W
$\dot{Q} * \text{in}, i$	Total heat input to control volume or node i	W
$\dot{Q} * \text{int}, i$	Internal heat input to node i	W
$\dot{Q} * \text{out}, i$	Heat removed from control volume or node i	W
q	Dynamic pressure	Pa
q_A	Earth albedo heat flux	$W\,m^{-2}$
$q * \text{aero}$	Aerothermal heat flux	$W\,m^{-2}$
q_{albedo}	Earth-reflected solar heat flux	$W\,m^{-2}$
q_{IR}	Earth-emitted infrared heat flux	$W\,m^{-2}$
q_{solar}	Direct solar heat flux	$W\,m^{-2}$
$\dot{q}$	Heat flux	$W\,m^{-2}$
$\dot{q} * \text{FM}$	Free-molecular heat flux	$W\,m^{-2}$
R	Universal gas constant	$J\,mol^{-1}\,K^{-1}$
R_E	Earth radius	m
$R * ij$	Radiative coupling factor between nodes i and j	m^2
r	Number of Morris trajectories	–
r_{CG}	Position vector of the spacecraft centre of gravity	m
r_i	Position vector of surface centroid i	m
r_{orbit}	Orbital radius	m
S_∞	Free-stream molecular speed ratio	–
S_{cold}	Combined cold-case Earth-radiation score	–
$s_k(t)$	Time-dependent switching factor of subsystem k	–
T	Temperature; simulated time interval where stated	K or s
T_i	Temperature of control volume or node i	K
T_∞	Free-stream neutral temperature	K
T_w	Wall temperature	K
T_{HS}	Heat-shield temperature	K
$T_{interface,hot}$	Hot-case structure-to-subsystem interface temperature	K

Symbol	Definition	Unit
T_{req}	Structure-to-payload interface temperature requirement	K
T_{RSP}	Radiative-side-panel temperature	K
T_{SP}	Solar-panel temperature	K
T_{STR}	Structure-node temperature	K
T_{struc}	Structure temperature	K
$T_{subsystem,cold}$	Cold-case subsystem temperature	K
T_{SUP}	Support-structure temperature	K
T_{SYS}	Subsystem-node temperature	K
T_{tgt}	Margin-adjusted deterministic design target	K
T_0	Reference temperature used for linearisation	K
t	Time	s
t_0	Reference epoch in ECSS Earth-radiation formulation	day
U	Spacecraft velocity relative to the neutral atmosphere	$\mathrm{m\,s^{-1}}$
U_{orb}	Spacecraft orbital velocity	$\mathrm{m\,s^{-1}}$
U_{rel}	Spacecraft velocity relative to the atmosphere	$\mathrm{m\,s^{-1}}$
U_{wind}	Zonal wind component	$\mathrm{m\,s^{-1}}$
V_{corot}	Atmospheric co-rotation speed	$\mathrm{m\,s^{-1}}$
V_{wind}	Meridional wind component	$\mathrm{m\,s^{-1}}$
v_r	Most probable molecular re-emission speed	$\mathrm{m\,s^{-1}}$
x_i	Input parameter i in screening or uncertainty propagation	variable
$\mathbf{x}$	Vector of input parameters or uncertain inputs	variable
Y	Model output used in Morris screening	K
$Y_{\ell 0}$	Zonal spherical harmonic of degree ℓ and order zero	–
z	Standardised distance from propagated mean	–
α	Angle of attack; solar absorptivity where stated	rad or –
α_i	Solar absorptivity of surface or node i	–
α_{sp}	Side-panel solar absorptivity	–
$\bar{\alpha}$	Thermal accommodation coefficient	–
β	Yaw angle	rad
γ	Specific heat ratio	–
γ_i	Cosine of incidence angle in Sentman formulation	–
Δh	Smoothing interval for altitude-dependent switching	m
Δ_i	Morris finite perturbation applied to input parameter x_i	variable
ΔT	Temperature deviation from reference case	K
ΔT_{margin}	Temperature margin reserved between design target and requirement	K
ε	Infrared emissivity	–
ε_i	Infrared emissivity of surface or node i	–
ε_{sp}	Side-panel infrared emissivity	–

Symbol	Definition	Unit
η_i	Optical absorption coefficient used for flux-to-heat conversion	–
η_{pld}	Payload duty cycle	–
λ	Molecular mean free path	m
μ_i^*	Mean absolute elementary effect of input parameter x_i	–
ρ	Atmospheric mass density; albedo coefficient in ECSS notation where stated	$\mathrm{kg\,m^{-3}}$ or –
ρ_∞	Free-stream atmospheric density	$\mathrm{kg\,m^{-3}}$
$\rho_{\max}$	Hot-case upper atmospheric density envelope	$\mathrm{kg\,m^{-3}}$
$\rho_{\min}$	Cold-case lower atmospheric density envelope	$\mathrm{kg\,m^{-3}}$
σ	Stefan-Boltzmann constant; standard deviation where stated	$\mathrm{W\,m^{-2}\,K^{-4}}$ or variable
σ_i	Standard deviation of elementary effects of input parameter x_i	–
θ	Incidence angle	rad
θ_i	Incidence angle of surface i	rad
θ_{dist}	Wind-induced yaw disturbance angle	rad
$\theta_{\text{dist,max}}$	Maximum wind-induced yaw disturbance angle	rad
φ	Geocentric latitude	rad
$\omega_\oplus$	Earth angular velocity	$\mathrm{rad\,s^{-1}}$
$\omega_{\alpha,0}$	Undamped natural frequency of pitch motion	Hz
$\omega_{\beta,0}$	Undamped natural frequency of yaw motion	Hz
ω_β	Yaw disturbance angular frequency	Hz

1
Introduction and Literature Review

1.1. Introduction

Very Low Earth Orbit (VLEO), typically defined as the band of orbital altitudes below approximately 450 km, has emerged as a regime of growing interest for next-generation Earth-observation and scientific missions [1]. Operating closer to the Earth's surface improves the geometric performance of optical payloads, increases communication link margins, and shortens data-downlink latencies relative to conventional Low Earth Orbit (LEO) [1]. These advantages are particularly relevant to small-satellite platforms, which can deliver high-performance observations at comparatively low system mass, volume, and launch cost [1].

The same proximity to the upper atmosphere that makes VLEO attractive, however, introduces a set of engineering challenges that differ substantially from those encountered in conventional LEO. Below roughly 200 km, atmospheric density increases by several orders of magnitude relative to typical LEO conditions and exhibits strong, partially predictable variability driven by solar and geomagnetic activity, causing aerodynamic drag and aerothermal heating to become design drivers for both attitude control and thermal management [2, 3]. Aerothermal heat fluxes scale approximately with the product of atmospheric density and the cube of the relative velocity, and high-fidelity predictions place stagnation-region fluxes on the order of several kilowatts per square metre near 120 km, a regime in which heating cannot be neglected as it can at higher altitudes [2]. In parallel, atomic oxygen (AO), the dominant atmospheric species in the lower thermosphere, drives accelerated surface erosion and gradual drift in the optical properties of exposed thermal-control surfaces [3]. The residual atmosphere further introduces heat-transfer mechanisms and aerodynamic forces that are negligible at higher altitudes [3, 4]. Spacecraft operating in VLEO must therefore withstand a coupled combination of aerodynamic, aerothermal, and material-degradation effects that remain only partially characterised by current models [1, 3].

In this context, the German Aerospace Center (DLR) is conducting a study to evaluate the feasibility of commercial-off-the-shelf (COTS) platforms for sustained operations in the lower thermosphere. The present thesis is carried out within the scope of that study and centres on a conceptual 16U CubeSat configuration designed to operate in the 130-300 km altitude band, with the dual objective of generating flight-representative data for the validation of atmospheric-density, aerodynamic, and thermal-load models in this regime and of characterising the operational limits of COTS components at these altitudes.

Architecturally, the conceptual platform adopts a predominantly COTS-based 16U CubeSat configuration. An aluminium primary structure integrates the avionics, communications, propulsion and power subsystems, together with deployable solar panels. A ram-facing honeycomb heat shield, additively manufactured from a carbon-filled high-temperature polymer composite, shields the windward face from direct aerodynamic heating and atomic-oxygen exposure. Scientific instruments, dedicated to in-situ measurement of atomic-oxygen flux and the thermospheric particle environment, are accommodated in the ram-facing section of the spacecraft. With the exception of the custom heat shield and these payload instruments, the platform relies on commercially available subsystems whose required temperature ranges impose tight constraints on the achievable thermal environment. The CubeSat form factor further restricts the mass, volume, and electrical power available for thermal-control hardware, limiting the design toolbox available to address these constraints under the additional thermal loading [5].

Against this background, this thesis investigates the thermal feasibility of operating a representative 16U CubeSat within the lower thermosphere, and the implications for the design of its thermal control system. To establish the basis for this investigation, the following section reviews current thermal-modelling practice and TCS-design approaches for spacecraft operating in this regime, drawing on both flown missions and recent mission concepts. The gaps identified through this review then motivate the research objectives and questions formulated in Subsection 1.2.4.

1.2. Literature Review and Problem Statement

The following review surveys the state of thermal modelling and thermal-control design for spacecraft operating in VLEO. The discussion is organised around three questions: which missions have flown or been studied in this regime and what thermal data they provide; how the environmental and spacecraft responses are modelled in current practice; and what thermal-control concepts have been proposed for CubeSat-class spacecraft below 300 km. The review concludes by drawing these threads together into the problem statement and the research questions that guide the remainder of this work.

1.2.1. Prior VLEO missions and the state of thermal characterisation

Operations below approximately 300 km remain comparatively rare in the small-satellite literature. Among non-CubeSat platforms, ESA's Gravity Field and Steady-State Ocean Circulation Explorer (GOCE) flew at a mean altitude of ~260 km between 2009 and 2013 [6], and JAXA's Super Low Altitude Test Satellite (SLATS) descended to 167 km in 2017 [2, 7]. Direct in-flight measurements of aerothermal heating below 150 km nonetheless still come only from the 1976 Atmosphere Explorer-C campaign of Caruso and Naegeli [8], in which instrumented tungsten wires were carried through a 130 km perigee and used to validate the classical free-molecular heating formulation. Table 1.1 summarises the thirteen missions identified in the literature, spanning minimum altitudes from approximately 50 to 350 km across both CubeSat and non-CubeSat platforms.

More recent activity has focused on mission-concept studies rather than flown hardware. Adam et al. [9] described the RAMSES configuration at 214 km; Canosa Ybarra et al. [10] reported the SigmaSat 274 km design; Paugnat et al. [11] examined IonSat at altitudes between 250 and 280 km; Selvadurai [12] proposed ARCADE for operations near 350 km with an actively cooled infrared payload; La Marca et al. [13] presented the EarthNEXT mission targeting 302 km; and at the lower-altitude end, Hoffmann et al. [4] examined sustained operation at 115 km in the Daedalus assessment study. CubeSat-class platforms that did operate in or transit through the regime, such as Bailet et al.'s QARMAN [14], EntrySat [15], and missions that will, such as SOURCE [16], have focused primarily on demise survivability and propulsion-system characterisation. Long-duration thermal performance has not been addressed within this body of work. Across this body of work, thermal results are typically limited to steady-state estimates or to radiator sizing for a fixed assumption on incoming heat flux, and no validated, mission-level CubeSat thermal analysis below 200 km has yet been reported in the open literature. Separately, Table 1.1 identifies only one 16U-class platform among the reviewed missions, the concept-level EarthNEXT study [13], whose published description is limited to a preliminary mission definition at 302 km. The thermal behaviour of a 16U-class spacecraft, and in particular its behaviour at the sub-200 km perigee altitudes, therefore remains unaddressed in the open literature.

Table 1.1: Reviewed VLEO and re-entry missions.

Mission	Author(s)	Year	Launch	Regime	Altitude [km]	Satellite Type
Atmospheric Explorer C	Caruso et al. [8]	1976	1973	VLEO	130	Not CubeSat
GOCE	Floberghagen et al. [6]	2011	2009	VLEO	~260	Not CubeSat
QARMAN	Bailet et al. [17]	2011	2020	Re-entry	50–70	3U CubeSat
EntrySat	Martella et al. [15]	2017	2019	Re-entry	~100	3U CubeSat
SASSI2	Pikus et al. [18]	2017	2019	VLEO	~100	3U CubeSat
RAMSES	Adam et al. [9]	2019	N/A	VLEO	214	Not CubeSat
IonSat	Paugnat et al. [11]	2020	2026	VLEO	250–280	6U CubeSat
Daedalus	Hoffmann et al. [4]	2020	N/A	VLEO	115	Not CubeSat
SigmaSat	Canosa Ybarra et al. [10]	2022	N/A	VLEO	274	8U CubeSat
SOURCE	Galla et al. [16]	2022	2026	Re-entry	<200	3U+ CubeSat
SLATS	Kimoto et al. [7]	2022	2017	VLEO	167.4	Not CubeSat
ARCADE	Selvadurai [12]	2023	2023	VLEO	350	27U CubeSat
EarthNext	Marca et al. [13]	2025	TBD	VLEO	302	16U CubeSat

1.2.2. Modelling the VLEO environment and the spacecraft response

The thermal environment in VLEO is shaped by four distinct physical contributions: atmospheric properties, aerothermal heating, Earth-radiation fluxes, and the conductive and radiative coupling internal to the spacecraft. Each of these is treated in current practice at a different and largely uncoordinated level of fidelity.

The thermosphere between roughly 100 and 450 km is characterised by strong spatio-temporal variability, with density fluctuations reaching ~40% at 200 km [3]. The empirical NRLMSISE-00 model of Picone et al. [19], recommended as a baseline by ISO 14222 [20], remains the standard atmospheric input for VLEO design and is typically used together with the DTM-2013 and JB2008 thermosphere models [21, 22]. Its accuracy below approximately 250 km is nevertheless limited: Kimoto et al. [7] found that NRLMSISE-00 overestimates SLATS-measured density by roughly 30% and atomic-oxygen flux by approximately 50%. Density itself responds strongly to solar EUV activity, geomagnetic disturbances, semi-annual cycles, and high-altitude winds [3, 20], and these variability sources propagate directly into the predicted aerothermal flux. Existing VLEO design studies, however, generally represent the atmosphere through a single nominal profile, and only a small subset acknowledges the implied uncertainty bounds [2, 16].

Aerothermodynamics in the free-molecular regime characteristic of VLEO can be estimated through a hierarchy of methods that trade fidelity against computational cost. At the lower-cost end, analytical formulations rooted in kinetic theory provide closed-form predictions at negligible computational cost. The principal examples are the free-molecular flow theory of Sentman [23] for aerodynamic forces and the related heat-transfer treatment of Schaaf and Chambre [24], of which the AE-C campaign of Caruso and Naegeli [8] remains the only in-flight validation. Their applicability is, however, restricted to simple convex geometries in which multi-bounce particle and surface interactions can be neglected, and their accuracy degrades as the flow transitions out of the strictly free-molecular regime at the lowest VLEO altitudes [25, 26]. At the high-fidelity end, the Direct Simulation Monte Carlo (DSMC) method of Bird [27] is the established standard for rarefied aerothermodynamics [28]. Hsieh et al. [29] applied DSMC to a 2U CubeSat at 200 km, Kaiser [30] used the PICLas code on the SOURCE configuration at 150 km, and the ultra-LEO study of Huang et al. [2] reports stagnation-region fluxes of up to 5100 W m^{-2} at 120 km. DSMC's accuracy comes at the cost of substantial computational time, which limits its use in design-loop iterations [26]. A range of intermediate methods has therefore been developed to bridge the analytical and DSMC tiers. Jin et al. [31] combined the free-molecular kinetic-theory equations with a test-particle Monte Carlo (TPMC) treatment of secondary particle and surface interactions to recover DSMC-comparable accuracy on complex geometries at a fraction of the cost. Walsh et al. [32] trained DSMC-informed surrogates to support shape-optimisation studies of VLEO platforms. Mehta et al. [33] developed an open-source panel-based hypersonic and rarefied-flow tool. More recently, Sethuraman et al. [26] proposed an improved analytical method that combines ray-tracing-style particle tracing with the free-molecular kinetic model, reportedly recovering DSMC-level accuracy on complex geometries above 120 km at greatly reduced run-time.

A consistent finding across DSMC and reduced-order studies is the strong dependence of local heat flux on the spacecraft's attitude. Hsieh et al. [29] and Kaiser [30] both report substantial shifts of peak local flux with even small AoA excursions, and Jiang et al. [34] show comparable AoA sensitivity in the drag and surface-load distribution. The majority of spacecraft-level thermal models nevertheless adopt uniform or stagnation-only boundary heating, and small damped attitude oscillations, expected even for aerodynamically stable configurations, have, to the author's knowledge, not been incorporated into any published spacecraft-level thermal analysis.

Earth-radiation fluxes are treated at different levels of fidelity in literature. The simplest representation, recommended by Gilmore [35], replaces albedo and outgoing longwave radiation (OLR) by constant orbital averages of ≈ 0.30 and ≈ 230 W m^{-2}, respectively. The ECSS long-term averaging method, derived from the Earth-radiation pressure model of Knocke et al. [36] and codified in [37], provides latitude-resolved monthly averages. More recently, Sasaki [38] has proposed higher-fidelity, CERES-derived zonal-harmonics expansions, although these have rarely been applied in published VLEO design studies. Because VLEO spacecraft view a larger apparent Earth disk than conventional LEO platforms, the choice between these formulations can introduce orbit-averaged surface-temperature differences of several degrees, particularly for high-inclination orbits that repeatedly traverse polar latitudes [38]. Despite this sensitivity, reviewed VLEO thermal-design studies have not, to the author's knowledge, considered the effect of this modelling choice on predicted-temperature uncertainty.

Once the external loads have been defined, the spacecraft thermal response is computed through one of several methods that span a hierarchy of fidelity and cost, analogous to the aerothermodynamic landscape discussed above. Closed-form analytical formulations derived from the lumped-capacitance balance, exemplified by the spinning-spacecraft solutions of Oshima et al. [39] and Gadalla [40] and the Fourier-based approach of Pérez-Grande et al. [41], remain useful for conceptual sizing, as in the SigmaSat feasibility study [10], but cannot represent complex geometries or non-uniform loading and are therefore unsuitable for detailed engineering design. Higher-fidelity predictions rely on the numerical discretisation of the spacecraft energy balance, and three competing families are in routine use. The finite difference method (FDM), formalised by Özişik [42], is computationally efficient and has been successfully applied at the CubeSat scale by Malde et al. [43]. The finite element method (FEM), most commonly deployed in its Galerkin form, accommodates complex geometries, multilayer assemblies, and temperature-dependent properties, as illustrated by Sundu and Döner [44] and by Díaz-Aguado et al. [45] on the FASTRAC nanosatellite. The finite volume method (FVM), demonstrated on a 1U CubeSat by Morsch Filho et al. [46], enforces strict energy conservation and is particularly well suited to problems coupling conduction with convective or aerothermal effects.

Lumped-parameter thermal-network models, such as those used by Corpino et al. [5], occupy an intermediate modelling tier between closed-form analytical estimates and detailed numerical thermal models. They remain the dominant tool for early CubeSat trade studies, because they retain the speed of the analytical models while accommodating arbitrary node topologies. At the upper end of computational efficiency for repeated evaluations, surrogate and data-driven techniques such as response surfaces [47], Kriging-based methods, and machine-learning predictors have been used to accelerate thermal optimisation [48–50].

Across all of these tiers, defensible predictions in VLEO additionally require attitude-dependent aerothermal boundary conditions and a credible representation of the conductive interfaces that dominate the internal temperature distribution. Thermal contact conductance (TCC) is widely recognised as the single largest uncertainty source in vacuum thermal models [51, 52]. Blanchete and Bah [52] recently reported temperature shifts of tens of kelvin between bounding TCC values for a CubeSat-scale assembly. Despite this sensitivity, formal uncertainty quantification is rarely reported in the VLEO design literature, whether through Monte Carlo sampling, polynomial chaos expansion, or the generalised Statistical Error Analysis method of Gómez-San Juan et al. [53]. Most published studies rely on conservative margin allocation instead.

1.2.3. Thermal control concepts for VLEO operations

Thermal-control concepts proposed for VLEO range from purely passive surface treatments to integrated two-phase loops. Passive concepts dominate the reviewed mission designs. Adam et al. [9] specified aluminium and Kapton finishes and a dedicated radiative panel for RAMSES at 214 km; the SOURCE configuration [16] combines AO-resistant coatings with multi-layer insulation (MLI) for its short-duration descent below 200 km; Bailet et al. [14] used a P50 cork ablator on QARMAN; and Hoffmann et al. [4] stacked a metallic heat shield, phase-change radiators, and conductive thermal straps for the Daedalus 115 km concept; although SLATS was not a CubeSat, the satellite overview accompanying Kimoto et al. [7] shows a heat shield mounted ahead of the ram-facing structure. La Marca et al. [13] additionally reported an extensive coating qualification campaign for combined AO, UV, and thermal-cycling exposure. Active concepts are far less common at the CubeSat scale and most often reduce to resistive heating for eclipse survival, as on both SOURCE [16] and EntrySat [15]. The exceptions are Selvadurai's ARCADE concept [12], which adopts a single-phase mechanically pumped fluid loop in support of an infrared payload, and the ultra-LEO concept of Huang et al. [2], which combines a loop heat pipe with a thermoelectric cooler and was validated in a thermal-vacuum facility reproducing aerothermal loading between 110 and 160 km.

Across the reviewed concepts, the proposed thermal-control configurations are qualified against nominal operating assumptions. Conservative margins are applied to absorb the residual modelling uncertainty, which is not propagated explicitly through the analysis. To the author's knowledge, no passive thermal-control configuration for a 16U-class CubeSat operating at sub-150 km perigee has been quantitatively verified through formal uncertainty propagation.

1.2.4. Synthesis, Problem Statement, and Research Questions

Taken together, the reviewed literature reveals three gaps. First, no 16U CubeSat has been studied in sustained operation below 200 km, and no CubeSat-class thermal analysis in this regime has been validated against flight data. Second, no clear consensus exists on how key VLEO-specific effects, such as aerothermodynamics and Earth-radiation properties, should be represented in a VLEO thermal-analysis framework, and how these choices affect the resulting modelling uncertainty. Third, uncertainties in both the environmental inputs and the design itself, such as thermal contact conductance, are rarely propagated through the analysis. Passive thermal-control configurations proposed for comparable VLEO missions are correspondingly qualified against nominal assumptions using conservative margins, without formal propagation.

While flight validation of the thermal analysis was deemed outside the scope of this work, since it would require in-orbit telemetry from a flown mission, this thesis addresses the remaining gaps by answering the following two research questions:

RQ1.

How do the VLEO-specific environmental conditions and spacecraft configuration determine the thermal response of a 16U CubeSat at low VLEO perigee, and how can these contributions be represented in a physics-consistent thermal-analysis framework suitable for design assessment?

RQ2.

What passive thermal-control configuration meets the structure-to-payload interface requirement of the conceptual 16U CubeSat platform at the minimum-perigee operating point, and what residual probability of requirement violation remains once the dominant environmental and modelling uncertainties are propagated through the selected design?

1.3. Reference Mission and Spacecraft

This section defines the mission profile, spacecraft configuration, and subsystem-level constraints used throughout the remainder of the thesis. The purpose is to fix, in a single reference location, the orbital cases, external geometry, allowable temperature ranges, and internal heat loads that subsequent chapters refer back to. Detailed subsystem descriptions and the full power-budget breakdown are deferred to Appendix A.

1.3.1. Mission profile

The conceptual platform is launched into a circular 550 km Sun-synchronous orbit (SSO) for commissioning. An impulsive manoeuvre subsequently lowers the perigee establishing an elliptical orbit. From this point the spacecraft performs a perigee-lowering burn of approximately 10 km every 16 revolutions, while the apogee is held at 500 km, until reaching a minimum perigee of 130 km. The resulting altitude evolution is shown in Fig. 1.1. The progressive reduction of the perigee causes the spacecraft to spend an increasing fraction of each orbit in the denser thermosphere, and the lowest-altitude segments correspondingly drive both the aerothermal and the internal-dissipation thermal loads addressed in this work.

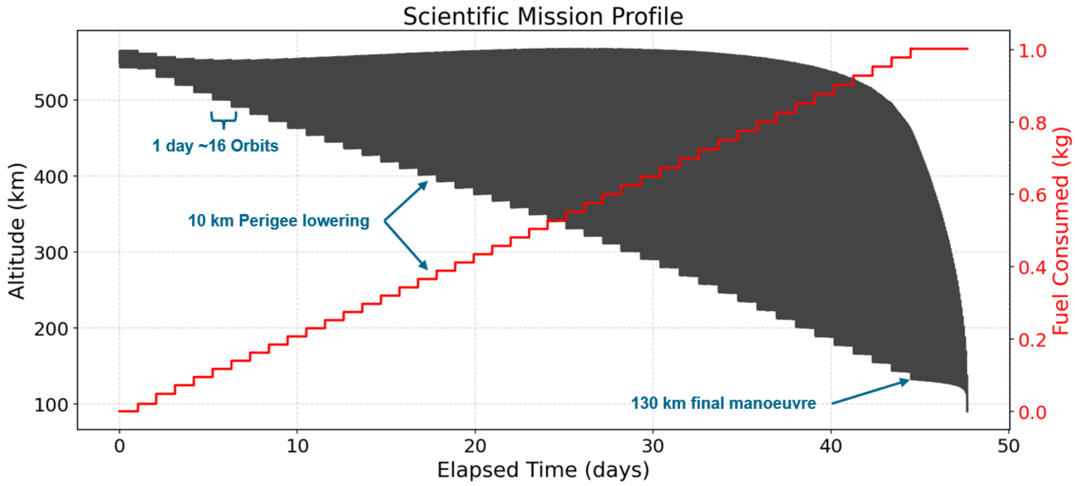

Figure 1.1: Simulated altitude and fuel-mass consumption over the reference mission timeline (elapsed time in days). Perigee is lowered in 10 km burns every 16 revolutions from the 550 km SSO commissioning orbit to the 130 km minimum-perigee condition, with apogee varying between 500-550 km. Cumulative fuel mass consumed is shown on the right-hand axis.[54]

Two SSO configurations are considered to bound the orbital thermal environment. Both have an inclination of 98° and are evaluated at the spring-equinox solar geometry with a solar constant of 1367 W m^{-2} [35]. The first is a dawn-dusk orbit with a Local Time of the Ascending Node (LTAN) of 06:00, which remains nearly continuously sunlit and defines the thermal hot case. The second is an LTAN 12:00 orbit, which exhibits eclipse durations of approximately 35-40 min per revolution and defines the thermal cold case. The two illumination patterns are illustrated in Fig. 1.2.

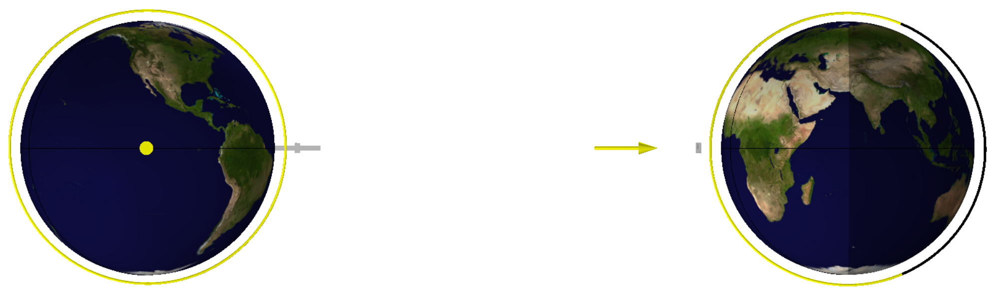

(a) SSO 550 km, LTAN 06:00 (dawn-dusk). Yellow orbit arc: sunlit; yellow arrow: Sun-vector direction.

(b) SSO 550 km, LTAN 12:00. Yellow orbit arc: sunlit; black: eclipse; yellow arrow: Sun-vector direction.

Figure 1.2: Illumination conditions for the two Sun-synchronous orbit configurations at an altitude of 550 km. Sunlit and eclipse arcs of the orbit track are indicated in yellow and black, respectively, while the yellow arrow denotes the direction of the Sun vector. The LTAN 06:00 (dawn-dusk) orbit remains continuously sunlit, whereas the LTAN 12:00 orbit exhibits significant eclipse periods during each revolution.

1.3.2. Subsystem Thermal Requirements

The thermal design is sized against two categories of subsystem-level requirements. The first is the operational temperature range of each design-driving component, summarized in Table 1.2. The second is the maximum allowable structure-to-subsystem interface temperature for the same components, which is equally design-driving and is reported in full in Appendix A (Table A.4). The operational ranges set the cold-case constraint derived below, while the interface allocations set the hot-case constraint. Both sets of values are taken from publicly available manufacturer datasheets for the representative COTS components retained on the conceptual platform, and are reproduced in full in Appendix A.3. The battery pack is locally heated in flight and is therefore not part of the bus-level cold-case constraint. The PEEK-based heat shield operates well within its material range under all considered cases and is not part of the hot-case constraint.

Table 1.2: Operational temperature ranges of the design-driving subsystems retained for the thermal analysis.

Subsystem	Operational range [°C]
On-board computer (OBC)	−30 to +85
ADCS controller	−30 to +85
Payload instruments	−40 to +85
Bi-propellant thruster	−16 to +31
Solar panels	−101 to +101

Two derived requirements follow from this set and are used directly in the thermal sizing of the present work. The hot-case constraint is the structure-to-subsystem interface temperature requirement, bounded by the strictest interface allocation across the subsystems retained on the platform (Table C.4):

$$T_{\text{interface,hot}} \leq +50\ ^\circ\text{C}. \tag{1.1}$$

The cold-case constraint is set by the lowest chip-level operating limit among the subsystems coupled conductively to the primary structure, namely the OBC and the ADCS controller:

$$T_{\text{subsystem,cold}} \geq -30\ ^\circ\text{C}. \tag{1.2}$$

A 10 °C uncertainty margin is applied to each thermal requirement during deterministic sizing. This margin is selected to be consistent with the expected maturity of a phase-B thermal analysis, for which ECSS specifies the TCS design as "a coarse design, with overall spacecraft TMM where critical items are modelled explicitly" [55]. The corresponding design targets are therefore +40 °C for the bus structure in the hot case and −20 °C for the cold-case lower-temperature limit. The 10 °C difference between each design target and the underlying requirement is retained as the dispersion budget propagated in the reliability assessment in Chapter 4.

The bi-propellant thruster has an operational range narrower than the bus thermal envelope sized above and cannot be accommodated within the same passive configuration. It is assumed thermally insulated from the primary structure in the present analysis, and its operating-range compliance is identified as an open point requiring a dedicated thermal analysis once the thruster hardware is specified.

1.3.3. Internal power dissipation

Internal dissipation is treated through three operational modes:

- **Cold Mode (Safe Mode):** all non-essential subsystems off, critical systems at minimum power. This mode bounds the lower limit of internal heat generation and is used for cold-case survival assessment.
- **Hot Mode (Science Mode):** all payloads and supporting subsystems active during scientific measurement at low perigee. This mode bounds the upper limit of internal dissipation and is used for hot-case sizing.
- **Nominal Mode:** a time-varying profile used in transient analyses, defined by the duty-cycle assumptions of each operational segment along the orbit.

The total dissipated power in each stationary mode is summarised in Table 1.3. The full subsystem-level breakdown used to derive these totals is reproduced in Appendix A.

Table 1.3: Total internal heat dissipation in the stationary operating modes used as bounding inputs to the thermal analysis.

Operating mode	Total dissipated power [W]
Cold Mode (Safe Mode)	7.7
Hot Mode (Science Mode)	69.4
Nominal Mode	duty-cycle dependent

For transient simulations, the Nominal Mode is constructed from the orbit-resolved duty cycle of each subsystem. At the current stage of the design, the power budget permits the dawn-dusk (LTAN 06:00) orbit to spend approximately 97% of each revolution in payload acquisition mode, with the remainder allocated to daily downlink and one orbit-lowering manoeuvre per day. The LTAN 12:00 orbit is permitted to spend approximately 67% of each revolution in payload acquisition mode, 29% in safe mode, and the remaining 4% in manoeuvre and downlink activities. These duty cycles are inputs to the thermal analysis, and their feasibility under the resulting temperature predictions is itself one of the outputs evaluated in Chapter 4.

1.4. Methodology

The thesis follows a simulation-based methodology organised around the workflow shown in Fig. 1.3. Mission Analysis and Environmental Modelling establish the orbital, atmospheric, aerothermal, Earth-radiation, and attitude boundary conditions. These enter the Analysis and Design Loop, in which the thermal modelling framework, design space, and design-and-verification steps are iterated for each candidate configuration. Only the parameters collected in the Design Space, namely surface properties and thermal interfaces, are varied across iterations. All other quantities entering the loop are recomputed at each pass but are not themselves design variables.

Chapter 2 first develops the environmental input chain. The atmospheric state, aerothermal heat flux, aerodynamic stability behaviour, and Earth-radiation field are evaluated in conservative hot- and cold-case directions and consolidated into altitude-, attitude-, and time-dependent thermal boundary conditions.

Chapter 3 then defines the thermal modelling framework used to propagate these boundary conditions into spacecraft temperatures. A lumped-parameter LTspice network provides rapid system-level checks and heat-path interpretation, while a finite-element COMSOL model resolves the spacecraft temperature field and serves as the primary design-assessment model. Together these correspond to the Geometrical and Thermal Mathematical Model blocks of Fig. 1.3.

Chapter 4 uses these models to develop and verify the passive bus thermal configuration, corresponding to the Design and Verification block of Fig. 1.3. The unmodified baseline is first assessed to identify the limiting heat-transfer mechanism. Morris elementary-effects screening then ranks the influential inputs and separates environmental drivers from actionable design levers. The retained design parameters, the same surface-property and thermal-interface parameters collected in the Design Space of Fig. 1.3, are sized through targeted sweeps, and the selected configuration is finally assessed through surrogate-based uncertainty propagation to quantify the residual probability of exceeding the structure-to-payload interface requirement. For the resulting configuration, the maximum payload duty cycle compatible with this requirement is also determined.

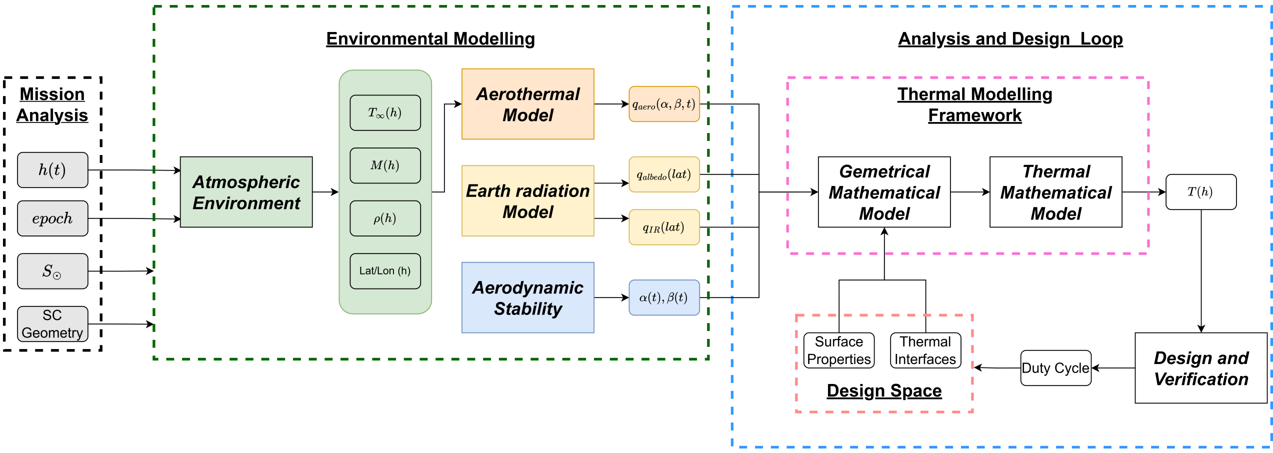

Figure 1.3: Overview of the analysis workflow. Mission Analysis (black) defines the orbital and configuration inputs used throughout: the altitude profile $h(t)$, epoch, solar irradiance $S_\odot$, and spacecraft geometry. Environmental Modelling (green) evaluates the atmospheric and positional state (NRLMSISE-00-based $T_\infty(h)$, $M(h)$, $\rho(h)$, and sub-satellite latitude/longitude) and propagates it into three parallel branches: the Aerothermal Model ($q_{aero}(\alpha,\beta,t)$), the Earth Radiation Model ($q_{albedo}(lat)$, $q_{IR}(lat)$), and Aerodynamic Stability ($\alpha(t)$, $\beta(t)$), the latter additionally driven by spacecraft geometry. These four boundary-condition streams enter the Analysis and Design Loop (blue), re-evaluated for each candidate configuration. Within it, the Thermal Modelling Framework (pink) combines the boundary conditions with the current design in the Geometrical Mathematical Model and solves the Thermal Mathematical Model for the temperature field $T(h)$. Design and Verification assesses $T(h)$ against the structure-to-payload interface requirement and determines the achievable payload duty cycle, returning both to the Design Space (red). Surface properties and thermal interfaces are the only parameters varied across iterations, closing the loop into the Geometrical Mathematical Model for the next pass.

2

VLEO Environment

This chapter defines the VLEO environmental inputs used by the thermal models. The relevant heat-transfer contributions are summarised in Fig. 2.1. Direct solar radiation is treated as a prescribed orbital input, while internal dissipation and radiative rejection are governed by the spacecraft configuration and thermal model. The chapter therefore focuses on the external contributions whose magnitude and variability are specific to the VLEO regime: aerothermal heating, Earth albedo, and Earth infrared radiation.

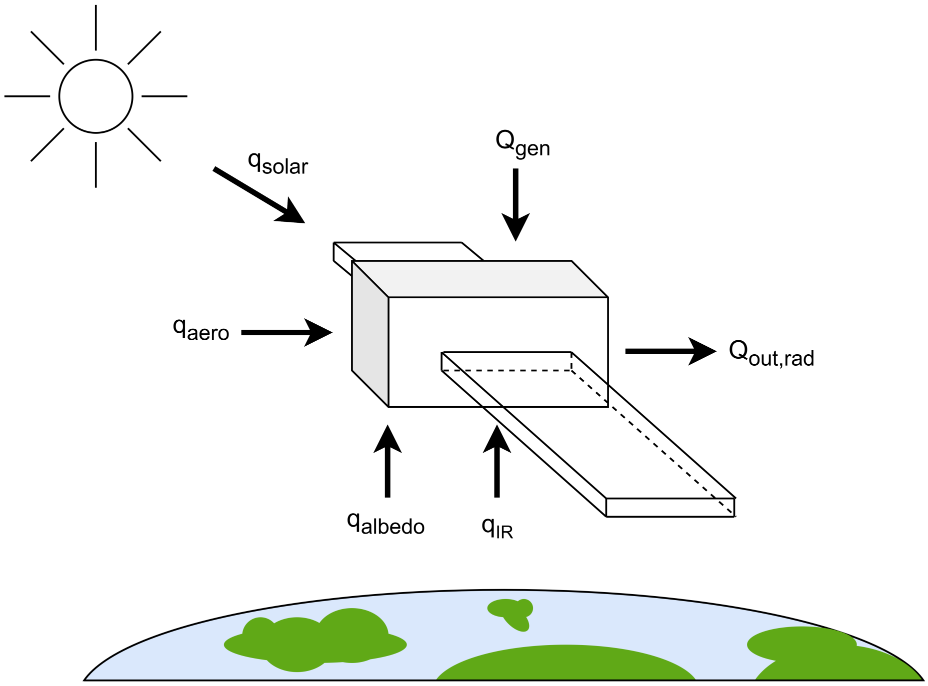

Figure 2.1: Schematic of the external and internal heat contributions acting on a CubeSat in VLEO. Lowercase q denotes heat flux $\mathrm{W\,m^{-2}}$: incoming solar radiation q_{solar}, Earth albedo q_{albedo}, Earth-emitted infrared q_{IR}, and aerothermal heating q_{aero}. Uppercase Q denotes heat flow W: internal dissipation Q_{gen} and radiative rejection $Q_{\mathrm{out,rad}}$.

The chapter follows the construction of the environmental input chain. Section 2.1 defines the atmospheric conditions, Section 2.2 establishes the flow regime, and Sections 2.3 and 2.4 derive the aerothermal and aerodynamic inputs required by the thermal model. Section 2.5 defines the Earth-radiation inputs, and Section 2.6 consolidates the resulting contributions into the thermal analysis cases used in the following chapters.

2.1. Atmospheric Environment

The atmospheric state sets both the magnitude and the variability of the aerothermal heat flux on spacecraft operating in VLEO, since this flux scales primarily with the local mass density and the spacecraft velocity relative to the neutral flow [2, 3]. Accurate modelling of the upper atmosphere, together with appropriate uncertainty margins, is therefore essential for defining conservative thermal design cases in this regime.

In accordance with ECSS space-environment recommendations [56], the neutral atmosphere is modelled using the NRLMSISE-00 empirical atmospheric model. NRLMSISE-00 provides altitude-dependent neutral temperature, species composition, and total mass density from ground level to the exosphere. The model belongs to the Mass Spectrometer and Incoherent Scatter (MSIS) family. It enforces hydrostatic equilibrium to ensure self-consistent temperature-density profiles, and improves on earlier MSIS versions through additional drag-derived constraints, updated accelerometer and ground-based datasets, and an explicit treatment of anomalous oxygen to better represent density at high altitudes [19]. The general applicability and limitations of empirical upper-atmosphere models of this class are described in ISO 14222 [20].

The NRLMSISE-00 model requires the specification of epoch, geodetic position, altitude, and indices describing solar and geomagnetic activity. Solar activity is characterised through the daily 10.7 cm solar radio flux $F_{10.7}$ and its 81-day running mean $F_{10.7a}$, while geomagnetic activity is represented by the daily planetary geomagnetic index A_p.

Future solar activity levels were obtained from the NASA Marshall Space Flight Center (MSFC) solar activity forecast [57, 58]. These forecasts build on the established eleven-year solar cycle and provide percentile-based predictions of the future $F_{10.7}$ radio flux. The best-estimate forecast corresponds to the 50th percentile, while the 95th and 5th percentile curves are obtained by adding or subtracting twice the standard deviation from the mean prediction. The probability that the true radio flux lies outside the 5th-95th percentile envelope is therefore approximately 10%. A representative example of such percentile-based $F_{10.7}$ forecasts is shown in Figure 2.2.

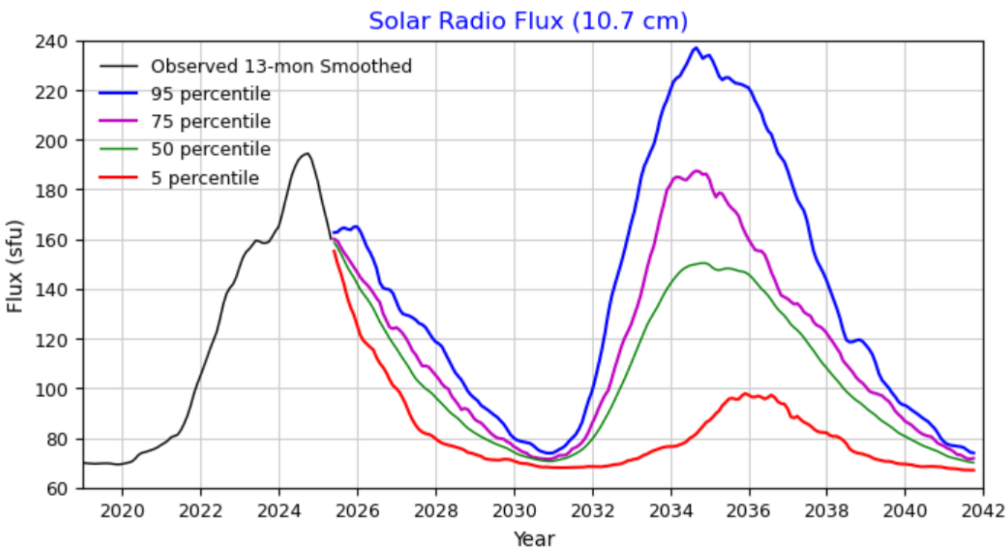

Figure 2.2: NASA MSFC percentile-based forecast of the solar radio flux at 10.7 cm ($F_{10.7}$). The 50th-percentile curve is the best-estimate baseline; the 5th- and 95th-percentile curves bound the forecast by subtracting/adding twice the forecast standard deviation. Within the assumed 2027–2029 mission timeframe, the hot case is taken at the 95th percentile in January 2027 ($F_{10.7} \approx 135.1$) and the cold case at the 5th percentile in August 2029 ($F_{10.7} \approx 71.1$).[57]

To define bounding atmospheric conditions for thermal and aerothermal analysis, the percentile-based framework adopted in earlier CubeSat entry and VLEO studies was applied [30]. The hot-case atmosphere was defined using the 95th percentile $F_{10.7}$ forecast together with the corresponding high-activity A_p index, while the cold-case atmosphere was defined using the 5th percentile forecast and the matching low geomagnetic activity level. Physical consistency between solar and geomagnetic activity was preserved by using percentile-matched datasets in each case.

2.1.1. Definition of Hot and Cold Atmospheric Cases

The mission is assumed to launch and operate between January 2027 and December 2029. Based on the adopted NASA MSFC forecasts, the hot case corresponds to January 2027 with $F_{10.7} \approx F_{10.7a} = 135.1$ and $A_p = 22.5$, while the cold case corresponds to August 2029 with $F_{10.7} \approx F_{10.7a} = 71.1$ and $A_p = 6.7$. The corresponding evaluation epochs are 2027-01-01 02:37:40.79 (UTC) for the hot case and 2029-08-02 01:12:43.19 (UTC) for the cold case. These dates represent, respectively, the expected maximum and minimum solar activity levels within the mission timeframe. For both cases the daily and 81-day-averaged radio flux values are taken as equal, $F_{10.7} = F_{10.7a}$, since the percentile-based MSFC forecasts provide smoothed, long-term activity levels [20].

Latitude and Longitude Enveloping

All three atmospheric quantities entering the aerothermal and aerodynamic formulations, mass density ρ_∞, neutral temperature T_∞, and mean molar mass M vary with latitude and longitude, driven by diurnal solar heating, geomagnetic energy deposition, and the resulting expansion and contraction of the thermosphere [3, 19]. Figure 2.3 shows the relative spatial variation of each quantity along the subsolar meridian ($\lambda \approx -40°$ at the hot-case epoch) at $h \in \{125, 250, 500\}$ km. Two trends are apparent from Fig. 2.3. First, spatial variability decreases with decreasing altitude across all three quantities, as is evident from Fig. 2.3 and consistent with established thermospheric behaviour [3, 19]. Second, density exhibits the largest relative variation at every altitude: at $h = 125$ km, $\rho_r = \rho_\infty/\bar{\rho}_\infty$ spans 0.83–1.09, while $T_r = T_\infty/\bar{T}_\infty$ spans 0.90–1.11 and $M_r = M/\bar{M}$ remains within 0.97–1.05. The limited variation in M_r at $h = 125$ km reflects the proximity of this altitude to the homopause (~100 km), above which diffusive separation of species gradually increases with altitude [20].

To obtain conservative, ground-track-independent profiles, the NRLMSISE-00 model was evaluated on a global latitude-longitude grid at each altitude between 500 km and 101 km in steps of 1 km, following the methodology of Kaiser [30]. For the hot case, the maximum density across all grid points was retained as the upper envelope $\rho_{\max}(h)$. For the cold case, the minimum density defined the lower envelope $\rho_{\min}(h)$. The enveloping procedure is applied to density only, as it is the dominant source of spatial variability. An independent latitude/longitude sweep of T_∞ and M is not performed, since their extrema need not coincide spatially with those of density. Combining independently-swept worst cases would yield a less probable, physically decorrelated scenario. Instead T_∞ and M are evaluated directly at the density-extrema locations used to define the envelope. The selected locations vary with altitude and do not represent an actual orbit profile.

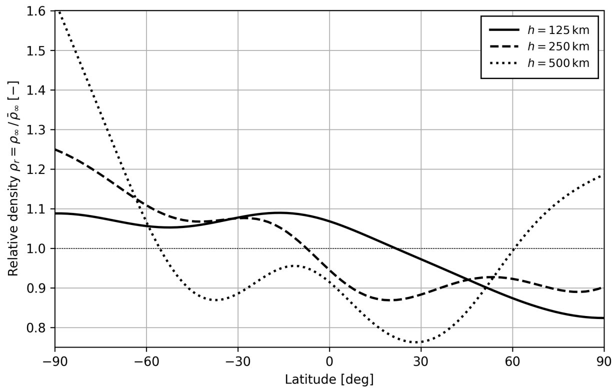

(a) Relative density $\rho_r = \rho_\infty/\bar{\rho}_\infty$.

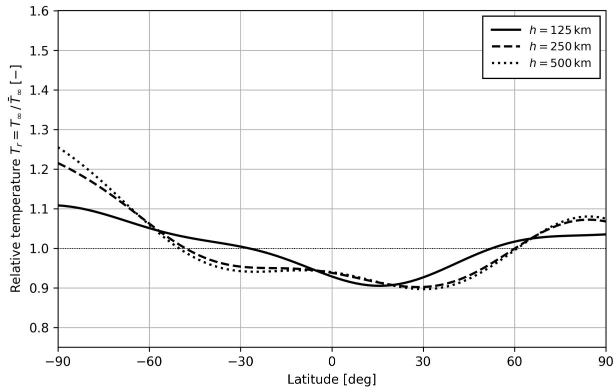

(b) Relative temperature $T_r = T_\infty/\bar{T}_\infty$.

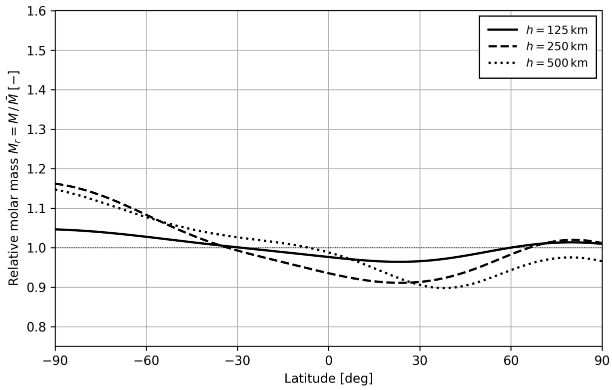

(c) Relative molar mass $M_r = M/\bar{M}$.

Figure 2.3: Spatial variability in mass density, neutral temperature, and mean molar mass along the subsolar meridian increases with altitude, evaluated using NRLMSISE-00 [19] at the hot-case epoch ($F_{10.7} = 135.1$, $A_p = 22.5$, January 2027, $\lambda \approx -40°$). At $h = 125$ km, close to the mission's minimum-perigee altitude: $\rho_r \in [0.83,\ 1.09]$; $T_r \in [0.90,\ 1.11]$; $M_r \in [0.97,\ 1.05]$.

Atmospheric Density Uncertainty

Additional density margins were applied following ISO 14222 [20]. A −15% reduction was applied to the cold-case profile, corresponding to the estimated mean-activity model uncertainty of NRLMSISE-00 stated in ISO 14222 [20]. A +100% increase was applied to the hot-case profile, corresponding to the short-term and local-scale density enhancements that can exceed the climatological model output under disturbed conditions [20]. In each case only the direction increasing the severity of the case at hand was retained.

Separate uncertainty margins are not applied to T_∞ or M. Both quantities enter the aerothermal heat flux only through the ratio T_∞/M, raised to the 3/2 power in the prefactor of Eq. (2.9), and via the most probable speed entering the speed ratio S_∞ of Eq. (2.4), shared with the aerodynamic stability model. At the density-extremal locations used to define the envelope, $h = 125$ km gives a $(T_\infty/M)^{3/2}$ deviation of −17.2% for the hot case and +13.8% for the cold case. The hot-case deviation is small relative to the +100% density margin applied above [20]. The cold-case deviation acts in the opposite direction to the −15% density margin and is comparable in magnitude, so it should not be read as additional conservatism beyond that margin. Furthermore, ISO 14222 provides uncertainty estimates for atmospheric density only [20]. The standard gives no corresponding bounds for T_∞ or M, and no standards-based justification therefore exists for applying analogous margins to these quantities.

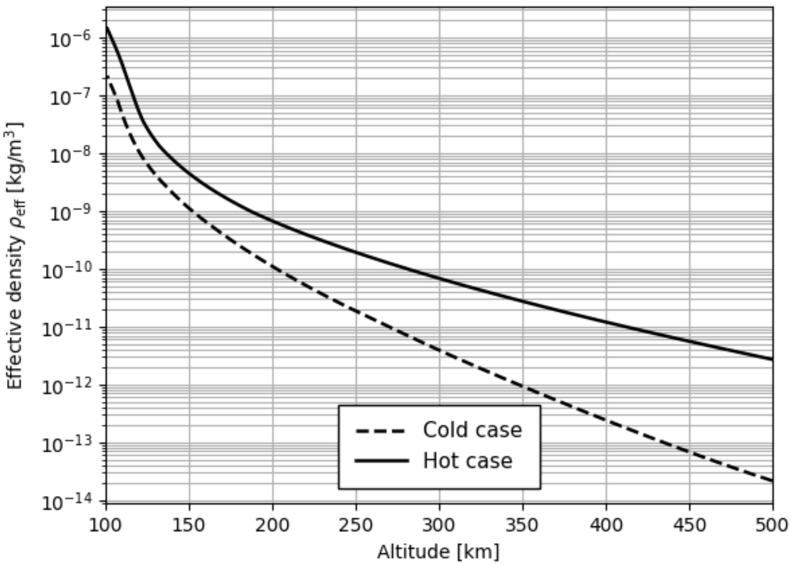

Figure 2.4: Hot- and cold-case effective atmospheric mass density ρ_{eff} versus altitude: NRLMSISE-00 [19] evaluated on a global latitude-longitude grid at the percentile-based hot/cold epochs, taking the grid maximum (hot) and minimum (cold) as the spatial envelope [30], with ISO 14222 [20] margins of +100% (hot) and −15% (cold) applied on top.

2.2. Flow Regime

With the atmospheric environment defined, the flow regime around the spacecraft must be characterised, since it determines which modelling framework is appropriate for computing aerodynamic and aerothermal loads. At VLEO altitudes the ambient density is several orders of magnitude lower than in conventional aeronautical applications. Intermolecular collisions consequently become increasingly rare, and the continuum assumption underlying classical fluid mechanics progressively breaks down [27].

The degree of rarefaction is quantified through the Knudsen number K_n, which Bird [27] defines as

$$K_n = \frac{\lambda}{l_{\mathrm{ref}}}, \tag{2.1}$$

where λ denotes the mean free path of the gas molecules and l_{ref} is a characteristic length of the spacecraft. The mean free path is the average distance travelled by a molecule between successive collisions and is given by kinetic gas theory, following Bird [27] as

$$\lambda = \frac{k_B T}{\sqrt{2}\,\pi d^2 p}, \tag{2.2}$$

where k_B is the Boltzmann constant, T the local gas temperature, d the effective molecular diameter, and p the static pressure. Using the ideal-gas relation $p = nk_BT$, the mean free path may equivalently be written as $\lambda = 1/(\sqrt{2}\,\pi d^2 n)$, which is the form most convenient when the number density n is taken directly from an atmospheric model output. In the thermosphere, atomic oxygen is the dominant species, but for the purpose of estimating K_n a representative N_2 molecular diameter is adopted [3].

Flow regimes are commonly classified according to the Knudsen number as summarised in Table 2.1. For $K_n \ll 1$ the flow is in the continuum regime and is well described by the Navier-Stokes equations. For $K_n \gg 1$ the flow is in the free-molecular regime, where intermolecular collisions can be neglected and the surface response is governed entirely by gas-surface interactions. The intermediate range, in which neither limit is strictly valid, is referred to as the transitional regime and requires direct solution of the Boltzmann equation, typically through particle-based methods such as Direct Simulation Monte Carlo (DSMC) [27, 28].

Table 2.1: Classification of flow regimes based on the Knudsen number.

Knudsen number	Flow regime
$K_n > 10$	Free molecular
$0.01 < K_n < 10$	Transitional
$K_n < 0.01$	Continuum

Figure 2.5 shows the computed Knudsen number as a function of altitude for the hot- and cold-case atmospheric profiles defined in Section 2.1, using a characteristic length corresponding to the spacecraft bus dimension along the nominal flow direction $l_{\text{ref}} = 0.4\,\text{m}$. The flow remains firmly in the free-molecular regime across most of the considered altitude range. Only at the lowest perigee altitudes does the Knudsen number approach the transitional boundary, crossing $K_n = 10$ at 131.6 km in the hot case and 117.3 km in the cold case.

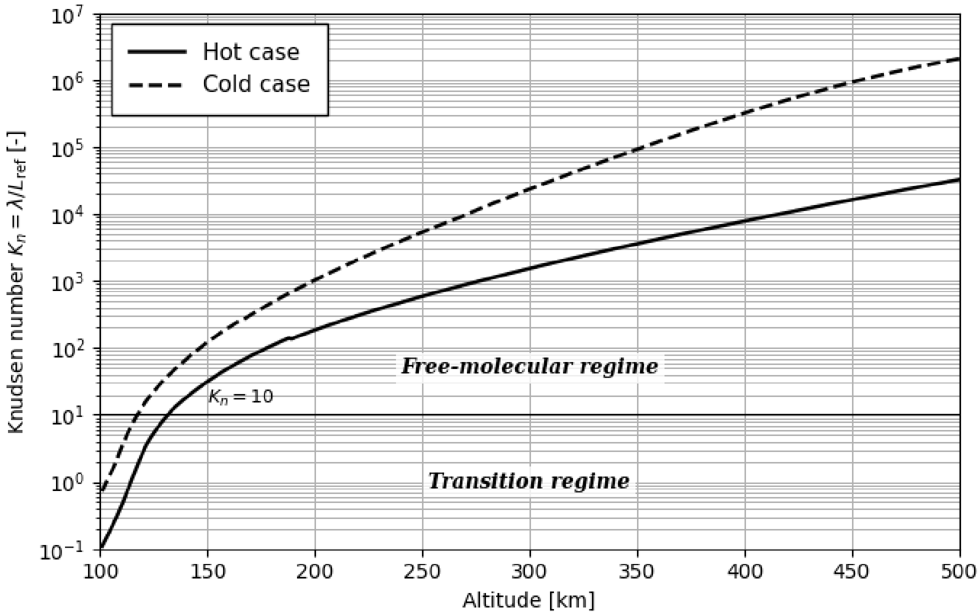

Figure 2.5: Knudsen number as a function of altitude for the hot- and cold-case atmospheric profiles. Horizontal lines indicate the regime boundaries defined in Table 2.1. The $K_n = 10$ boundary is crossed at 131.6 km (hot case) and 117.3 km (cold case), placing the mission's 130 km minimum perigee close to the transitional boundary in the hot case.

The computed Knudsen numbers include the conservative density margins introduced in Section 2.1 and therefore represent a lower bound on rarefaction. Since the mission remains at or above 130 km, the flow is rarefied throughout the relevant operating envelope and is free molecular over most of it. Free-molecular theory is therefore used for the analytical aerothermal and aerodynamic models in Sections 2.3 and 2.4. DSMC is retained as the reference method near the lowest altitudes, where the flow approaches the transitional boundary.

2.3. Aerothermal Loads

In the free-molecular regime established in Section 2.2, the aerothermal heat flux is governed by atmospheric density, temperature, relative velocity, spacecraft attitude, and gas-surface interaction properties. For the CubeSat geometry considered here, the incidence-angle dependence produces unequal, per-face-uniform heating that must be evaluated face by face.

An analytical free-molecular heat-transfer model is used as the primary design tool, as it enables rapid evaluation over altitude and attitude. DSMC results generated within the project are used as reference data to benchmark the analytical predictions at representative low-altitude conditions.

2.3.1. Analytical Free-Molecular Heat-Flux Model

In the free-molecular regime, continuum heat-transfer formulations are not applicable and the surface heat flux must be derived from kinetic theory. The free-molecular heating (FMH) formulation of Caruso [8] is adopted here. It gives the heat flux per unit surface area as

$$\begin{aligned}\dot{q}_{\mathrm{FM}} = \frac{\bar{\alpha}\,\rho_\infty}{\sqrt{2\pi}}\left(\frac{RT_\infty}{M}\right)^{3/2} \Bigg\{ &\left[S_\infty^2 + \frac{\gamma}{\gamma-1} - \frac{\gamma+1}{2(\gamma-1)}\frac{T_w}{T_\infty}\right] \\ &\times \left[\exp\left(-(S_\infty\cos\theta)^2\right) + \sqrt{\pi}\,S_\infty\cos\theta\,(1+\mathrm{erf}(S_\infty\cos\theta))\right] \\ &- \frac{1}{2}\exp\left(-(S_\infty\cos\theta)^2\right)\Bigg\},\end{aligned} \tag{2.3}$$

where $\bar{\alpha}$ is the thermal accommodation coefficient, ρ_∞ the free-stream atmospheric density, T_∞ the neutral temperature, M the mean molar mass of the gas, γ the specific heat ratio, T_w the surface temperature, and θ the incidence angle between the inward surface normal and the incoming flow direction. The free-stream speed ratio S_∞ is

$$S_\infty = \frac{U}{C_{mp}}, \qquad C_{mp} = \sqrt{\frac{2RT_\infty}{M}}, \tag{2.4}$$

with U the spacecraft velocity relative to the neutral atmosphere and R the universal gas constant.

Following Llop [3], the atmosphere is assumed to co-rotate with the Earth, so the local co-rotation speed is

$$V_{\mathrm{corot}} = \omega_\oplus\, r_{\mathrm{orbit}}\cos\phi, \tag{2.5}$$

where $\omega_\oplus$ is the Earth angular velocity, $r_{\mathrm{orbit}} = R_E + h$ the orbital radius, and ϕ the spacecraft latitude. The latitude at each orbit point is taken directly from the atmospheric scenario determined in Section 2.1. To bound the uncertainty associated with the spacecraft flight direction relative to the co-rotating atmosphere, hot- and cold-case relative velocities are defined as

$$U_{\mathrm{hot}} = U_{\mathrm{orb}} + V_{\mathrm{corot}}, \qquad U_{\mathrm{cold}} = U_{\mathrm{orb}} - V_{\mathrm{corot}}, \tag{2.6}$$

with U_{orb} the spacecraft orbital velocity. The hot case is the conservative upper bound on relative velocity, which maximises the free-molecular heating and the cold case is the corresponding lower bound.

Following Moe and Moe [59], the thermal accommodation coefficient $\bar{\alpha}$ quantifies the degree to which incident gas molecules exchange energy with the surface during collision. It is defined as

$$\bar{\alpha} = \frac{E_i - E_r}{E_i - E_w}, \tag{2.7}$$

where E_i and E_r are the energy fluxes of the incident and reflected molecules, and E_w is the energy flux that would result if all impinging molecules were reflected in complete thermal equilibrium with the wall at temperature T_w. The coefficient varies between 0 for purely specular reflection (no energy exchange) and 1 for fully diffuse reflection in thermal equilibrium with the surface. A fully accommodating surface is assumed throughout this work,

$$\bar{\alpha} = 1, \tag{2.8}$$

which is a conservative assumption for thermal design in the free-molecular regime.

Two further simplifications are introduced to reduce the model complexity without affecting its angular dependence. First, for orbital flight the ratio T_w/T_0 is of order 10^{-2}, so the corresponding term in Eq. (2.3) contributes only marginally and is neglected. Second, the gas is treated as diatomic with constant thermodynamic properties so that γ can be assumed constant. Under these assumptions, the simplified free-molecular heat-flux expression used in the remainder of the work reads

$$\begin{aligned} \dot{q}_{\mathrm{FM}} = \frac{\bar{\alpha}\,\rho_\infty}{\sqrt{2\pi}} \left(\frac{RT_\infty}{M}\right)^{3/2} \Bigg\{ & \left[S_\infty^2 + \frac{\gamma}{\gamma-1}\right] \\ & \times \left[\exp\left(-(S_\infty \cos\theta)^2\right) + \sqrt{\pi}\, S_\infty \cos\theta\,(1 + \mathrm{erf}(S_\infty \cos\theta))\right] \\ & - \frac{1}{2} \exp\left(-(S_\infty \cos\theta)^2\right) \Bigg\}. \end{aligned} \tag{2.9}$$

Face-wise Heat-Flux Evaluation

To evaluate aerothermal heating on individual spacecraft surfaces, Eq. (2.9) is applied on a face-by-face basis. Each external surface is treated as a flat plate characterised by a constant inward unit normal vector and a surface area.

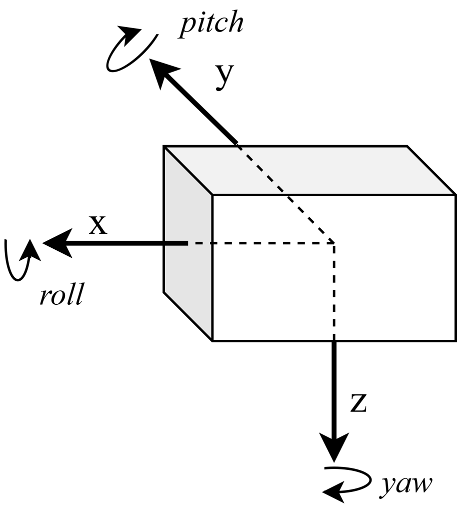

Figure 2.6: Body-fixed reference frame used in the aerothermal and aerodynamic analysis: the x-axis is aligned with the ram direction, the z-axis with nadir, and the y-axis completes a right-handed frame, with the origin at the spacecraft's geometric centre. Rotations about y and z correspond to the angle of attack α and yaw angle β used in Eq. (2.10); roll about x is shown for completeness but is not modelled.

The spacecraft body-fixed reference frame is defined with the x-axis aligned with the nominal spacecraft velocity vector and the z-axis pointing nadir. The y-axis is then chosen to complete a right-handed orthogonal coordinate frame.

The unit vector of the spacecraft velocity relative to the ambient atmosphere, expressed in the body-fixed frame and parameterised by the signed angle of attack α and the signed yaw angle β, is

$$\hat{\mathbf{u}}(\alpha,\beta) = \begin{bmatrix} \cos\alpha\cos\beta \\ \sin\beta \\ \sin\alpha\cos\beta \end{bmatrix}, \qquad \alpha,\beta \in (-\pi,\pi), \tag{2.10}$$

so that $\alpha = \beta = 0$ corresponds to nominal flight with the flow along $+x$. The spacecraft is modelled as a rectangular prism with six external faces, complemented by deployed solar panels treated as flat plates aligned with the spacecraft top surface. The inward unit normals of the six body faces are

$$\begin{aligned} \hat{\mathbf{n}}_{+x} &= [-1\;0\;0]^T, & \hat{\mathbf{n}}_{-x} &= [1\;0\;0]^T, \\ \hat{\mathbf{n}}_{+y} &= [0\;-1\;0]^T, & \hat{\mathbf{n}}_{-y} &= [0\;1\;0]^T, \\ \hat{\mathbf{n}}_{+z} &= [0\;0\;-1]^T, & \hat{\mathbf{n}}_{-z} &= [0\;0\;1]^T. \end{aligned} \tag{2.11}$$

For each surface i, the incidence angle θ_i between the inward surface normal and the incoming flow is

$$\cos\theta_i(\alpha,\beta) = \hat{\mathbf{n}}_i \cdot \big(-\hat{\mathbf{u}}(\alpha,\beta)\big), \tag{2.12}$$

so that $\cos\theta_i = 1$ corresponds to normal incidence and $\cos\theta_i = 0$ to grazing flow, as shown in Figure 2.7.

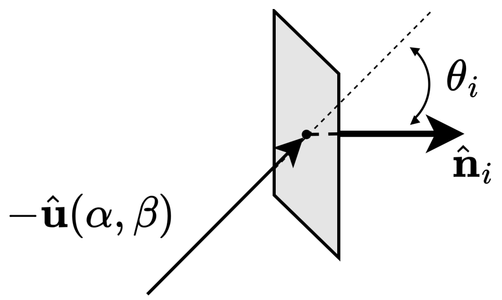

Figure 2.7: Definition of the incidence angle θ_i between the inward surface normal $\hat{\mathbf{n}}_i$ and the incoming flow direction $-\hat{\mathbf{u}}(\alpha,\beta)$, as used in Eq. (2.12).

In the original Caruso formulation [8] the incidence angle is measured from the surface plane rather than from the normal. To use a single angle convention throughout this chapter, the Caruso incidence factor is rewritten as $\sin\theta_{C,i} = \cos\theta_i$, where $\theta_{C,i}$ denotes the Caruso angle and θ_i the angle measured from the inward surface normal.

Surfaces oriented away from the incoming flow are conservatively assumed to receive the same aerothermal heating as surfaces parallel to flow, so negative values of $\cos\theta_i$ are clipped to zero:

$$\cos\theta_{i,\mathrm{eff}} = \max\!\big(0,\; \hat{\mathbf{n}}_i \cdot \big(-\hat{\mathbf{u}}(\alpha,\beta)\big)\big)\,. \tag{2.13}$$

The free-molecular heat flux on surface i is then obtained by evaluating Eq. (2.9) with $\cos\theta_{i,\mathrm{eff}}$:

$$\dot{q}_{\mathrm{FM},i}(\alpha, \beta) = \dot{q}_{\mathrm{FM}}\big(\cos\theta_{i,\mathrm{eff}}(\alpha, \beta)\big)\,, \tag{2.14}$$

with the atmospheric properties ρ_∞, T_∞ and the relative speed U evaluated as functions of altitude. The total aerothermal heat rate into surface i is

$$\dot{Q}_i(\alpha, \beta) = \dot{q}_{\mathrm{FM},i}(\alpha, \beta)\, A_i, \tag{2.15}$$

with A_i the area of the corresponding face or solar panel. This formulation supports the evaluation of aerothermal heating on all spacecraft surfaces as a function of attitude, including fixed offsets and oscillatory motion, while retaining the physical mechanisms responsible for finite heating on surfaces near grazing incidence.

Heat-Flux Trends

Representative heat-flux distributions obtained from the analytical model are shown in Fig. 2.8 as functions of altitude and angle of attack. At a fixed altitude of 130 km, the ram-facing $+x$ surface remains dominant across the considered attitude range, while the lateral $\pm z$ surfaces exhibit a windward switching behaviour with angle of attack. With increasing altitude, the ram-face heat flux decreases strongly, and the hot-case atmosphere produces consistently higher loads than the cold case at every altitude.

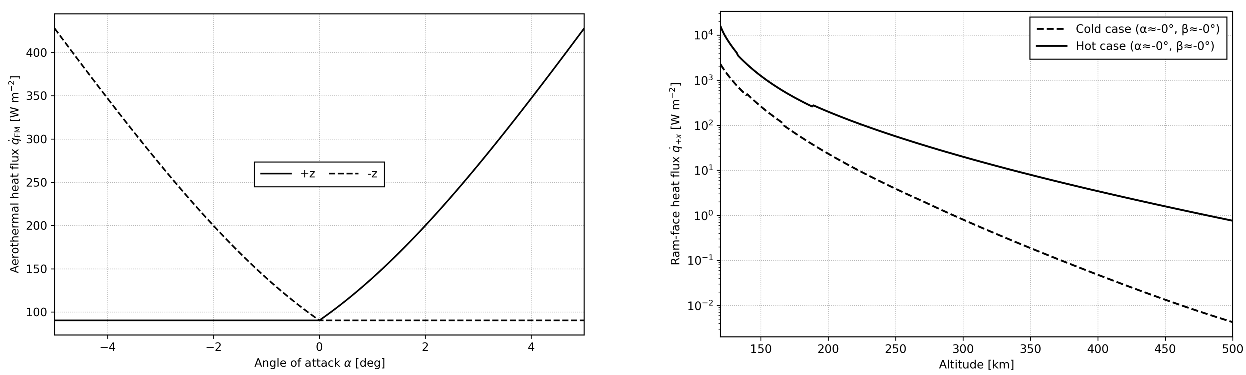

(a) Heat flux versus angle of attack at $h = 130\,\mathrm{km}$.

(b) Ram-face ($+x$) heat flux at $\alpha, \beta = 0°$ versus altitude for the hot and cold cases.

Figure 2.8: Representative free-molecular heat-flux trends from the analytical Caruso model [8]. In (a), the leeward surface plateaus at the same flux as the grazing case ($\alpha = 0°$) rather than decreasing further, reflecting the incidence-clipping assumption of Eq. (2.13) that leeward and grazing surfaces receive equal, non-zero heating. In (b), the ram-facing surface dominates the aerothermal load at low altitude, falling to the order of the solar constant ($\approx 1367\,\mathrm{W\,m^{-2}}$) by approximately 150 km in the hot case.

2.3.2. Comparison with DSMC Reference Data

The reference DSMC results used in this work were generated with the *PICLas* framework [60, 61], a three-dimensional particle-based solver developed at the University of Stuttgart. Within PICLas, intermolecular collisions are handled using the Variable Hard Sphere (VHS) model, and gas-surface interactions are treated through diffuse re-emission at the surface temperature, consistent with the fully accommodating assumption adopted in Section 2.3.1. Mesh refinement is performed by recursive octree subdivision to maintain resolution with respect to the local mean free path [27, 61]. The PICLas simulations were carried out by M. Fertig and M. Ertl (German Aerospace Center, DLR) as part of the project framework (personal communication, 2025). In the present work the outputs are used exclusively as reference values for benchmarking the analytical model.

The analytical heat-flux predictions are compared with the DSMC reference at identical atmospheric conditions, altitudes, and angles of attack, for the ram-facing surface ($+x$), a lateral surface ($+z$), and a side surface ($+y$). The results are shown in Fig. 2.9 for two representative altitudes, h = 180 km and h = 140 km. Both validation altitudes lie within the free-molecular regime identified in Section 2.2. The mission's 130 km perigee sits close to the transitional boundary (Kn = 10 at 131.6 km hot-case, 117.3 km cold-case), so the analytical model is expected to deviate most from DSMC precisely in the regime closest to the lowest operating altitude, beyond what is captured by the present comparison.

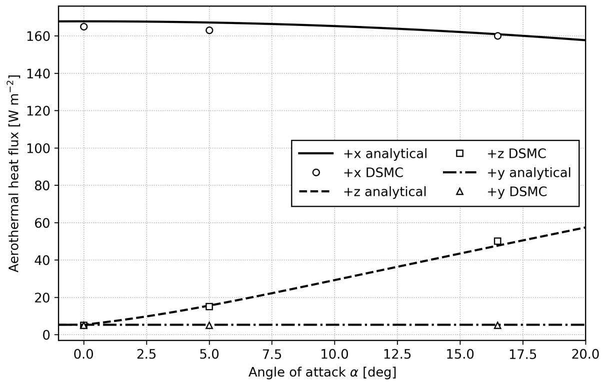

(a) h = 180 km.

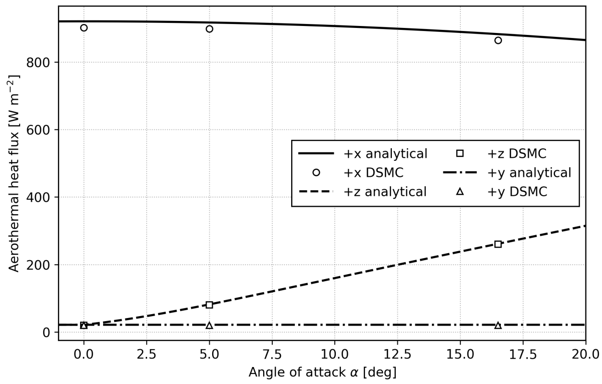

(b) h = 140 km.

Figure 2.9: Analytical Caruso heat-flux predictions [8] (solid lines) agree closely with PICLas DSMC reference data (markers, M. Fertig and M. Ertl, DLR, personal communication, 2025) as a function of angle of attack, for the ram-facing ($+x$), lateral ($+z$), and side ($+y$) surfaces at h = 180 km (a) and h = 140 km (b). Relative error stays below 2.5% on the $+x$ surface and within 6-7% on the $+z$ and $+y$ surfaces across both altitudes.

For the ram-facing $+x$ surface, the analytical model reproduces the DSMC reference closely across all considered conditions. The relative error stays below 2.5% at both altitudes, with the analytical model slightly overpredicting the DSMC value in most cases. The largest absolute deviation occurs at the highest-flux case (h = 140 km, $\alpha = 0°$, $\dot{q} \sim 920\ \mathrm{W\,m^{-2}}$) and is of order 19 $\mathrm{W\,m^{-2}}$.

For the lateral $+z$ and $+y$ surfaces, which spend most of the pitch range near grazing incidence, the relative error is larger. At $\alpha = 0°$ it reaches approximately 6 to 7% at both altitudes. For the $+z$ surface, the deviation decreases below approximately 1% as the angle of attack rotates the surface toward the flow. The $+y$ surface, which remains in persistent near-grazing orientation under pitch variations, retains an essentially constant offset of approximately 6 to 7% across the entire pitch range, indicating a systematic offset between the two formulations for surfaces in this orientation. The absolute deviation on both lateral surfaces stays below approximately 2.4 $\mathrm{W\,m^{-2}}$, one order of magnitude below the ram-face deviations.

To place these results in a broader context, the nominal ram-face heat flux at $\alpha = 0°$ is compared with values reported in the open literature for similar VLEO conditions. Figure 2.10 collects published free-molecular heat-flux estimates from Huang et al. [2], Kaiser [30], Hsieh et al. [29], Pikus et al. [18], Fujita and Noda [62], and Sethuraman et al. [26] across the altitude range 120 to 200 km, together with the hot- and cold-case predictions of the present model. The sources span a range atmospheric assumptions, so direct numerical agreement is not expected. The comparison serves to verify that the present predictions are consistent in order of magnitude and altitude scaling with the peer-reviewed corpus.

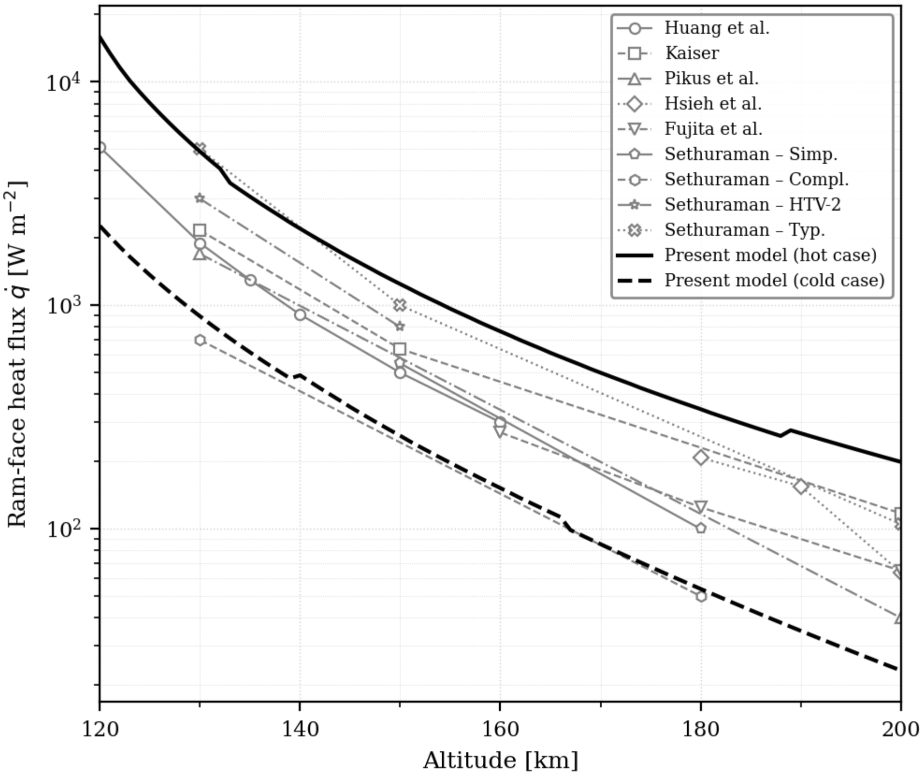

Figure 2.10: Ram-face free-molecular heat-flux predictions from the present analytical model (hot and cold cases, solid and dashed black lines) compared with published estimates from the literature [2, 18, 26, 29, 30, 62] as a function of altitude at nominal attitude ($\alpha = 0°$). Sources span a range of spacecraft geometries and atmospheric assumptions; exact numerical agreement is not expected, and the comparison instead targets consistency in magnitude and altitude scaling. The model envelope brackets all but one of the nine literature cases, with a single point falling modestly below the cold-case bound.

Across the full altitude range considered (120 to 200 km), the hot- and cold-case envelope of the present model (solid and dashed black lines in Fig. 2.10) brackets nearly all of the collected literature estimates. Huang et al. [2] provide the most extensive single dataset, decreasing from 5100 W m^{-2} at h = 120 km to 300 W m^{-2} at h = 160 km; these values remain within the envelope over the entire span and track its altitude scaling. At h = 130 km, where the largest number of independent cases coincide (six of the nine plotted), the values of Huang et al. (1900 W m^{-2}), Pikus et al. [18] (1700 W m^{-2}, DSMC, 3U CubeSat), Kaiser [30] (2167 W m^{-2}, PICLas, SOURCE configuration), and the 'HTV-2' and 'Typ.' cases of Sethuraman et al. [26] (3000 and 5000 W m^{-2}) all fall within the predicted bounds. At higher altitudes, the SLATS estimates of Fujita and Noda [62] (270, 125, 65 W m^{-2} at h = 160, 180, 200 km) and the values of Hsieh et al. [29] (208, 155, 64 W m^{-2} at h = 180, 190, 200 km) likewise lie between the hot- and cold-case curves. The single exception is the 'Compl.' case of Sethuraman et al. [26], which falls modestly below the cold-case lower bound at both reported altitudes (700 W m^{-2} at h = 130 km and 50 W m^{-2} at h = 180 km); the 'Simp.' case from the same source, by contrast, remains within the envelope. None of the nine cases exceeds the hot-case upper bound. Given the diversity of atmospheric assumptions underlying the literature ensemble, this single exception at the lower bound does not compromise the overall consistency between the present envelope and the published corpus in order of magnitude and altitude scaling.

The analytical model therefore matches the DSMC reference within a few percent on the surfaces that carry the dominant aerothermal load, and within a constant small offset on surfaces in grazing orientation where the absolute fluxes are an order of magnitude smaller. The literature comparison further confirms that the hot-case predictions are conservative and physically consistent with the published corpus. A single evaluation of the analytical model requires on the order of milliseconds, against several days for an equivalent DSMC case [26]. The analytical model is therefore adopted as the aerothermal heat-flux model used in the remainder of the thesis.

2.4. Aerodynamic Stability

In the VLEO regime, aerodynamic forces and moments influence both spacecraft attitude behaviour and the panel-level distribution of the aerothermal load. The attitude envelope used by the transient thermal model in Section 2.3 is therefore set by the combination of the spacecraft's passive aerodynamic stability and the residual disturbances acting on it. This section develops the analytical free-molecular force model used to compute drag, lift, and the pitch and yaw moments, benchmarks it against DSMC reference data, and applies it in a linearised stability analysis. The outputs propagated into the thermal model are the pitch and yaw natural frequencies and a conservative oscillation amplitude.

2.4.1. Analytical Free-Molecular Force Model

In the free-molecular regime established in Section 2.2, the molecular mean free path exceeds the characteristic spacecraft length. Intermolecular collisions can therefore be neglected on the scale of the spacecraft, so that gas-surface interactions dominate the aerodynamic load. Aerodynamic forces therefore arise from momentum exchange between incident atmospheric particles and the spacecraft surfaces. The present analysis adopts the classical Sentman formulation [23] as presented by Sutton [63], which provides closed-form expressions for surface force coefficients as a function of the free stream speed ratio S_∞ and the surface incidence angle θ. Each external surface is modelled as a flat plate with diffuse reflection and full thermal accommodation. The total aerodynamic force and moment on the spacecraft are obtained by summing the contributions of all exposed surfaces.

Reference Frame and Kinematics

The reference frame used for the aerodynamic analysis is the same body-fixed frame introduced in Section 2.3.1, with the x-axis along the nominal velocity direction, the z-axis pointing nadir, rotations about the y-axis corresponding to pitch motion, and rotations about the z-axis corresponding to yaw motion. The attitude perturbation is described by the pitch angle α, corresponding to rotation about the body-fixed y-axis, and the yaw angle β, corresponding to rotation about the body-fixed z-axis or, equivalently, a lateral deviation of the incoming flow direction in the body frame. Both angles enter the aerodynamic model through the relative-flow unit vector defined in Eq. (2.10). The corresponding nominal attitude is shown in Fig. 2.6.

The unit vector of the spacecraft velocity relative to the atmosphere, expressed in the body frame, is given by Eq. (2.10). The unit vectors of the drag and lift directions are

$$\hat{\mathbf{e}}_D = -\hat{\mathbf{u}}, \qquad \hat{\mathbf{e}}_L = \frac{\hat{\mathbf{y}} \times \hat{\mathbf{u}}}{||\hat{\mathbf{y}} \times \hat{\mathbf{u}}||}, \tag{2.16}$$

so that positive lift acts in the positive z-direction at $\alpha, \beta = 0$. It should be stated that this is purely a sign convention fixed by Eq. (2.16) and carries no physical meaning beyond labelling the force component orthogonal to $\hat{\mathbf{e}}_D$.

Sentman Free-Molecular Force Formulation

The Sentman free-molecular force formulation [23], presented in the notation of Moe and Moe [59] and following the compact representation of Sutton [63], gives the incremental drag and lift coefficients for surface element i as

$$C_{D,i}A_{\text{ref}} = A_i \left[\frac{P_i}{\sqrt{\pi}} + \gamma_i Q Z_i + \frac{\gamma_i v_r}{2U} \left(\gamma_i \sqrt{\pi} Z_i + P_i \right) \right], \tag{2.17}$$

$$C_{L,i}A_{\mathrm{ref}} = A_i \left[l_i G Z_i + \frac{l_i v_r}{2U} \left(\gamma_i \sqrt{\pi} Z_i + P_i \right) \right], \tag{2.18}$$

where U is the spacecraft velocity relative to the neutral atmosphere. The auxiliary functions are

$$P_i = \frac{e^{-S_\infty^2 \gamma_i^2}}{S_\infty}, \qquad G = \frac{1}{2S_\infty^2}, \qquad Q = 1 + \frac{1}{2S_\infty^2}, \qquad Z_i = 1 + \mathrm{erf}(S_\infty \gamma_i),$$

with $\gamma_i = \cos\theta_i$, $l_i = \sin\theta_i$, and $\mathrm{erf}(x) = (2/\sqrt{\pi}) \int_0^x e^{-y^2}\, \mathrm{d}y$. The ratio between the most probable re-emission speed v_r and the incident bulk speed U is

$$\frac{v_r}{U} = \sqrt{\frac{2}{3}\left[1 + \bar{\alpha}\left(\frac{3RT_w}{U^2} - 1\right)\right]}, \tag{2.19}$$

where T_w is the wall temperature and $\bar{\alpha}$ the energy accommodation coefficient. Consistent with the gas-surface interaction assumption adopted in Section 2.3, the accommodation coefficient is set to

$$\bar{\alpha} = 1. \tag{2.20}$$

This value is required for the closed-form Sentman formulation used here, which is valid only for completely diffuse reflection [23, 64].

Total Force Coefficients The total drag coefficient is the scalar sum of the surface contributions [63],

$$C_D = \sum_i C_{D,i}, \tag{2.21}$$

and the corresponding drag force is

$$\mathbf{F}_D = q\, A_{\mathrm{ref}}\, C_D\, \hat{\mathbf{e}}_D, \tag{2.22}$$

with $q = \frac{1}{2}\rho U^2$ the dynamic pressure. Lift contributions are not collinear and must be summed vectorially,

$$C_L \hat{\mathbf{e}}_L = \sum_i C_{L,i}\, \hat{\mathbf{e}}_{L,i}, \tag{2.23}$$

with the lift direction of surface element i defined as

$$\hat{\mathbf{e}}_{L,i} = \frac{(\hat{\mathbf{e}}_D \times \hat{\mathbf{n}}_i) \times \hat{\mathbf{e}}_D}{\|(\hat{\mathbf{e}}_D \times \hat{\mathbf{n}}_i) \times \hat{\mathbf{e}}_D\|}. \tag{2.24}$$

Flat-Plate Coefficient Behaviour The drag and lift coefficients predicted by the Sentman formulation for a single flat plate are shown in Fig. 2.11 as functions of the incidence angle θ, evaluated for $\bar{\alpha} = 1$, $T_w = 300\,\text{K}$, and $h = 250\,\text{km}$. The left panel shows the range $0° \le \theta \le 90°$, corresponding to direct molecular impingement, and the right panel the extended range $90° \le \theta \le 105°$ near and beyond grazing incidence. Drag dominates across the investigated range: at $\theta \approx 90°$, $C_D \approx 0.072$ exceeds $C_L \approx 0.011$ by a factor of approximately 6.5, and the gap widens further at lower incidence angles, where $C_L \to 0$ as $\theta \to 0°$ by the symmetry of normal incidence on a flat plate. Finite drag and lift values persist for $\theta > 90°$, reflecting the thermal velocity distribution of the impinging molecules [3]. These contributions are small in magnitude but are retained in the formulation as they may become non-negligible for slender geometries or large moment arms. The qualitative behaviour of both coefficients is consistent with the DSMC-computed flat-plate force coefficients reported by Fujita and Noda [62] at 160 and 300 km for accommodation coefficients spanning 0.7 to 1.0: at normal incidence and $\bar{\alpha} = 1$ their data yield $C^N \approx$ 2.0-2.1, the tangential coefficient C^T peaks near $\theta \approx 45°$, and the altitude sensitivity between 160 and 300 km is small, all consistent with the present predictions.

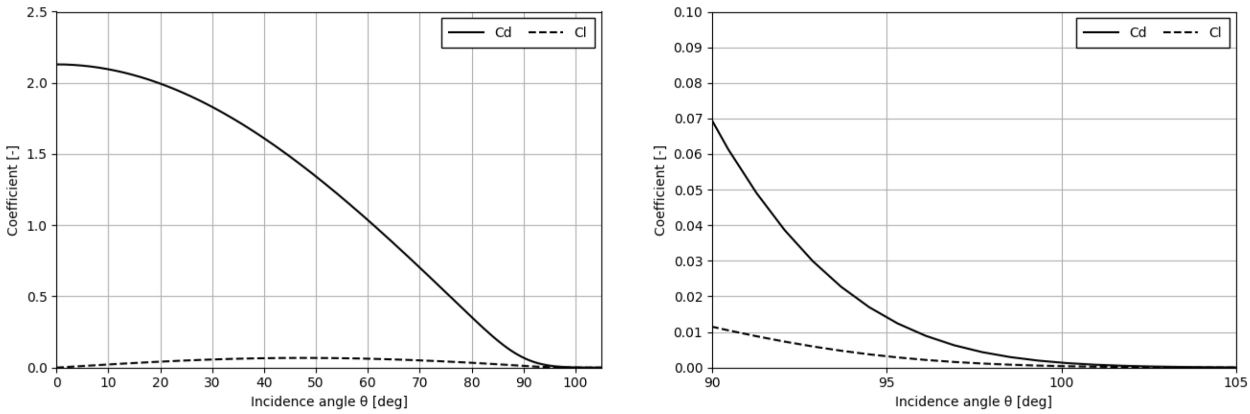

(a) $0° \le \theta \le 90°$.

(b) $90° \le \theta \le 105°$.

Figure 2.11: Flat-plate drag and lift coefficients predicted by the Sentman free-molecular model [23] for $\bar{\alpha} = 1$ and $T_w = 300\,\text{K}$ at $h = 250$ km. Drag dominates across the investigated incidence range, exceeding lift by a factor of approximately 6.5 near grazing incidence ($\theta = 90°$). Finite values persist beyond grazing ($\theta > 90°$) due to the thermal velocity spread of impinging molecules.

2.4.2. Comparison with DSMC Reference Data

The analytical Sentman model is compared with DSMC results obtained from the PICLas framework introduced in Section 2.3.2, at identical atmospheric conditions, altitudes, and angles of attack (M. Fertig and M. Ertl, DLR, personal communication, 2025). The comparison is shown in Fig. 2.12 for two representative altitudes, $h = 180$ km and $h = 140$ km.

For the drag force F_D, the analytical model consistently underpredicts the DSMC reference across all conditions. At $h = 180$ km the relative error is between approximately 7.7% and 12.7%, corresponding to absolute differences in the range 2.8×10^{-4} to 6.5×10^{-4} N. At $h = 140$ km the relative error is between approximately 8.9% and 15.7%, with absolute differences up to 3.8×10^{-3} N. The maximum relative error is reached at $\alpha = 5°$, $h = 140$ km. The bias is consistent in sign across all evaluated cases, and the analytical model preserves the trend of F_D with both angle of attack and altitude.

For the lift force F_L, the absolute differences are between 3.5×10^{-5} and 1.3×10^{-3} N across all evaluated cases. At non-zero angles of attack the relative error stays between approximately 5.5% and 13.1%, with the analytical model underpredicting the DSMC value. At $\alpha = 0°$, F_L vanishes by the mirror symmetry of the platform about the x-z plane, so both the DSMC ($\sim 10^{-6}$ N) and analytical ($\sim 10^{-5}$ to 10^{-4} N) values are residuals rather than converged predictions. The relative error is not reported at this point, as it diverges for a vanishing reference value. The absolute difference, $O(10^{-5}\,\text{N})$, is consistent with the deviations observed at non-zero angles of attack.

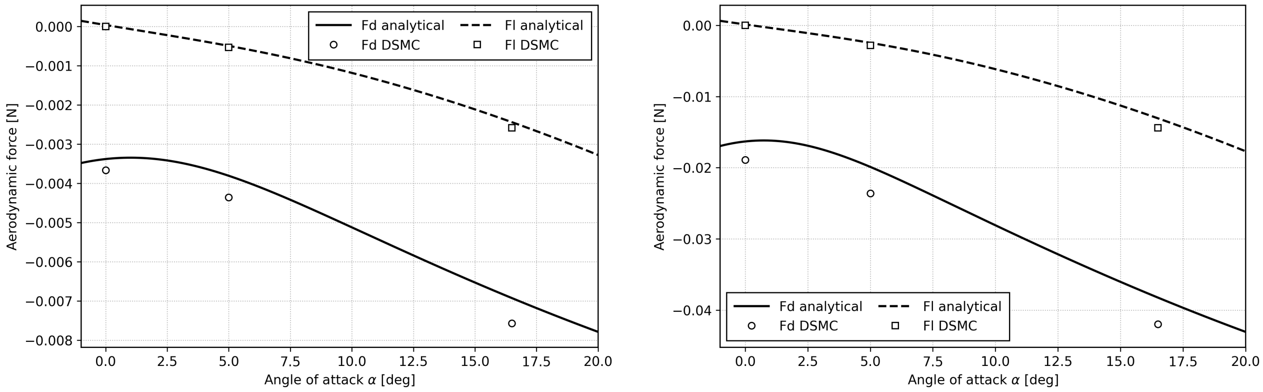

(a) $h = 180\,\mathrm{km}$. **(b)** $h = 140\,\mathrm{km}$.

Figure 2.12: Analytical Sentman drag (F_D, solid) and lift (F_L, dashed) predictions [23, 63] compared with PICLas DSMC reference data (markers; M. Fertig and M. Ertl, DLR, personal communication, 2025) as a function of angle of attack at $h = 180$ km (a) and $h = 140$ km (b). The analytical model systematically underpredicts both F_D and F_L by approximately 6-16% across the evaluated range, while preserving the trend of both quantities with angle of attack and altitude.

The systematic underprediction of F_D is consistent with the limitations of a face-by-face analytical formulation, which does not resolve secondary particle-surface interactions between adjacent panels and therefore omits the additional momentum flux carried by reflected or shadowed molecules. This mechanism is corroborated by Ramesh et al. [64], who compared a TPMC panel-element formulation against DSMC for a VLEO CubeSat-class satellite with deployed solar panels and found a comparable sign-consistent drag underprediction, attributing it to multiple reflections unresolved by face-by-face summation. The pitch and yaw moments $M_y(\alpha)$ and $M_z(\beta)$ used in the linearised stability analysis derive from the same per-panel formulation and were not independently benchmarked against DSMC.

From a thermal perspective, this consistent underprediction of F_D is conservative. A higher true drag force implies a larger restoring aerodynamic moment and hence a higher natural oscillation frequency. As demonstrated in Fig. 2.13, higher yaw-oscillation frequencies yield lower steady-state maximum temperatures, since the panel dwell time at windward incidence decreases. The bias in F_D therefore drives the oscillation frequency used in the thermal model towards lower values, resulting in a more severe and conservative bound on the side-panel temperature. On this basis, the analytical model is retained as the primary tool for the stability analysis, with the F_D and F_L deviations reported above defining the expected error bound on the resulting aerodynamic coefficients.

2.4.3. Linearised Stability and Oscillation Behaviour

The aerodynamic force model provides the nonlinear pitch and yaw moments $M_y(\alpha, \beta)$ and $M_z(\alpha, \beta)$, which are the basis for the stability analysis. The aerodynamic moment vector about the centre of gravity is obtained by summing the contributions of all surface forces,

$$\mathbf{M}(\alpha, \beta) = \sum_i (\mathbf{r}_i - \mathbf{r}_{\mathrm{CG}}) \times \mathbf{F}_i, \tag{2.25}$$

where $\mathbf{r}_i$ is the position vector of the surface centroid and $\mathbf{r}_{\mathrm{CG}}$ the centre-of-gravity location. The pitch and yaw moments used in the stability analysis are the components of $\mathbf{M}$ along the body-fixed y- and z-axes, $M_y = \mathbf{M} \cdot \hat{\mathbf{y}}$ and $M_z = \mathbf{M} \cdot \hat{\mathbf{z}}$.

Linearised Stability

The stability analysis is performed separately for small pitch and yaw perturbations about the nominal equilibrium orientation, defined by $\alpha = 0$ and $\beta = 0$. This separation is consistent with the implementation, where the aerodynamic moment is evaluated once by perturbing the pitch angle α while keeping $\beta = 0$, and once by perturbing the yaw angle β while keeping $\alpha = 0$. The local incidence angle of each surface remains a dependent quantity through $\cos\theta_i(\alpha, \beta)$ and is therefore not used as an independent linearisation variable. By the mirror symmetry of the platform geometry about the x-z plane, the cross-coupling derivatives $\partial M_y/\partial\beta\big|_0$ and $\partial M_z/\partial\alpha|_0$ vanish at the equilibrium orientation, so the pitch and yaw modes are decoupled to leading order. This decoupling relies on the assumption of perfect geometric and mass symmetry, with the solar panels and centre of gravity nominally aligned with the x-z plane. In practice, manufacturing and integration tolerances introduce small asymmetries that break this idealisation. A persistent offset of this kind couples the pitch and yaw moments and biases the equilibrium attitude away from $\alpha = \beta = 0$, which can manifest as a constant roll or a steady non-zero angle of attack rather than a pure oscillation about the nominal orientation. Such offset-driven, quasi-static attitude bias and the corresponding modification of the thermal loading footprint are not captured by the present linearised, symmetric formulation and are left for future analysis.

For the pitch mode, the aerodynamic moment about the body-fixed y-axis is linearised as

$$M_y(\alpha, \beta = 0) \approx \left.\frac{\partial M_y}{\partial\alpha}\right|_{\alpha=0,\beta=0} \alpha. \tag{2.26}$$

For the yaw mode, the aerodynamic moment about the body-fixed z-axis is linearised as

$$M_z(\alpha = 0, \beta) \approx \left.\frac{\partial M_z}{\partial\beta}\right|_{\alpha=0,\beta=0} \beta. \tag{2.27}$$

The corresponding aerodynamic stiffness parameters are defined as

$$k_\alpha = -\frac{1}{q}\left.\frac{\partial M_y}{\partial\alpha}\right|_{\alpha=0,\beta=0}, \qquad k_\beta = -\frac{1}{q}\left.\frac{\partial M_z}{\partial\beta}\right|_{\alpha=0,\beta=0}, \tag{2.28}$$

where $q = \frac{1}{2}\rho U^2$ is the dynamic pressure. Positive values of k_α and k_β correspond to restoring aerodynamic moments and therefore to static stability about the respective axis. The derivatives are evaluated numerically using central finite-difference schemes,

$$\left.\frac{\partial M_y}{\partial\alpha}\right|_{\alpha=0,\beta=0} \approx \frac{M_y(+\Delta\alpha, 0) - M_y(-\Delta\alpha, 0)}{2\,\Delta\alpha}, \tag{2.29}$$

$$\left.\frac{\partial M_z}{\partial\beta}\right|_{\alpha=0,\beta=0} \approx \frac{M_z(0, +\Delta\beta) - M_z(0, -\Delta\beta)}{2\,\Delta\beta}. \tag{2.30}$$

Natural Frequency and Critical Damping

The linearised rotational dynamics are treated as two independent single-axis perturbation problems. For pitch motion about the body-fixed y-axis, the governing equation is

$$I_{yy}\,\ddot{\alpha} + q\,k_\alpha\,\alpha = 0, \tag{2.31}$$

where I_{yy} is the principal moment of inertia about the y-axis. For yaw motion about the body-fixed z-axis, the corresponding equation is

$$I_{zz}\,\ddot{\beta} + q\,k_\beta\,\beta = 0, \tag{2.32}$$

where I_{zz} is the principal moment of inertia about the z-axis. The undamped natural frequencies of the pitch and yaw modes are therefore

$$\omega_{\alpha,0} = \sqrt{\frac{q\, k_\alpha}{I_{yy}}}, \qquad \omega_{\beta,0} = \sqrt{\frac{q\, k_\beta}{I_{zz}}}. \tag{2.33}$$

The corresponding critical damping coefficients are

$$c_{\alpha,\mathrm{crit}} = 2\sqrt{I_{yy}\, q\, k_\alpha}, \qquad c_{\beta,\mathrm{crit}} = 2\sqrt{I_{zz}\, q\, k_\beta}. \tag{2.34}$$

Eqs. 2.31-2.34 define the natural frequencies and critical damping coefficients governing the pitch and yaw dynamics. Evaluated as functions of altitude and atmospheric state, they characterise the spacecraft's aerodynamic stability in the VLEO regime and, through the resulting oscillation frequency, define part of the attitude variation propagated into the thermal model. Before this frequency is evaluated as a function of altitude, its thermal significance must first be established.

Figure 2.13 shows the periodic-steady-state peak nodal temperature of a representative 1 mm thick aluminium side panel (ε = 0.20) under Caruso aerothermal heating only [8], as a function of yaw-oscillation frequency at h_p = 130 km. The static reference (f = 0 Hz, $T_{\max}$ = 25.63 °C) corresponds to fixed leeward incidence, where the panel receives only the residual aerothermal flux from the thermal velocity distribution of the impinging molecules. At low but nonzero frequency, the oscillation period is long relative to the panel's thermal response time, so the panel temperature tracks the slowly varying incidence angle and swings between a hot value near windward dwell and a cold value near leeward dwell; the peak temperature rises by up to 29.3 °C above the static reference as a result. As the frequency increases, the oscillation period becomes short relative to the panel's thermal response time, so the panel's thermal mass increasingly averages the alternating windward and leeward flux, and the peak temperature converges to a frequency-independent asymptote near 40.4 °C above approximately 10^{-2} Hz.

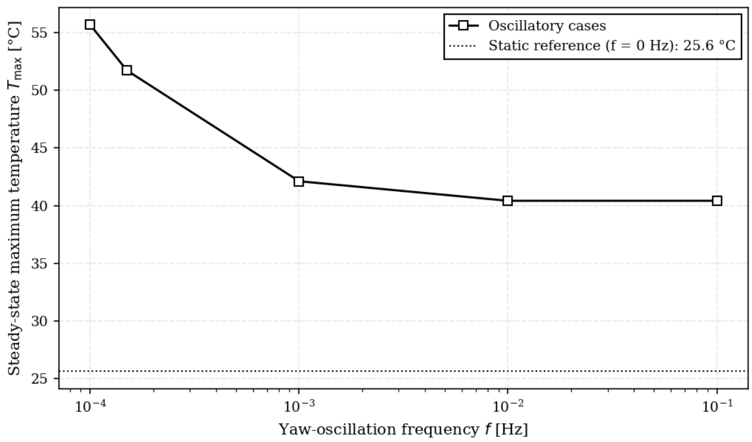

Figure 2.13: Periodic-steady-state peak temperature of a representative 1 mm-thick aluminium side panel (ε = 0.20, modelled as a single lumped-capacitance node) under Caruso aerothermal heating [8], as a function of yaw-oscillation frequency at h_p = 130 km, evaluated over the final 50 % of a 16-orbit simulation to exclude the initial transient. Peak temperature decreases monotonically with frequency, from a maximum at the lowest oscillation frequency to a frequency-independent asymptote of 40.4 °C above approximately 10^{-2} Hz, as higher frequencies reduce panel dwell time at windward incidence.

From this frequency dependence, the hot case is modelled as undamped yaw oscillation about the disturbance angle θ_{dist} at frequency $\omega_{\beta,0}$, since oscillation only ever raises the panel temperature above the static reference (Figure 2.13). The cold case is modelled as non-oscillatory, held at the static attitude with the heat shield ram-facing, the orientation that minimises side-panel heating and therefore bounds the cold case from below.

The yaw-oscillation frequency propagated into the hot-case model follows from the linearised stability analysis (2.33), evaluated for the configuration that maximises the aerodynamic stiffness k_β ((2.28)): the centre of gravity at its most forward allowable location, with the solar panels positioned as far aft as possible. This configuration is chosen to make the platform as aerodynamically stable as possible, which, by Figure 2.13, also minimises the oscillatory side-panel heating.

Evaluated as a function of altitude for this configuration, the resulting pitch and yaw natural frequencies are shown in Figure 2.14.

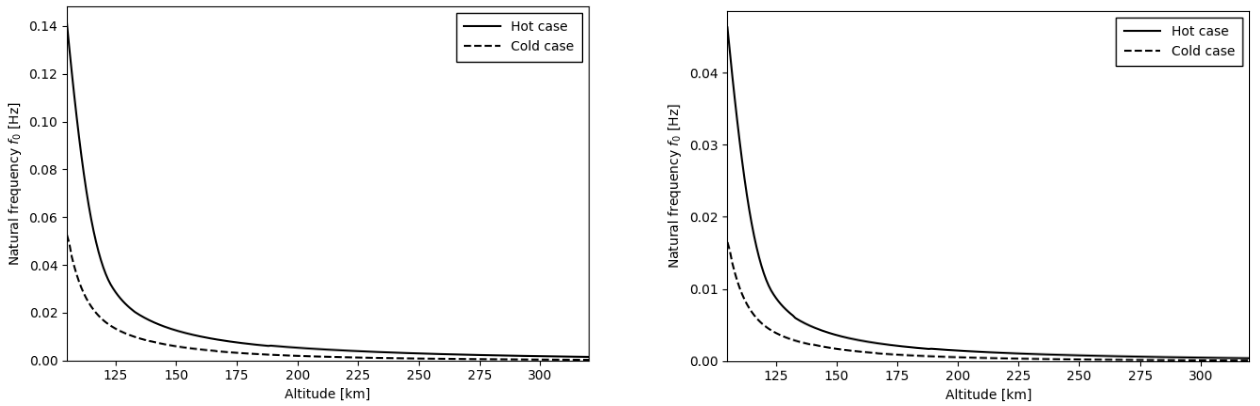

(a) Pitch motion.

(b) Yaw motion.

Figure 2.14: Altitude-dependent natural frequency of the linearised aerodynamic motion in pitch and yaw, evaluated for the centre-of-gravity-forward, solar-panel-aft configuration that maximises the aerodynamic stiffness within the design envelope. Near the minimum-perigee altitude, both natural frequencies remain above the 10^{-2} Hz threshold identified in Fig. 2.13, placing the yaw-driven thermal response in the frequency-converged regime; at higher altitudes within the envelope the natural frequency falls below this threshold.

The oscillation amplitude used in the thermal model is derived from the wind-induced horizontal flow disturbance. The HWM wind model [65] provides only horizontal zonal and meridional wind components, which tilt the apparent flow direction within the horizontal plane and therefore act as a yaw excitation in the body frame, not as a pitch excitation. Vertical winds, which would drive the pitch mode, are not represented in HWM, so the pitch mode is treated as non-oscillatory in the thermal analysis and only the yaw oscillation is propagated. For each altitude, latitude, and longitude point in the hot and cold atmospheric datasets, the yaw disturbance angle is computed as

$$\theta_{\text{dist}} = \tan^{-1}\left(\frac{\sqrt{U_{\text{wind}}^2 + V_{\text{wind}}^2}}{U_{\text{rel}}}\right), \tag{2.35}$$

where U_{wind} and V_{wind} are the zonal and meridional wind components and U_{rel} is the spacecraft velocity relative to the atmosphere. The maximum disturbance angle obtained across the evaluated hot and cold atmospheric datasets is $\theta_{\text{dist,max}} = 1.23°$, which is adopted as a conservative, altitude-independent yaw oscillation amplitude for the transient thermal simulations. Together with the altitude-dependent yaw natural frequency $\omega_{\beta,0}(h)$, this amplitude defines the aerodynamic attitude variation case propagated into the thermal model.

2.5. Earth Albedo and Infrared Radiation

Earth-reflected solar radiation (albedo) and Earth-emitted infrared (IR) radiation are contributors to the external thermal environment of spacecraft in low Earth orbit. Albedo refers to the fraction of incident solar radiation reflected by the Earth-atmosphere system, while Earth IR denotes the thermal radiation emitted by the surface and atmosphere as a result of their finite temperature. Both act as external heat inputs to spacecraft surfaces and are commonly treated as boundary conditions in thermal analysis [56].

In early-phase thermal design these fluxes are routinely represented using constant, globally averaged values. Such representations are derived from long-term observations and describe mean conditions rather than short-term or local extremes. ECSS recognises that spatial and temporal variations in Earth radiative properties may become relevant for specific mission classes, in particular polar or near-polar orbits and components with low thermal inertia, in which case a uniform representation does not adequately capture the distribution of external heat loads over an orbit [56].

This section investigates three Earth radiation models of increasing fidelity, namely a constant model, the ECSS long-term averaging method, and the CERES-derived zonal harmonics model of Sasaki [38, 56]. The three models are first compared against each other to characterise their structural differences, then quantitatively benchmarked against CERES observations, and finally evaluated through a simplified thermal test case. The outcome is the selection of the Earth radiation model and the definition of the hot- and cold-case Earth radiation conditions used in the spacecraft-level thermal analysis.

2.5.1. Constant Earth Radiation Model

The constant Earth radiation model is the simplest representation of the planetary radiative environment. Earth albedo and IR are assumed to be spatially and temporally uniform, and are described using fixed, globally averaged values. Following ECSS recommendations, the average Earth albedo is taken as $a_a = 0.3$ and the average Earth IR as 230 $\mathrm{W\,m^{-2}}$ [56]. ECSS further notes that for an orbiting spacecraft the albedo can vary between 0.05 over open ocean and 0.6 over regions with high cloud cover or polar ice, while Earth IR can range from about 150 to 350 $\mathrm{W\,m^{-2}}$, with diurnal IR variations of up to approximately 20% over desert regions [56].

By construction, the constant model cannot capture latitudinal gradients, seasonal asymmetries, or short-term fluctuations. It is nevertheless retained as a baseline in this work for two reasons. First, its simplicity makes it well-suited for benchmarking the higher-fidelity models. Second, it remains widely adopted in preliminary thermal design and represents the most common starting point against which an upgrade to more detailed Earth radiation models must be justified [38].

2.5.2. ECSS Long-Term Averaging Method

To address the limitations of the uniform model while retaining analytical simplicity, ECSS recommends a long-term averaging method when spatial and seasonal variations are expected to be relevant [56]. The method exploits the observation that the longitudinal dependence of the Earth radiative properties is comparatively small and is largely averaged out by Earth rotation, while the latitudinal and seasonal dependences are not. The Earth-atmosphere system is therefore described through latitude- and time-dependent albedo and infrared emissivity functions expanded in low-order Legendre polynomials, with an annual cosine term representing the seasonal cycle. The formulation follows Knocke et al. [36].

The solar albedo $\rho(\varphi, t)$ and infrared emissivity $\varepsilon(\varphi, t)$ are given by

$$\rho(\varphi, t) = \rho_0 + \rho_1 \cos\big[\omega(\mathrm{JD} - t_0)\big]\, P_1(\sin\varphi) + \rho_2\, P_2(\sin\varphi), \tag{2.36}$$

$$\varepsilon(\varphi, t) = \varepsilon_0 + \varepsilon_1 \cos\big[\omega(\mathrm{JD} - t_0)\big]\, P_1(\sin\varphi) + \varepsilon_2\, P_2(\sin\varphi), \tag{2.37}$$

where φ is the geocentric latitude, JD is the Julian date, t_0 corresponds to 22 December 1981, and $\omega = 2\pi/365.25\ \text{day}^{-1}$. The functions $P_1(x) = x$ and $P_2(x) = \frac{1}{2}(3x^2 - 1)$ are the first- and second-order Legendre polynomials. The coefficients are taken from ECSS as $\rho_0 = 0.34$, $\rho_1 = 0.10$, $\rho_2 = 0.29$ for albedo and $\varepsilon_0 = 0.68$, $\varepsilon_1 = -0.07$, $\varepsilon_2 = -0.18$ for infrared emissivity [56]. Equivalent coefficient sets have been reported in recent applications of the method to spacecraft thermal uncertainty analyses [66].

The method represents time-averaged behaviour, short-term or localised extremes are not resolved. It does, however, capture the dominant latitudinal and seasonal trends, and is computationally inexpensive compared to data-driven alternatives.

2.5.3. Sasaki Zonal Harmonics Model

As a higher-fidelity, data-driven alternative, Sasaki proposed a representation of Earth albedo and outgoing longwave radiation (OLR) obtained by fitting long-term observations from the Clouds and the Earth's Radiant Energy System (CERES) project [38]. The objective is to capture the dominant spatial and seasonal variability while retaining a computationally efficient analytical form.

CERES Reference Data

The CERES project produces top-of-atmosphere (TOA) flux estimates from scanning radiometers operating in a shortwave channel (0.3 to 5 µm), a total channel (0.3 to 200 µm), and either a window or longwave channel covering the thermal infrared range [67]. CERES instruments have flown on several platforms, including TRMM, Terra, Aqua, Suomi-NPP, and NOAA-20, mostly in Sun-synchronous orbits to ensure global coverage [68, 69]. Raw radiances are converted into TOA fluxes through Angular Distribution Models that account for viewing geometry and surface-type-dependent anisotropy [70, 71]. The SYN1deg products combine observations from multiple platforms, including geostationary satellites, to provide temporally and spatially complete global datasets [72, 73].

In this work, Sasaki's fit is based on the daily CERES SYN1deg-Day Edition 4.1 product, which provides OLR and albedo coefficient on a $1° \times 1°$ grid over more than two decades [74]. These data show a strong equator-to-pole gradient with additional seasonal modulation: OLR is larger in equatorial regions and decreases toward the poles, while the albedo coefficient shows the opposite trend. During polar night, no shortwave radiation is incident and the CERES product does not assign an albedo value.

Model Formulation

The OLR and albedo distributions are expanded in zonal spherical harmonics (order $m = 0$ only), under the same rotation-averaging argument used by ECSS. Seasonal variability is added through a first-order annual harmonic. The OLR is

$$e(\varphi, \mathrm{DOY}) = \sum_{\ell=0}^{L} \left[E_{\ell,0} + E_{\ell,1}\cos x + E_{\ell,2}\sin x\right] Y_{\ell 0}(\varphi), \tag{2.38}$$

and the albedo coefficient is

$$a(\varphi, \mathrm{DOY}) = \sum_{\ell=0}^{L} \left[A_{\ell,0} + A_{\ell,1}\cos x + A_{\ell,2}\sin x\right] Y_{\ell 0}(\varphi), \tag{2.39}$$

with annual phase $x = (2\pi/365)(\mathrm{DOY} - 1)$ and zonal harmonics

$$Y_{\ell 0}(\varphi) = \sqrt{\frac{2\ell + 1}{4\pi}}\, P_\ell\big(\sin\varphi\big)\,. \tag{2.40}$$

Model Order Selection

Sasaki evaluated expansions up to $L = 8$ and found that increasing L beyond approximately 4 yields diminishing improvements in RMSE for both OLR and albedo, with the sensitivity to L being most pronounced for high-inclination orbits [38]. The present work therefore adopts the fourth-degree ($L = 4$), first-order time-variant zonal harmonics model. This choice is consistent with the near-polar orbit ($i \approx 90°$) considered here, which traverses high-latitude regions every revolution and is therefore most sensitive to the latitudinal structure that $L = 4$ resolves. The coefficients employed are reproduced from Tables 2 and 3 of [38] and are listed in Tables 2.2a and 2.2b for reproducibility.

Table 2.2: Sasaki zonal harmonics coefficients for OLR and albedo ($L = 4$). [38]

(a) OLR coefficients.

ℓ	$E_{\ell,0}$	$E_{\ell,1}$	$E_{\ell,2}$
0	846.5127	-11.71082	-5.007011
1	9.35777	-38.16887	-11.64853
2	-84.64440	-4.148072	-1.399822
3	12.61828	-24.18759	-6.167508
4	-22.10449	6.610096	2.565666

(b) Albedo coefficients.

ℓ	$A_{\ell,0}$	$A_{\ell,1}$	$A_{\ell,2}$
0	1.135866	0.02805531	0.004863929
1	-0.02568558	0.1560044	0.02306313
2	0.4129173	0.03064179	0.01157133
3	-0.04970988	0.03961453	0.02163336
4	0.1107932	0.002905725	-0.01802534

Like the ECSS method, the Sasaki model represents long-term average Earth radiation behaviour, and does not resolve short-term or highly localised extremes. It does, however, capture a stronger equator-to-pole gradient and more realistic seasonal asymmetries than the ECSS formulation, which is the basis for the comparison developed in the following subsection.

2.5.4. Earth Radiation Model Comparison

The model selection is based on three comparisons: the latitudinal structure of the predicted fluxes, the area-weighted RMSE against CERES observations, and the resulting temperature response of a simplified thermal test case.

Yearly-Mean Latitudinal Profiles

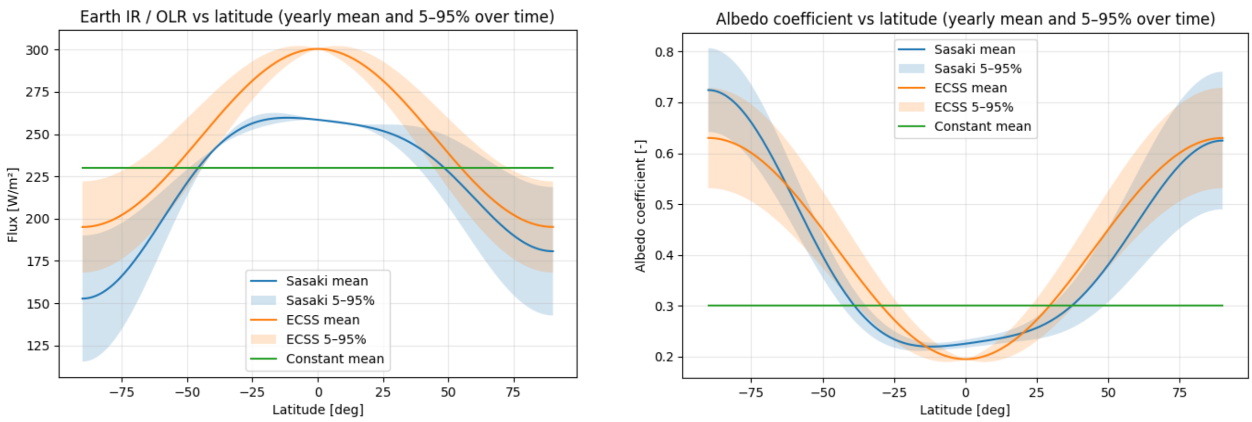

(a) Outgoing longwave radiation (OLR).

(b) Albedo coefficient.

Figure 2.15: Yearly-mean latitudinal profiles of OLR and albedo coefficient predicted by the constant model [56], the ECSS long-term averaging method [36, 56], and the Sasaki zonal harmonics model [38]. Solid lines denote yearly means; shaded regions indicate the 5 to 95% temporal percentile range over one year for the ECSS and Sasaki cases. The gap between ECSS and Sasaki is smaller than the gap to the constant model overall, but is largest in the polar regions traversed each orbit by a near-polar mission.

Figure 2.15 shows the yearly-mean latitudinal profiles of OLR and albedo predicted by the three models, together with their 5 to 95% temporal envelopes for the ECSS and Sasaki cases. The three models exhibit structurally different behaviour. The constant model assumes spatially invariant fluxes (230 $\mathrm{W\,m^{-2}}$ and $a = 0.3$) and therefore reproduces neither the latitudinal gradient nor any seasonal variation. The ECSS method introduces a symmetric equator-to-pole variation, with peak OLR near the equator and minimum albedo there, together with a moderate seasonal spread that grows with latitude. The Sasaki model also follows the expected equator-to-pole gradient, but with a smaller equatorial amplitude in OLR and a stronger seasonal asymmetry at high latitudes, reflecting structure embedded in the CERES-derived harmonic coefficients. The differences between ECSS and Sasaki are smaller than the gap to the constant model, but remain non-negligible in polar regions, which is precisely the part of the latitudinal range traversed every orbit by a near-polar mission.

Quantitative Comparison with CERES

To quantify the deviation of each analytical model from the observational reference, area-weighted RMSE values are computed against the zonal-mean CERES distribution [67, 68] for a representative set of equinox and solstice dates in 2022. The RMSE uses cosine-latitude weighting,

$$\mathrm{RMSE} = \sqrt{\frac{\sum_i w_i \left(f_{\mathrm{model},i} - f_{\mathrm{CERES},i}\right)^2}{\sum_i w_i}}, \qquad w_i = \cos\varphi_i, \tag{2.41}$$

so that each latitude band contributes in proportion to its physical surface area. Albedo errors are computed only where CERES provides a valid albedo coefficient, since shortwave radiation is physically absent during polar night and no reflectivity can be defined there [38]. OLR is defined regardless of solar illumination and no masking is required.

Across the four reference dates (vernal and autumnal equinoxes and the two solstices of 2022), the Sasaki model consistently produces the smallest RMSE for both quantities. For the albedo coefficient, the Sasaki RMSE remains at or below approximately 0.03, against approximately 0.05 to 0.07 for ECSS and in excess of 0.10 for the constant model. For OLR, the Sasaki RMSE lies between approximately 12 and 16 $\mathrm{W\,m^{-2}}$, compared to between 30 and 34 $\mathrm{W\,m^{-2}}$ for both ECSS and the constant representation. The signed percent error against CERES, evaluated at the same dates, confirms that the constant model exhibits monotonic latitude-dependent bias of order ±50% near the poles; that ECSS retains a structured seasonal bias near the solstices, overpredicting equatorial OLR and underpredicting mid-latitude albedo asymmetries; and that the Sasaki model has no large-scale systematic bias, with residuals within approximately ±20% for albedo and approximately ±10 to ±15% for OLR over most latitudes. The resulting plots of these analysis can be found in Appendix B.

Impact on Surface Temperature

The incident Earth radiation flux is ultimately consumed by the spacecraft thermal model, where it interacts with optical properties, thermal mass, and conductive coupling. Since these characteristics are component-dependent, no single thermal configuration is universally representative. To isolate the influence of the Earth radiation model itself, a simplified thermal test case is used: a 1×1 m aluminium plate of 1 mm thickness ($\rho = 2700\ \mathrm{kg\,m^{-3}}$, $c_p = 900\ \mathrm{J\,kg^{-1}\,K^{-1}}$, $k = 201\ \mathrm{W\,m^{-1}\,K^{-1}}$), nadir-pointing, with front-face optical properties $\alpha = \varepsilon = 0.9$ and an adiabatic rear face. The configuration is illustrated in Fig. 2.16.

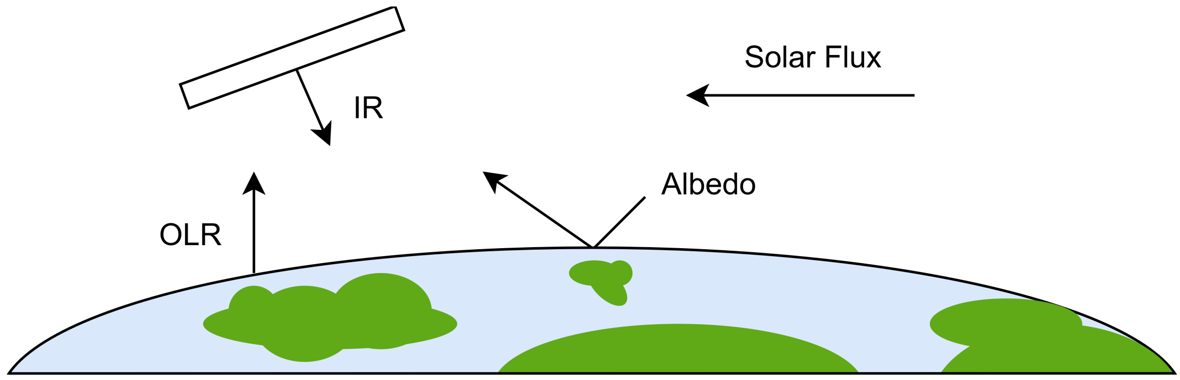

Figure 2.16: Simplified thermal test configuration: a nadir-pointing flat plate exposed to direct solar radiation, OLR, and albedo flux, radiating to space and with an adiabatic rear face [38].

Transient simulations are performed on circular polar orbits ($i = 90°$) for altitudes between 100 and 500 km, with aerothermal heating switched off so that only the Earth radiation contribution differs between the four runs. Three reference dates are evaluated to sample seasonal variability: the winter solstice (21 December 2022), the summer solstice (21 June 2022), and an equinox case (6 March 2022). The main seasonal trends are summarised at 100 km altitude in Figs. 2.17 and 2.18. The complete altitude-dependent time-evolution and deviation curves are provided in Appendix B, with the temperature histories shown in Figs. B.9 to B.11 and the corresponding deviations from the CERES-driven reference case shown in Figs. B.12 to B.14. The inter-model ΔT spread does not decrease appreciably across this altitude range, since neither the OLR nor the albedo flux varies strongly with spacecraft altitude, whereas the aerothermal contribution decreases by several orders of magnitude over the same range (Fig. 2.8). Earth-radiation-model accuracy is therefore most consequential precisely where aerothermal heating is negligible.

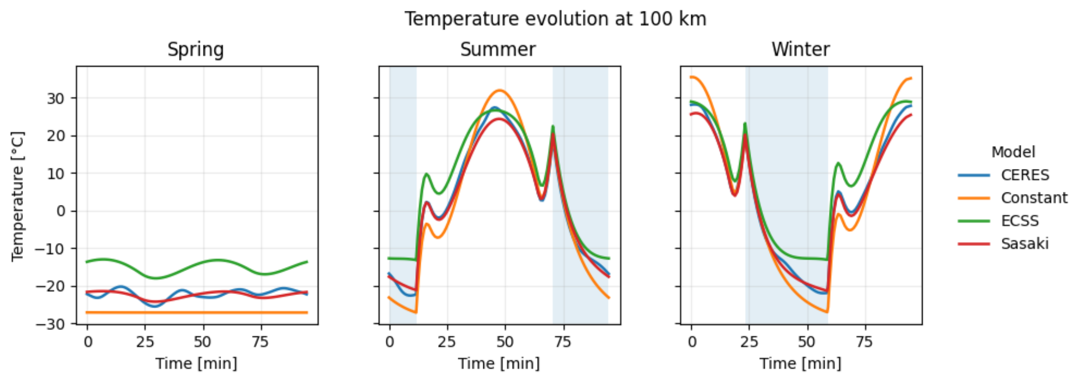

Figure 2.17: Aluminum plate temperature evolution over one orbit at the 100 km lower bound of the test sweep, for circular polar test orbits at each of the selected seasonal cases. The three panels compare the response obtained with the CERES-driven reference case [67, 68] and the simplified Earth radiation models (constant, ECSS [36, 56], and Sasaki [38]) for 6 March, 21 June, and 21 December 2022. Shaded regions indicate eclipse intervals.

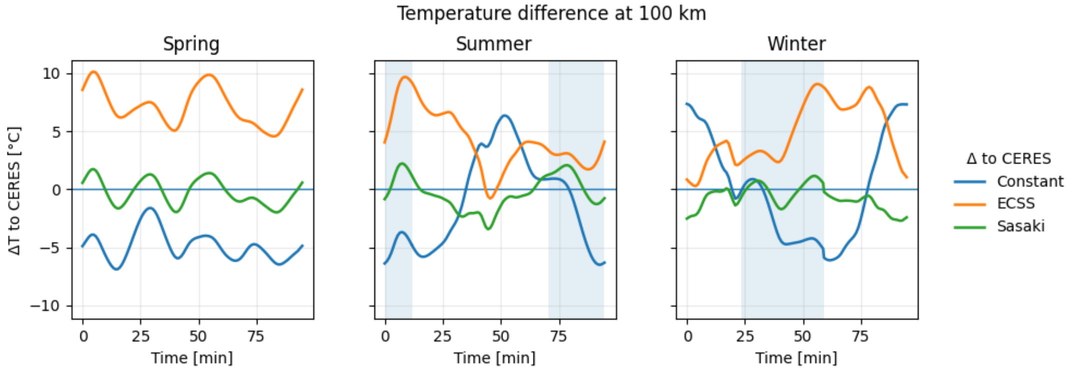

Figure 2.18: Aluminum plate temperature deviation ΔT relative to the CERES-driven reference case [67, 68], under the same 100 km test conditions as Fig. 2.17, for the selected seasonal cases. The three panels show the deviations for 6 March, 21 June, and 21 December 2022, allowing the seasonal dependence of the model differences to be compared directly. Shaded regions indicate eclipse intervals.

The maximum and mean absolute deviations of the plate surface temperature from the CERES-driven reference are summarised in Table 2.3. Across all three reference dates and the considered altitude range, the Sasaki model produces the smallest deviations: mean $|\Delta T|$ remains at or below approximately 1.2 °C, against between 3.8 and 5.2 °C for the constant model and between 4.6 and 7.2 °C for the ECSS method. The maximum instantaneous deviation reaches 3.6 °C for Sasaki at the summer solstice and remains below 2.9 °C at the other two reference dates, while the constant and ECSS models reach instantaneous peaks of between 7 and 11 °C across all dates. ECSS consistently exhibits the largest mean deviation across all three seasons, indicating a persistent systematic bias that the latitude-and-seasonal terms of the analytical formulation do not eliminate.

Table 2.3: Maximum and mean absolute temperature deviation $|\Delta T|$ of the simplified flat-plate test case relative to the CERES-driven reference [67, 68], for the constant model, the ECSS long-term averaging method [36, 56], and the Sasaki zonal harmonics model [38], evaluated over one orbit at each reference date.

	Winter solstice (21 Dec 2022)		Summer solstice (21 Jun 2022)		Equinox (6 Mar 2022)	
Model	Max [°C]	Mean [°C]	Max [°C]	Mean [°C]	Max [°C]	Mean [°C]
Constant	8.12	4.04	7.13	3.83	7.79	5.16
ECSS	10.09	5.26	10.23	4.63	10.79	7.21
Sasaki	2.86	1.02	3.62	1.20	2.15	1.05

The largest ΔT excursions in Figs. B.12 to B.14 coincide with the eclipse intervals, where neither solar nor albedo flux reaches the surface and the response is governed by OLR alone. Differences between the models in OLR representation therefore act undamped during these intervals, consistent with the OLR RMSE values reported in the previous subsection of approximately 12-16 W m^{-2} for the Sasaki model against approximately 30-34 W m^{-2} for ECSS and the constant representation. Outside eclipse, the temperature curves converge because the solar flux, the largest and most consistent term in the incident budget, outweighs the differences in both OLR and albedo.

The magnitudes shown in Table 2.3 are amplified by the low thermal capacity of the test plate. The thermal mass of the chosen 1 mm aluminium plate is $\rho\, c_p\, t \approx 2.4$ kJ m^{-2} K^{-1}, so the surface temperature equilibrates rapidly to instantaneous flux differences and inherits the full amplitude of any modelling error. Components with substantially larger thermal mass damp transient flux excursions through their heat capacity and consequently exhibit proportionally smaller temperature deviations. Conversely, thin elements such as radiators or deployable solar arrays respond on timescales comparable to the test plate and inherit deviations of similar order. The reported $|\Delta T|$ values therefore represent a conservative upper bound on the temperature impact of Earth radiation model selection, attained when the analysed component has a thermal time constant short compared to the orbital period.

A separate limitation concerns the completeness of the temporal sampling, independent of the thermal-mass scaling above. The three reference dates evaluated (the two solstices and the equinox of 2022) are a discrete sample of the annual cycle. Untested days could therefore produce larger deviations than reported in Table 2.3 for all three models. This has a direct consequence for model selection: against the 10 °C model-maturity margin adopted in Section 1.3, the ECSS model's sampled maximum (10.79 °C, equinox) already exceeds the budget on only three tested days, and the constant model's sampled maximum (8.12 °C) leaves comparatively little headroom for an untested worse day. The Sasaki model's sampled maximum (3.62 °C) retains substantially more headroom before approaching the same limit. This asymmetry in remaining margin, not only the smaller mean and peak deviations under the tested dates, supports adopting Sasaki as the baseline Earth radiation model.

Taken together, the structural comparison, the area-weighted RMSE against CERES, and the temperature response of the simplified test case all point to the same conclusion: the constant and ECSS representations introduce systematic thermal bias on the order of several degrees per orbit, while the Sasaki $L = 4$ zonal harmonics model reproduces the CERES-driven response within approximately 1 to 1.2 °C in the mean and 2 to 4 °C in the peak. The Sasaki model is therefore adopted as the baseline Earth radiation model in the remainder of this work.

2.5.5. Hot- and Cold-Case Earth Radiation Conditions

The Sasaki $L = 4$ formulation is used to construct the hot- and cold-case Earth radiation conditions propagated into the thermal model. In both cases, model uncertainty is incorporated by perturbing the nominal latitude-dependent profiles by the RMSE values reported in [38] for the $L = 4$ expansion,

$$\mathrm{RMSE}_{\mathrm{OLR}} = 31.38\ \mathrm{W\,m^{-2}}, \qquad \mathrm{RMSE}_{a} = 0.1148, \tag{2.42}$$

applied in the conservative direction for the case at hand.

Hot Case

For the hot case, the conservative direction corresponds to maximum radiative input from Earth, so the nominal Sasaki OLR and albedo coefficient are increased by their respective RMSE values. The hot-case epoch is chosen consistent with the atmospheric scenario of Section 2.1 as 2027-01-01 02:37:40.79 UTC. The resulting latitude-dependent profiles are shown in Fig. 2.19. The corresponding latitude-averaged values, suitable as equivalent inputs in lumped-parameter thermal networks that do not resolve the latitudinal dependence of the Earth radiation field, are $a_{\mathrm{hot}} = 0.5282$ and $\mathrm{OLR}_{\mathrm{hot}} = 250.73\ \mathrm{W\,m^{-2}}$.

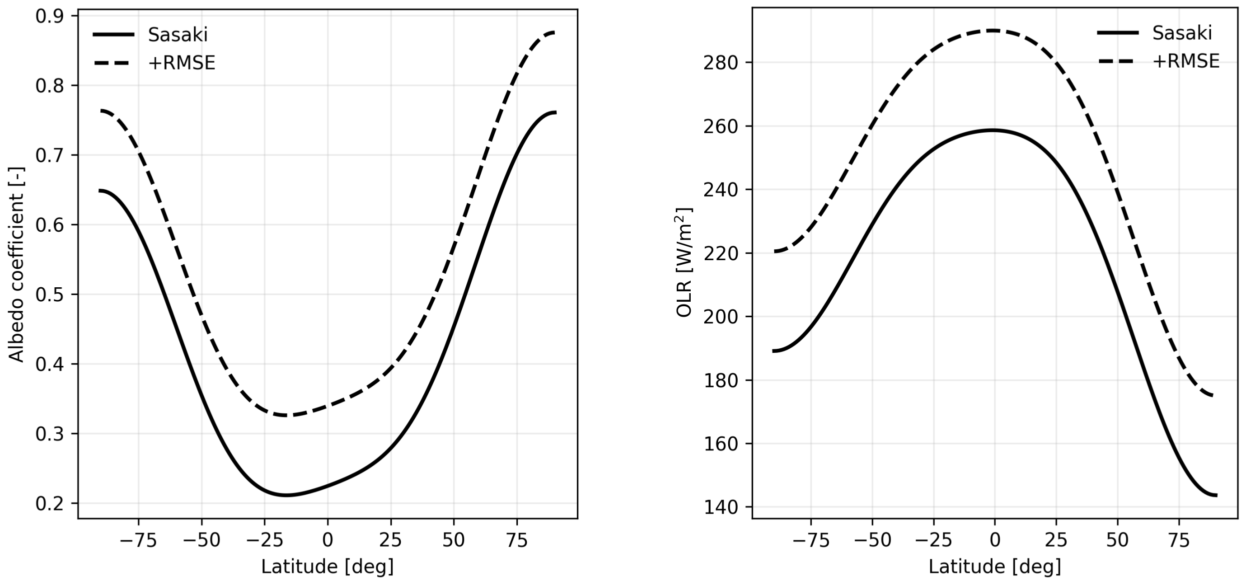

(a) Albedo coefficient. **(b)** Outgoing longwave radiation (OLR).

Figure 2.19: Latitude-dependent Earth radiation profiles defining the thermal hot-case boundary conditions on 2027-01-01. Solid curves are the nominal Sasaki $L = 4$ prediction [38]; dashed curves are the conservative upper bounds obtained by adding the model RMSE.

Cold Case

The cold case requires a combined selection criterion because the two radiative quantities have different physical dimensions and orders of magnitude: the albedo coefficient is dimensionless and bounded between 0 and 1, while OLR is expressed in $\mathrm{W\,m^{-2}}$ and varies between approximately 150 and 350 $\mathrm{W\,m^{-2}}$. A direct minimisation of the unweighted sum would be dominated by OLR and would not identify a physically representative cold day.

To define the cold case in a unit-independent way, the Sasaki albedo and OLR fields are first min-max normalised over the evaluated annual latitude-time domain,

$$\hat{x} = \frac{x - x_{\min}}{x_{\max} - x_{\min}}, \tag{2.43}$$

where x denotes either the albedo coefficient or the OLR. A combined cold-case score is then defined as the latitude-mean of the two normalised fields,

$$S_{\text{cold}}(\text{DOY}) = \frac{1}{N_\varphi} \sum_{i=1}^{N_\varphi} \frac{1}{2} \Big(\hat{a}_i(\text{DOY}) + \widehat{\text{OLR}}_i(\text{DOY}) \Big), \tag{2.44}$$

and the day of year that minimises S_{cold} is selected as the cold-case Earth-radiation epoch. The two terms are weighted equally, reflecting an a priori equal sensitivity to albedo and OLR uncertainty in the absence of a spacecraft-specific weighting.

Applying (2.44) to the mission year identifies 2027-04-08 (day of year 98, $S_{\text{cold}} = 0.5192$) as the cold-case Earth-radiation epoch; the seasonal evolution of S_{cold} is shown in Fig. 2.21. This date differs from the atmospheric cold-case epoch of 2029-08-02 identified in Section 2.1. The cold case is nevertheless propagated at the Earth-radiation epoch, because the cold reference orbit (550 km circular) lies at an atmospheric density several orders of magnitude lower than at the minimum-perigee hot case. The aerothermal contribution to the surface heat flux is therefore negligible at this altitude, and the cold case is driven by Earth radiation, eclipse cooling, and minimum internal dissipation. The nominal latitude-averaged Sasaki values on this date are $a = 0.3999$ and OLR = 221.51 $\mathrm{W\,m^{-2}}$.

The conservative cold-case profiles are obtained by subtracting the Sasaki RMSE from the nominal latitude-dependent profiles, with a lower clipping at zero to enforce physical bounds,

$$a_{\text{cold}}(\varphi) = \max\left(a_{\text{Sasaki}}(\varphi) - 0.1148,\ 0\right), \tag{2.45}$$

$$\text{OLR}_{\text{cold}}(\varphi) = \max\left(\text{OLR}_{\text{Sasaki}}(\varphi) - 31.38,\ 0\right). \tag{2.46}$$

The resulting profiles are shown in Fig. 2.20. Latitude-averaged values, used as equivalent lumped-parameter inputs, are $a_{\text{cold}} = 0.2851$ and $\text{OLR}_{\text{cold}} = 190.13\ \text{W}\,\text{m}^{-2}$.

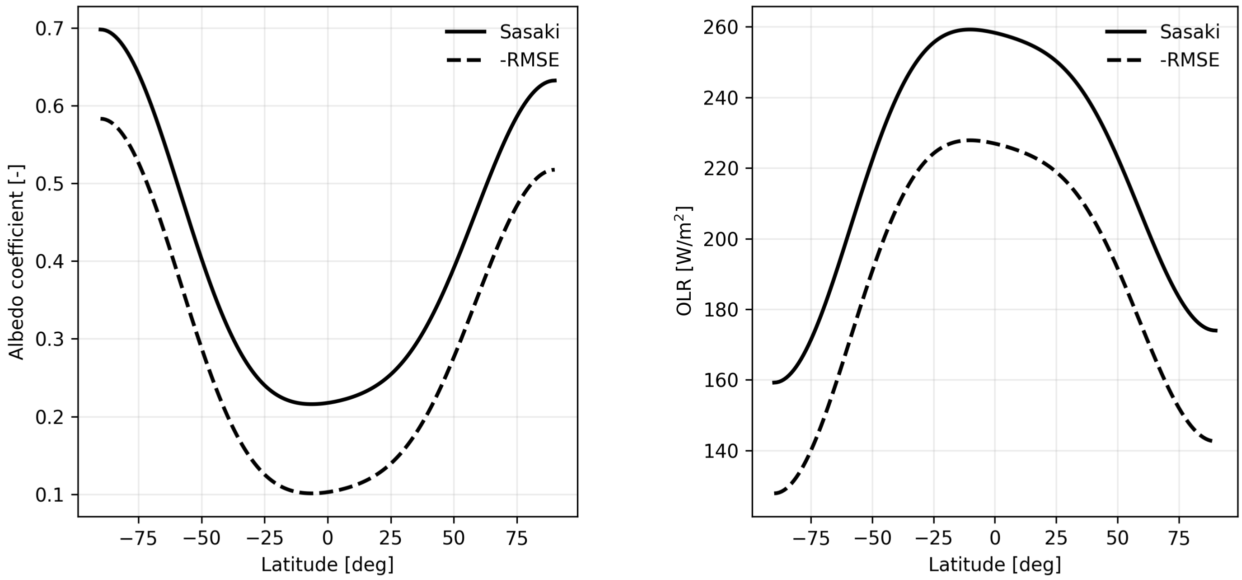

(a) Albedo coefficient.

(b) Outgoing longwave radiation (OLR).

Figure 2.20: Latitude-dependent Earth radiation profiles defining the thermal cold-case boundary conditions on 2027-04-08. Solid curves are the nominal Sasaki $L = 4$ prediction [38]; dashed curves are the conservative lower bounds obtained by subtracting the model RMSE and clipping at zero.

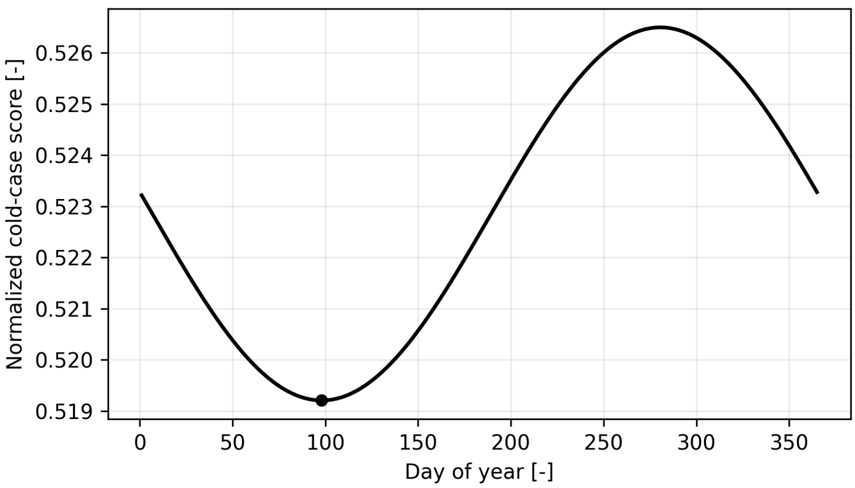

Figure 2.21: Combined cold-case score S_{cold} as a function of day of year for 2027, computed from min-max normalised Sasaki albedo and OLR fields [38]. The minimum value, attained on 2027-04-08, identifies the cold-case Earth-radiation epoch.

2.6. Analysis Cases

This section consolidates the environmental and operational inputs into the thermal analysis cases used in Chapters 3 and 4. The hot and cold cases are conservative envelopes: each combines the input direction that increases the severity of the corresponding thermal condition. The joint occurrence of all selected extremes is unlikely, but the resulting cases provide bounding inputs for deterministic thermal design.

Two case types are defined. Stationary reference cases provide fixed-point equilibrium estimates for rapid comparison, while transient orbital cases retain the time dependence of solar illumination, eclipse, altitude, attitude-dependent aerothermal heating, and operational power dissipation. The transient cases form the primary basis for the thermal performance assessment.

2.6.1. Stationary Reference Cases

The stationary reference cases combine the most adverse environmental and operational assumptions of the chapter at a fixed orbital point, with the spacecraft assumed to have reached steady state. The hot stationary case is evaluated at the minimum-perigee mission condition, $h_p = 130$ km in a dawn-dusk SSO, with conservative atmospheric, aerothermal, and Earth-radiation inputs and the highest-dissipation operational mode. The cold stationary case is evaluated at the 550 km circular reference orbit in eclipse, with the cold atmosphere, no direct solar input, and the lowest-dissipation operational mode. The parameter combinations are summarised in Table 2.4. The main thermal assessment is performed using the transient orbital cases of Section 2.6.2.

Table 2.4: Stationary thermal reference cases.

Case	Orbit	Environment	Power mode
Hot stationary	$h_p = 130$ km, dawn-dusk SSO	Hot atmosphere, hot aerothermal load, Sasaki Earth radiation with +RMSE	Payload acquisition
Cold stationary	550 km circular SSO, eclipse	Cold atmosphere, no direct solar input, Sasaki Earth radiation with −RMSE	Safe mode

2.6.2. Transient Orbital Cases

The transient orbital cases retain the time dependence of every input category developed in the preceding sections: orbital and mission parameters from Section 1.3; atmospheric state and aerothermal loading from Sections 2.1 and 2.3; Earth radiation from Section 2.5; attitude and stability from Section 2.4; and the operational power profile from Section 1.3.3. The complete case definitions are given in Table 2.5.

Table 2.5: Transient orbital thermal analysis cases. Parameters are grouped by input category, in the order in which they are developed in the present chapter.

Parameter	Hot transient case	Cold transient case
Orbital and mission state		
Reference orbit	$h_p = 130$ km, $h_a = 500$ km, SSO (LTAN 06:00), $i = 98°$	$h_a = h_p = 550$ km, SSO (LTAN 12:00) $i = 98°$
Mission phase	Minimum-perigee VLEO operation	Commissioning at reference altitude

Continued on next page

Table 2.5: Transient orbital thermal analysis cases. Continued.

Parameter	**Hot transient case**	**Cold transient case**
Eclipse condition	0 min per orbit	35-45 min per orbit
Atmosphere and aerothermal loading		
Atmospheric epoch	2027-01-01 02:37:40 UTC	2027-04-08
Solar and geomagnetic activity	$F_{10.7}$ = 135.1, A_p = 22.5	$F_{10.7}$ = 71.1, A_p = 6.7
Density profile	NRLMSISE-00, hot envelope, +100% ISO 14222 margin	NRLMSISE-00, cold envelope, −15% ISO 14222 margin
Relative velocity	$U_{\mathrm{orb}} + V_{\mathrm{corot}}$	$U_{\mathrm{orb}} - V_{\mathrm{corot}}$
Heat-flux model	Caruso simplified FMH, $\bar{\alpha} = 1$	Caruso simplified FMH, $\bar{\alpha} = 1$
Earth radiation		
Earth radiation model	Sasaki $L = 4$ zonal harmonics	Sasaki $L = 4$ zonal harmonics
Earth radiation epoch	2027-01-01	2027-04-08
Uncertainty treatment	Nominal + RMSE	Nominal − RMSE
Orbit-averaged albedo	$a = 0.5282$	$a = 0.2851$
Orbit-averaged OLR	$250.73\ \mathrm{W\,m^{-2}}$	$190.13\ \mathrm{W\,m^{-2}}$
Attitude and stability		
Nominal attitude	x-axis ram-aligned, z-axis nadir	x-axis ram-aligned, z-axis nadir
Pitch mode	Non-oscillatory	Non-oscillatory
Yaw mode	Oscillatory $\omega_{\beta,0}(h)$	Non-oscillatory
Yaw amplitude	$\theta_{\mathrm{dist,max}} = 1.23°$	–
Operational mode and power dissipation		
Operating mode	Nominal Mode	Safe mode
Internal dissipation	$f(Duty\,Cycle)$	7.7 W

The same set of case inputs is consumed by both the lumped-parameter thermal network and the finite-element thermal model presented in Chapter 3, with two implementation-level differences. First, the latitude-dependent Sasaki albedo and OLR fields are applied directly in the finite-element model, whereas the lumped-parameter model uses the orbit-averaged values listed in Table 2.5. Second, the aerothermal heat flux is applied face by face in the finite-element model through the geometry-resolved formulation of Section 2.3.1, whereas the lumped-parameter network applies an effective body-averaged free-molecular load derived from the same Caruso expression. In both representations, the yaw oscillation enters through the time-dependent incidence angles of the external surfaces, with $\beta(t) = \theta_{\mathrm{dist,max}} \sin(\omega_{\beta,0}(h)\,t)$ in the hot transient case and $\beta(t) = 0$ in the cold transient case.

These consolidated hot- and cold-case definitions complete the VLEO environmental input chain, providing the boundary conditions propagated into spacecraft temperatures by the thermal models developed in Chapter 3.

3
Thermal Modelling Methodology

Translating the VLEO-specific environmental loads and spacecraft characterisation of Chapters 1 and 2 into a temperature prediction requires representing conduction, radiation, environmental heat loads, and internal dissipation at a level of fidelity appropriate to each stage of the design process. This chapter develops the resulting modelling framework, which is then used for the thermal performance analysis and passive thermal-control sizing of Chapter 4. Section 3.1 first introduces the transient energy balance common to both modelling approaches. Section 3.2 then develops the lumped-parameter thermal network used for rapid system-level checks and heat-path interpretation. Section 3.3 defines the finite-element model used for spatially resolved design assessment, including the geometry, heat loads, thermal interfaces, and solver settings. Finally, Section 3.4 specifies the COMSOL radiation setup and selects the numerical view-factor method used in the spacecraft thermal analyses.

3.1. Thermal Balance Formulation

Both the lumped-parameter and finite-element thermal models are based on the transient thermal energy balance. For a thermal control volume, the temperature evolution is governed by the difference between the heat entering and leaving the control volume. Following the general formulation of Gilmore and Donabedian [35], this balance is written as

$$C_i \frac{\mathrm{d}T_i}{\mathrm{d}t} = \dot{Q}_{\mathrm{in},i} - \dot{Q}_{\mathrm{out},i}, \tag{3.1}$$

where C_i is the thermal capacitance of control volume i, T_i is its temperature, $\dot{Q}_{\mathrm{in},i}$ contains internally dissipated power and absorbed environmental heat loads, and $\dot{Q}_{\mathrm{out},i}$ contains heat removed from the control volume through conductive and radiative exchange.

In the lumped-parameter model, each control volume corresponds to a thermal node and Equation 3.1 is solved as a system of coupled ordinary differential equations. Conductive and radiative exchange are represented by equivalent coupling terms between nodes, while heat inputs are applied directly to the corresponding nodes.

In the finite-element model, the same energy balance is solved in spatially discretised form. Thermal capacitance and conduction are evaluated over the finite-element domains, while surface heat fluxes, surface-to-surface radiation, external radiation, and contact conductance are applied through boundary conditions and interface definitions. The finite-element model therefore retains the same physical balance as the lumped model, but resolves the temperature field within the spacecraft geometry instead of assigning one representative temperature to each thermal node.

3.2. Lumped-Parameter Thermal Network Model

The lumped-parameter thermal network provides the reduced-order branch of the modelling framework. The spacecraft is represented by isothermal nodes connected through equivalent conductive and radiative couplings, with environmental loads and internal dissipation applied as nodal heat inputs. The model is used for rapid case comparison, identification of dominant heat-transfer paths, and verification of the finite-element results.

3.2.1. Node Definition

Two nodal discretisations of the lumped-parameter thermal network are used in this work. The primary discretisation is the four-node network listed in Table 3.1, used for rapid case comparison and quick sizing during initial design iterations. The spacecraft is represented by the solar panels, the main structure, the heat shield, and one lumped node for all internal subsystems and payloads. Deep space is treated as a radiative sink at $T_s = 0\,\mathrm{K}$.

Table 3.1: Node definition of the four-node lumped-parameter thermal network.

Node	Node temperature	Represented spacecraft region
1	T_{SP}	Deployed solar panels
2	T_{STR}	Primary structure, side panels, and body-mounted solar panels
3	T_{SYS}	Internal subsystems and payloads, represented as one lumped node
4	T_{HS}	Heat shield

The structure node of the four-node network lumps four regions whose individual thermal roles are not interchangeable. The primary structure forms the conductive hub linking the remaining sub-nodes. The support structure carries the conductive path from the primary structure into the bus. The radiative side panels are the principal heat-rejection surfaces of the spacecraft. The body-mounted solar panels are externally absorbing surfaces. A secondary discretisation is therefore introduced for design and verification, in which the structure node is split into these four sub-nodes, while the solar-panel, heat-shield, and subsystem nodes are retained unchanged. The resulting seven-node network, shown in Figure 3.1, is used to verify the support-structure temperature obtained from the finite-element model. The temperature vector of the seven-node network is

$$\mathbf{T} = [T_{\mathrm{SP}}, T_{\mathrm{RSP}}, T_{\mathrm{BMSP}}, T_{\mathrm{STR}}, T_{\mathrm{SUP}}, T_{\mathrm{SYS}}, T_{\mathrm{HS}}]^{\mathrm{T}}\,, \tag{3.2}$$

where T_{RSP}, T_{BMSP}, T_{STR}, and T_{SUP} denote the radiative-side-panel, body-mounted solar-panel, primary-structure, and support-structure temperatures.

The general thermal balance of Equation 3.1 is applied to each node of either network. In steady-state analyses the node temperatures are constant, so $\mathrm{d}T_i/\mathrm{d}t = 0$ and the capacitance term vanishes. In transient analyses the thermal capacitance controls the response time of each node and is retained. For a node represented by a single equivalent material [35], the nodal thermal capacitance is defined as:

$$C_i = m_i c_{p,i}\,, \tag{3.3}$$

where m_i is the node mass and $c_{p,i}$ is the specific heat capacity. For nodes that lump several components or materials, the individual contributions are summed into one equivalent nodal capacitance.

For the four-node network the heat shield is assumed to be thermally isolated from the remaining spacecraft, exchanging heat only with deep space. The solar panels are assumed to be conductively isolated from the spacecraft structure, since the mechanical connection at the deployment hinge presents a small contact area and is therefore treated as a negligible conductive path relative to the panel-to-space radiative coupling. Direct radiative coupling between the spacecraft structure and the internal subsystem node is neglected, consistent with the assumption that internal heat transport between non-space-facing nodes is conduction-dominated. Under these assumptions, the resulting four-node topology is shown in Figure 3.3. The seven-node topology, shown in Figure 3.1, differs from the four-node topology by the local refinement of the structure node into four sub-nodes, and by the addition of a conductive coupling (G_{47}) between the heat shield and the primary structure node, relaxing the heat-shield isolation assumption adopted in the four-node network.

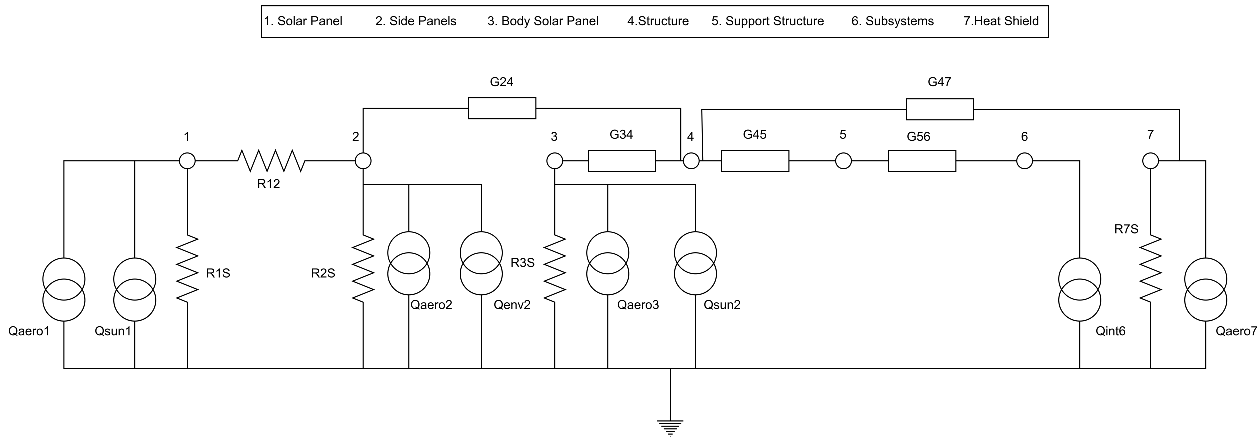

Figure 3.1: Seven-node lumped-parameter thermal network topology of the 16U CubeSat. Nodes represent: (1) deployed solar panel, (2) side panels, (3) body-mounted solar panel, (4) primary structure, (5) support structure, (6) subsystems, (7) heat shield. Zigzag elements denote radiative couplings (R_{ij}) to the space node at $T_s = 0\,\mathrm{K}$ (R_{1S}, R_{2S}, R_{3S}, R_{7S}) and between the solar panel and side panels (R_{12}); boxed elements denote conductive couplings (G_{ij}) between structural and internal nodes (G_{24}, G_{34}, G_{45}, G_{56}, G_{47}). Circular sources denote the applied heat loads, subscripted by node: Q_{aero} (aerothermal), Q_{sun} (direct solar), Q_{env} (Earth albedo and IR), and Q_{int} (internal dissipation) [W]. The structure (4) and support structure (5) receive no direct external heat load, coupled to the remaining nodes only through conduction.

3.2.2. Radiative and Conductive Coupling

Thermal coupling between lumped nodes is represented by equivalent conductive and radiative exchange terms. Conductive coupling accounts for heat transfer through structural connections, mechanical interfaces, and contact paths, while radiative coupling accounts for radiative exchange between spacecraft surfaces and radiative rejection to deep space. Each radiating surface is treated as diffuse-grey, such that the optical properties are assumed to be independent of direction and wavelength.

Using the thermal balance introduced in Equation 3.1, the lumped-node balance can be expanded as

$$C_i \frac{\mathrm{d}T_i}{\mathrm{d}t} = \dot{Q}_{\mathrm{in},i} - \sum_j R_{ij}\sigma\left(T_i^4 - T_j^4\right) - \sum_j G_{ij}\left(T_i - T_j\right), \tag{3.4}$$

where $\dot{Q}_{\mathrm{in},i}$ contains the internal dissipation and absorbed environmental heat loads applied to node i. The coefficient R_{ij} is the radiative coupling factor between nodes i and j, while G_{ij} is the equivalent conductive conductance between the same nodes. In this analysis, conductive heat exchange is assumed to occur through surface contact. The conductive coupling between nodes i and j, following Gilmore and Donabedian [35], is therefore evaluated as

$$G_{ij} = h_{c,ij} A_{ij}. \tag{3.5}$$

Here, $h_{c,ij}$ is the effective contact conductance in $\mathrm{W\,m^{-2}\,K^{-1}}$, and A_{ij} is the effective contact area in m^2. The contact conductance coefficient is selected based on ECSS-E-HB-31-01 Part 4A [75], which provides joint conductance values as a function of interface conditions and tightening torque.

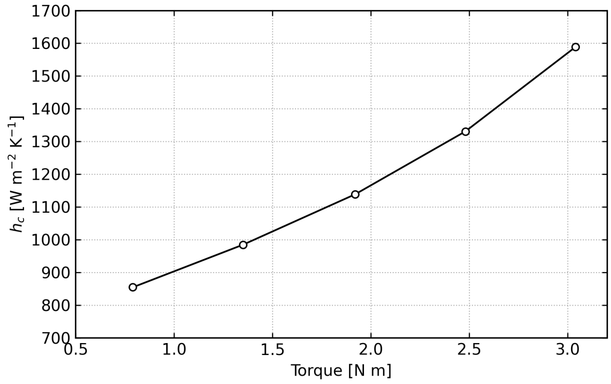

Figure 3.2: Contact conductance of joint interfaces as a function of tightening torque, adapted from ECSS-E-HB-31-01 Part 4A, Fig. 5.55. [75]

Radiative exchange is described through the radiative coupling factor R_{ij}. For diffuse-grey surfaces, Gilmore and Donabedian [35] define this factor as

$$R_{ij} = \varepsilon_i A_i B_{ij} = \varepsilon_j A_j B_{ji}, \tag{3.6}$$

where ε_i is the infrared emissivity of surface i, A_i is its radiating area, and B_{ij} is the Gebhart factor from surface i to surface j. The Gebhart factor B_{ij} represents the fraction of radiation leaving surface i that is ultimately absorbed by surface j, including multiple diffuse reflections. Following Gilmore and Donabedian [35], for an enclosure of N diffuse-grey surfaces it is defined as

$$B_{ij} = F_{ij}\varepsilon_j + \sum_{k=1}^{N} (1 - \varepsilon_k) F_{ik} B_{kj}, \tag{3.7}$$

where F_{ij} is the view factor from surface i to surface j. The view factor describes the purely geometrical part of the exchange and is defined as the fraction of radiation leaving surface i that directly reaches surface j. Following Gilmore and Donabedian [35], it can be expressed as

$$F_{i-j} = \frac{1}{A_i} \int \left(\int \frac{\cos\theta_i \cos\theta_j}{\pi S_{1-j}^2} dA_j \right) dA_i, \tag{3.8}$$

where dA_i and dA_j are differential elements within the two surfaces with S_{i-j} being the length of the connecting line between the two, while θ_i and θ_j are the respective angles between the connecting line and the normal vectors $\boldsymbol{n}_i$ and $\boldsymbol{n}_j$ to each surface [76].

The seven-node network retains four radiative links, shown in Figure 3.1: R_{12}, between the solar-panel and side-panel nodes, and R_{1S}, R_{2S}, and R_{7S}, rejecting heat from the solar-panel, side-panel, and heat-shield nodes to space. The heat shield and body-mounted solar panels have no radiative view of the remaining spacecraft and radiate only to deep space,

$$B_{7s} = B_{3s} = 1. \tag{3.9}$$

The solar-panel-to-side-panel view factor F_{12} is evaluated using the analytical relations of ECSS-E-HB-31-01 Part 1A [77], idealising both surfaces as rectangular planes sharing a common edge. The full derivation is given in Appendix C. The remaining, space-facing view factors then follow from view-factor summation. For the Earth view factor, needed for the albedo and infrared heat inputs, the spacecraft is treated as a small element viewing a spherical Earth, giving, for a nadir-facing surface [77],

$$F_{\mathrm{SC}\to E} = \left(\frac{R_E}{R_E + h} \right)^2, \tag{3.10}$$

where R_E is the Earth radius and h the spacecraft altitude; surfaces not aligned with nadir see a correspondingly reduced Earth view. These view factors, together with deep space as the remaining sink for the solar-panel and side-panel nodes, give the Gebhart factors of Equation 3.7, from which the radiative coupling factors used in the thermal network follow via Equation 3.6.

3.2.3. Environmental and Internal Heat Inputs

Environmental heat loads are applied to the network as nodal heat rates, obtained from the incident flux at each exposed node via

$$\dot{Q}_i = \eta_i q_i A_i, \tag{3.11}$$

using the solar absorptivity α_i for direct solar radiation and albedo, or the infrared emissivity ε_i for Earth infrared radiation. Earth albedo and infrared radiation are additionally scaled by each node's Earth view factor,

$$\dot{Q}_{E,i} = F_{E,i} A_{E,i} \left(\alpha_i q_A + \varepsilon_i q_{IR} \right), \tag{3.12}$$

with $F_{E,i}$ the Earth view factor and q_A, q_{IR} the albedo and infrared heat fluxes.

Aerothermal heating is evaluated face by face at the correct instantaneous angle of attack and yaw angle, and the resulting face-level heat rates are summed onto each node,

$$\dot{Q}_{aero,i} = \sum_f \eta_{i,f}\, q_{aero,i,f}(\alpha, \beta, t)\, A_{i,f}, \tag{3.13}$$

where the sum runs over the external faces assigned to node i. This preserves the correct total aerothermal input to each node. For stationary cases, all environmental inputs are prescribed as constants for the selected altitude, attitude, orbit condition, and thermal case. For transient orbital simulations they are time-dependent, $\dot{Q}_{env,i}(t)$. Internal dissipation follows the operational power modes of Section A.2 and is lumped onto the subsystem node in both networks, giving the total nodal heat input

$$\dot{Q}_{in,i} = \dot{Q}_{env,i} + \dot{Q}_{int,i}. \tag{3.14}$$

3.2.4. LTspice Implementation

The lumped thermal network is implemented in LTspice using an electrical-thermal analogy. Heat flow is represented by electrical current and thermal potential is represented by voltage. This allows the coupled nodal energy balances of the lumped model to be solved as an equivalent electrical circuit.

An important implementation choice is the definition of the network potential. This is required because conductive heat exchange is linear with temperature T, while radiative heat exchange depends on the fourth power of temperature, σT^4. To express all couplings in one network, the model can be formulated either in the linear potential T or in the radiative potential σT^4. In a linear-potential network, conductive couplings are implemented directly, while radiative couplings are linearised around a reference temperature. In a radiative-potential network, radiative couplings are implemented directly, while conductive couplings are converted to the radiative potential. The two potential definitions carry different physical units: T_i is expressed in K, while σT_i^4 is expressed in $\mathrm{W\,m^{-2}}$. This distinction propagates through every coupling and capacitance in the network, as summarised in Table 3.2. The choice depends on the dominant heat-transfer mechanism of the problem. For conduction-dominated systems, a linear-potential formulation is appropriate. For orbital thermal problems with strong radiative exchange, the radiative potential should be applied.

The linearisation is performed around a reference temperature T_0. For a coupling between two nodes, T_0 is chosen as a representative temperature of the connected nodes, commonly taken as the expected average temperature of the two node temperatures. Since the node temperatures are not known before the first solution is obtained, T_0 must initially be estimated.

Table 3.2: Thermal-electrical analogy used for the LTspice implementation.

Thermal quantity	Electrical equivalent	Relation	Units
Thermal potential	Voltage V	$V_i = T_i$ or $V_i = \sigma T_i^4$	K or $\mathrm{W\,m^{-2}}$
Heat flow $\dot{Q}$	Current I	$\dot{Q} \leftrightarrow I$	W
Heat input $\dot{Q}_i$	Current source	$\dot{Q}_i \leftrightarrow I_i$	W
Prescribed boundary potential	Voltage source or ground	$V_s = T_s$ or $V_s = \sigma T_s^4$	K or $\mathrm{W\,m^{-2}}$
Conductive coupling G_{ij}	Electrical conductance	potential-dependent	$\mathrm{W\,K^{-1}}$
Radiative coupling R_{ij}	Electrical conductance	potential-dependent	$\mathrm{m^2}$
Thermal capacitance C_i	Electrical capacitance	$C_i^{\mathrm{el}} = C_i$ for $V_i = T_i$	$\mathrm{J\,K^{-1}}$
		$C_i^{\mathrm{el}} = C_i/(4\sigma T_0^3)$ for $V_i = \sigma T_i^4$	$\mathrm{J\,m^2\,W^{-1}}$
Deep space	Ground	$V_s \approx 0$	K or $\mathrm{W\,m^{-2}}$

The coupling values implemented in LTspice depend on the selected network potential. If the network is formulated in the linear potential T, conductive couplings are used directly, while radiative couplings are linearised as

$$G_{ij}^{\mathrm{rad},T} = 4\sigma T_0^3 R_{ij}. \tag{3.15}$$

If the network is formulated in the radiative potential σT^4, radiative couplings are used directly, while conductive couplings are transformed as

$$R_{ij}^{\mathrm{cond},\sigma T^4} = \frac{G_{ij}^{\mathrm{cond}}}{4\sigma T_0^3}. \tag{3.16}$$

Here, T_0 is the reference temperature used for linearisation. The corresponding LTspice resistance is obtained from the reciprocal of the implemented conductance,

$$R_{ij}^{\mathrm{LTspice}} = \frac{1}{G_{ij}^{\mathrm{implemented}}}. \tag{3.17}$$

Since $\dot{Q}$ is represented by current in watts in both formulations, the units of $R_{ij}^{\mathrm{LTspice}}$ follow directly from the units of the network potential: $\mathrm{K\,W^{-1}}$ in the linear-potential network ($V_i = T_i$) and $\mathrm{m^{-2}}$ in the radiative-potential network ($V_i = \sigma T_i^4$). A representative LTspice implementation of the four-node transient network is shown in Figure 3.3. The circuit contains the radiative-potential resistances, the linearised conductive subsystem-to-structure coupling, the transformed nodal capacitances, and the applied environmental and internal heat-source terms.

For transient simulations, the initial thermal state is prescribed uniformly across all nodes, and time-varying heat loads are implemented as piecewise-linear current sources. The transient solver uses a modified trapezoidal integration scheme, which suppresses the numerical oscillations associated with standard trapezoidal integration while retaining its accuracy. The complete LTspice command syntax, current-source conventions, and integration-method selection are documented in Appendix D.

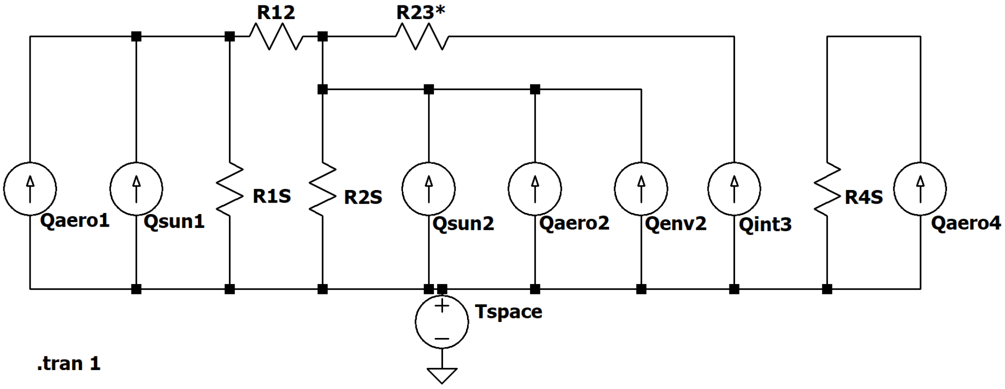

Figure 3.3: LTspice steady-state circuit implementation of the four-node lumped thermal network using a radiative-potential formulation ($V_i = \sigma T_i^4$). The circuit is solved with the transient solver (`.tran 1`) but reaches steady state instantaneously, since no capacitors are present. Radiative couplings (R_{12}, R_{1S}, R_{2S}, R_{4S}) appear directly as resistances $R_{ij}^{\text{LTspice}} = 1/R_{ij}$; the conductive structure-to-subsystem coupling (R_{23}^*, asterisked) is transformed to the radiative potential about the reference temperature T_0 following Equation 3.16. Environmental heat loads (solar, Earth infrared, albedo, aerothermal) and internal dissipation are applied as current sources at the corresponding nodes. Deep space is represented by an explicit voltage source referenced to circuit ground, $V_s = \sigma T_s^4 = 0$ for $T_s = 0\,\text{K}$.

3.2.5. Model Assumptions and Limitations

The lumped-parameter network relies on a set of simplifications that trade spatial and topological detail for computational speed and interpretability, consistent with the use of reduced-order network models for rapid case screening ahead of higher-fidelity analysis in the spacecraft thermal literature [50].

Each node is assumed isothermal, so one representative temperature describes an entire spacecraft region regardless of internal geometry, following directly from the lumped-capacitance formulation of Equation 3.3. Heat loads are correspondingly applied through an effective exposed area (Equation 3.11), or by summing face-level rates for multi-face nodes (Equation 3.13). Both operations preserve the total heat input but discard its spatial distribution and any internal conduction within the node. The predicted temperature is therefore a regional average.

The network topology carries further simplifications. Radiating surfaces are treated as diffuse-grey, and internal radiative exchange between non-space-facing nodes is neglected, consistent with conduction-dominated internal heat flow. The deployed solar panels are assumed conductively isolated, coupled only radiatively, since the deployment hinge presents a contact area small relative to the panel-to-space radiative path, and this should be reassessed if the hinge interface is redesigned. The four-node network further treats the heat shield as isolated, exchanging heat only with deep space and omitting its conductive path into the structure, which underestimates the peak structure temperature relative to a coupled interface. The seven-node network relaxes this isolation through the added heat-shield-to-structure conductance G_{47}.

Conductive interfaces are represented by equivalent conductances from interface-condition data via Equation 3.5 and the ECSS-E-HB-31-01 Part 4A values [75], substituting a single effective value for the true contact-pressure distribution, roughness, and mounting geometry. The conductance is therefore representative of the assumed interface type and torque condition, not the as-built contact area, and is only as reliable as the reference data's applicability to the actual hardware.

The LTspice implementation introduces a further linearisation. Since conductive coupling is linear in temperature while radiative coupling scales with T^4, a common network potential requires either linearising radiative couplings about a reference temperature T_0 or transforming conductive couplings to the radiative potential (Equation 3.15, Equation 3.16). As T_0 is unknown before a first solution is obtained, it is initially estimated from the expected node temperatures, so the linearised coupling is only as accurate as this estimate and as the choice of linearisation potential.

Material and optical properties are treated as constant across each node, consistent with the single-temperature representation implied by the node definition. Taken together, these simplifications place the lumped-parameter network at system level: useful for comparing thermal cases and identifying dominant heat paths, but not for resolving spatial gradients or localised design margins.

3.3. Finite-Element Thermal Model

The finite-element thermal model is implemented in COMSOL Multiphysics and forms the main design-assessment model. It resolves the spacecraft geometry as solid domains and boundary interfaces, allowing local temperature gradients, conductive heat paths, radiative exchange, and component-level heat dissipation to be represented explicitly. Stationary and transient studies are used depending on the thermal case.

The following subsections define the geometry, material properties, heat loads, thermal interfaces, and solver settings of the finite-element model.

3.3.1. Geometry, Component Representation, and Domain Material Properties

The COMSOL model represents the spacecraft with higher spatial detail than the lumped network, while retaining geometric simplifications where they do not control the main heat-transfer paths. Most internal subsystems are modelled as box-like domains based on their external dimensions and assigned effective bulk properties derived from mass budgets and available material data. This preserves their approximate thermal inertia and conductive behaviour without resolving their internal construction.

Higher geometric fidelity is retained for the primary structure, support structure, air intake, and heat shield. These components define the main conductive paths and external thermal interfaces of the spacecraft, and therefore have a stronger influence on the predicted temperature field. The dimensions, densities, thermal conductivities, and specific heat capacities used for the simplified component representations are listed in Table E.1 in Appendix E. Optical properties are assigned separately to the relevant external boundaries and are discussed in Section 3.4.

The resulting spacecraft geometry implemented in COMSOL is shown in Figure 3.4. The figure illustrates the complete spacecraft configuration, the simplified box-like representation adopted for the subsystem components, and the structural elements that are modelled with higher geometric fidelity.

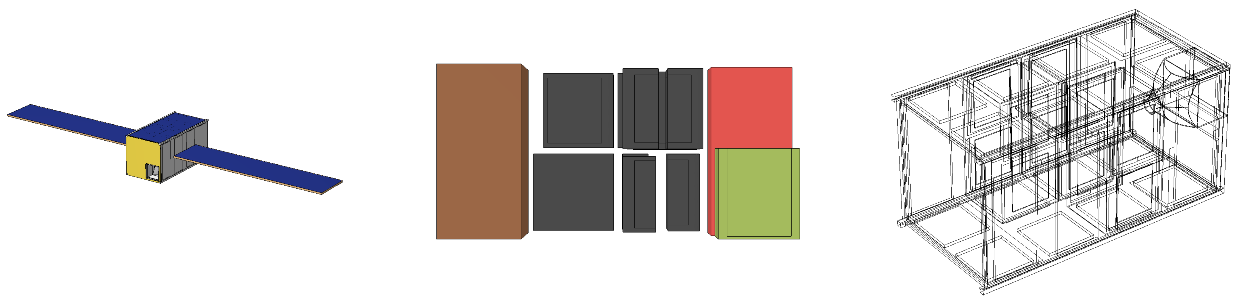

(a) Full spacecraft geometry including solar panels and external subsystems.

(b) Simplified box-like representation of subsystem components used in the thermal model.

(c) Structural elements modelled with higher geometric fidelity.

Figure 3.4: Geometric modelling approach of the 16U CubeSat thermal model implemented in COMSOL Multiphysics. (a) Full spacecraft configuration including deployed solar panels and all external-facing components. (b) Simplified box-like representation of internal subsystem components, in which each subsystem is assigned effective bulk material properties derived from mass-budget data. (c) Wireframe view of the structural elements retained with higher geometric fidelity (primary structure, support structure, air intake, and heat shield), which define the dominant conductive heat paths and external thermal interfaces of the model.

3.3.2. Internal and External Heat Load Implementation

Internal and external heat loads are applied to the COMSOL model as domain heat sources and boundary heat fluxes. Internal dissipation is assigned to the corresponding subsystem domains based on the power modes defined in subsection 1.3.3. For stationary analyses, the dissipated power is applied as a constant heat rate to each active component. For transient analyses, the same power terms can be made time-dependent to represent changes in operational state along the orbit. The internal power of subsystem k is applied to its corresponding COMSOL domain as a heat rate,

$$\dot{Q}_{\text{int},k} = P_k, \tag{3.18}$$

where P_k is the electrical power dissipated as heat in the selected operating mode. When the heat load is applied over a component volume, COMSOL distributes the heat rate over the selected domain. For time-dependent operation, the applied heat rate is multiplied by a switching factor,

$$\dot{Q}_{\text{int},k}(t) = s_k(t)P_k, \tag{3.19}$$

where $s_k(t)$ defines whether the subsystem is active. Smooth step functions are used instead of discontinuous switching in order to improve numerical convergence. For altitude-dependent operation, this is implemented using COMSOL's smoothed Heaviside function, for example

$$s_k(t) = \text{flc2hs}\left(h(t) - h_{\text{lim},k}, \Delta h\right), \tag{3.20}$$

where $h(t)$ is the spacecraft altitude, $h_{\text{lim},k}$ is the activation threshold, and Δh controls the smoothing interval. The switching threshold can be shifted slightly ahead of the intended activation point to avoid abrupt changes in the heat load. An example of the resulting operational indices is shown in Figure 3.5 for a case with a payload duty cycle of 25%.

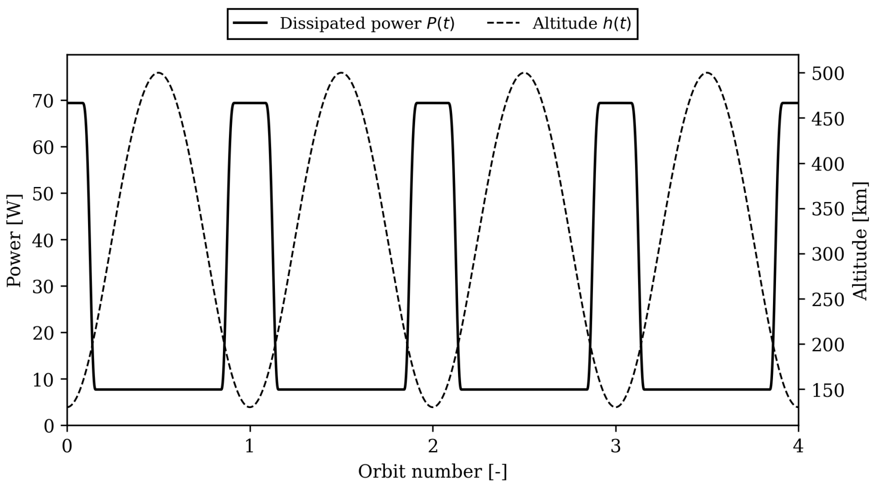

Figure 3.5: Altitude-dependent internal power dissipation for a representative transient orbit simulation with a duty cycle of 25%. The solid line shows the dissipated power $P(t)$ of the lumped subsystem node, switching smoothly between the Cold Mode (Safe Mode) and Hot Mode (Science Mode) power levels of Table 1.3, following the smoothed Heaviside function of Equation 3.20; the dashed line shows the corresponding altitude history $h(t)$ (right axis).

External heat loads are applied to the outer spacecraft surfaces. The model includes direct solar radiation, Earth albedo, Earth infrared radiation, and aerothermal heating. Solar radiation, albedo, and Earth infrared radiation are implemented through the orbital thermal load functionality and the planet-property definition in COMSOL. The Earth albedo and outgoing long-wave radiation distributions are imported from the environmental model and assigned as latitude-dependent radiative properties of the planet. COMSOL then discretises the visible Earth cap into radiation source points and computes the incident radiative flux on the spacecraft surfaces [78].

Aerothermal heating is applied separately as an inward boundary heat flux on the externally exposed surfaces. For each spacecraft face, heat-flux data are imported as interpolation functions of altitude, angle of attack, and yaw angle:

$$q_{\text{aero},f} = q_{\text{aero},f}\left(h, \alpha, \beta\right), \tag{3.21}$$

where f denotes the corresponding spacecraft face. These functions are evaluated during the simulation and applied as general inward heat flux boundary conditions,

$$-\mathbf{n}\cdot\mathbf{q} = q_{\text{aero},f}\left(h(t), \alpha(t), \beta(t)\right), \tag{3.22}$$

where $\mathbf{n}$ is the outward surface normal and $\mathbf{q}$ is the conductive heat-flux vector. This implementation allows the aerothermal load to vary with the instantaneous orbital altitude and spacecraft attitude.

For stationary analyses, the external heat loads are evaluated at the selected thermal case and applied as constant source terms or boundary fluxes. For transient analyses, they are evaluated as time-dependent functions of orbital state, attitude, and environmental conditions.

3.3.3. Boundary Conditions and Thermal Interfaces

External heat loads are applied to the exposed spacecraft surfaces as described in subsection 3.3.2. Surface-to-surface radiation is modelled using diffuse-grey surfaces and a two-band spectral representation with a band separation at 2.5 μm. Solar absorptivity is used for direct solar radiation and Earth albedo, while infrared emissivity is used for thermal emission and Earth infrared radiation. The assigned optical properties and radiation settings are given in Section 3.4.

COMSOL offers several approaches for representing thermal contact conductance at bolted or bonded interfaces, including the Cooper-Mikic-Yovanovich (CMY) and Mikic elastic analytical models, or the direct specification of a fixed conductance value [78]. Blanchete and Bah [52] applied the CMY and Mikic models to a 3U CubeSat in COMSOL and validated the predictions against thermal balance test data, finding that CMY outperformed Mikic in accuracy and satisfied-margin coverage, though both models require detailed geometric and material inputs, such as surface roughness, asperity slope, contact pressure, and microhardness. While the interface geometry and bolt torque are known, the surface roughness, asperity slope, and microhardness required by the CMY and Mikic models remain uncharacterised at this design stage. The ECSS-E-HB-31-01 Part 4A [75] reference conductance values are therefore used instead, consistent with the approach adopted for the lumped-parameter network. Applying the CMY model once the interface hardware is finalised is identified as a candidate refinement for future work.

3.3.4. Solver Settings and Study Types

The finite-element thermal model is solved using stationary and time-dependent studies in COMSOL. Stationary studies are used to evaluate bounding hot and cold cases and to obtain equilibrium temperature fields for selected altitude, attitude, power-mode, and environmental conditions. These solutions provide first estimates of component temperature margins and can also be used as initial conditions for transient simulations.

Time-dependent studies are used to resolve the orbital thermal response when internal power dissipation and external heat loads vary with time. The transient simulations are solved with an implicit time-dependent solver using the backward differentiation formula (BDF) method [42]. This method is suitable for the thermal problem because the model contains large thermal capacitances, radiative nonlinearities, and time-dependent heat inputs, which can lead to stiff transient behaviour.

For the transient studies, the output times are prescribed over the simulated orbital period, while the solver is allowed to take intermediate time steps as required for convergence and accuracy. A relative tolerance of 10^{-2} is used. The BDF order is limited to a maximum order of 2, with a minimum order of 1, and consistent initialization is performed using a backward Euler step. The solver settings used for the transient analyses are summarised in Table 3.3.

Table 3.3: Main COMSOL time-dependent solver settings used for transient thermal analyses.

Setting	Value
Solver type	Implicit
Time integration method	BDF
Steps taken by solver	Intermediate
Relative tolerance	10^{-2}
Maximum BDF order	2
Minimum BDF order	1
Consistent initialization	Backward Euler
Error estimation	Exclude algebraic

The physical plausibility of the solutions is checked by inspecting the resulting temperature ranges, verifying that critical components remain within expected thermal limits, and comparing the global temperature trends with the lumped-parameter thermal network. For transient analyses, additional checks are performed on the timing of temperature extrema relative to the applied orbital heat-load and power profiles.

3.4. Radiation Modelling in COMSOL

This section defines the COMSOL radiation setup. It first states the optical-property assumptions, then compares the available view-factor methods and selects the final numerical configuration used in the spacecraft thermal analyses.

3.4.1. Surface Optical Property Assumptions

As previously stated, all spacecraft surfaces are modelled as diffuse-grey. A two-band spectral representation is adopted, with the separation between the solar and ambient infrared wavelength bands set to 2.5 μm. The optical properties assigned to the different spacecraft components in the initial configuration of the model are listed in Table E.2 in Appendix E. These values are the starting point for the analyses of Chapter 4, with several of them treated as design variables and revised once the configuration is selected.

Optical properties for printed circuit board (PCB) structures and solar cells are taken from Finckenor and de Groh [79]. The antenna components are manufactured from Rogers 4003C. No optical property data could be identified in the literature for this material, so its properties are approximated using values reported for Astroquartz, which exhibits comparable radiative behaviour in spacecraft applications [35].

Given these assigned optical properties, the radiative heat exchange between spacecraft surfaces and the surrounding environment must be evaluated numerically. The Surface-to-Surface Radiation interface in COMSOL provides several approaches that differ in their treatment of spectral properties, view-factor evaluation, and computational cost. These approaches are described in the following subsection.

3.4.2. Radiation Modelling Approaches and Numerical Options

Surface-to-surface radiation in COMSOL is solved using a radiosity formulation, in which the radiative exchange between boundaries is determined from the surface optical properties and the corresponding view factors [80]. For the spacecraft geometry, these view factors are evaluated numerically.

COMSOL provides three view-factor evaluation methods [78]: direct area integration, hemicube, and ray shooting. Direct area integration evaluates the mutual irradiation between surface elements by directly integrating the view-factor relation introduced in Equation 3.8. Surface elements facing away from each other are omitted, but obstruction by intermediate surfaces is not considered. Consequently, this method does not account for shadowing [78].

The hemicube method evaluates visibility by projecting the surrounding geometry onto the faces of a virtual hemicube [80]. This enables mutual obstruction and shadowing to be included. The main numerical parameter is the radiation resolution. In three-dimensional models, the number of pixels on each hemicube face scales with the square of the selected resolution. Increasing the resolution improves the representation of small surfaces and shadowing effects, but increases the computational cost of the view-factor evaluation.

The ray shooting method evaluates irradiation by tracing rays from each surface element into the surrounding geometry [76]. This method can account for directional properties, specular reflection, transmission, and angular-dependent behaviour. The radiation resolution controls the number of rays launched from each element, while the tolerance, number of local adaptations, and maximum number of reflections define the accuracy and cost of the ray trajectories.

Radiation groups are used to restrict radiative exchange to selected boundary sets [78]. This reduces the number of surface-to-surface interactions and can therefore decrease computational cost. The grouping must be physically justified, since excluding relevant radiative paths can affect both the temperature field and the heat-balance consistency.

The radiation setup is selected by comparing the available view-factor evaluation methods in terms of view-factor convergence and solution time. This comparison is described in the following subsection.

3.4.3. Radiation Method Comparison

The radiation method comparison is performed to select an efficient numerical setup for the subsequent spacecraft thermal analyses. The comparison is based on a representative view-factor quantity to assess the radiation discretisation independently of the transient thermal response of the full model.

The convergence quantity is chosen as the area-averaged ambient view factor of the body panel adjacent to the solar panels. In the notation of the lumped radiation model, this corresponds to the view factor from surface 2 to the ambient space node,

$$F_{2S} = \overline{F}_{\mathrm{amb},2} = \langle \mathrm{otl.Famb2} \rangle_{A_2} \,. \tag{3.23}$$

Here, `otl.Famb2` is COMSOL's ambient view-factor variable for side 2 of the selected surface, and $\langle \cdot \rangle_{A_2}$ denotes an area average over that surface. The value represents the fraction of radiation leaving surface 2 that is not intercepted by other participating radiative boundaries and is therefore associated with the ambient environment.

The finite-element mesh is kept constant for all comparison cases, such that differences in the extracted view factor and solution time can be attributed to the radiation method and its numerical settings. Three radiation methods are compared: hemicube, ray shooting, and direct integration. For the hemicube method, the radiation resolution is varied from 32 to 1024. For ray shooting, the resolution is varied from 4 to 512. For direct integration, the integration order is varied from 0 to 8. All comparisons were performed on a single workstation (Intel Core i9-14900K, 24 cores / 32 threads, COMSOL restricted to 16 cores), so absolute solution times are hardware- and configuration-specific, though the relative ranking between methods is expected to hold generally.

For each setting, the selected view factor and the corresponding solution time are recorded. The objective is to identify the fastest radiation setup that provides a converged value of $\overline{F}_{\mathrm{amb},2}$ while retaining the radiation modelling capabilities required for the spacecraft model.

Adaptive resolution coarsening is not included in the comparison. Although this option can reduce the effective hemicube resolution and lower memory demand, the present comparison focuses on the influence of the radiation method and nominal radiation resolution on view-factor convergence and solution time [78]. The resulting comparison is shown in Figure 3.6.

The hemicube method reaches a stable value of

$$\overline{F}_{\mathrm{amb},2} \approx 0.79716$$

at low computational cost: refining the discretization from 32 to 512 subdivisions changes the extracted view factor by only 9×10^{-5} (0.011 %), while the solution time increases from 36 s to only 48 s. Ray shooting converges to a comparable value ($\overline{F} = 0.79741$, 0.031 % above the hemicube result), but only after 9341 s, 56 times the hemicube solution time. Direct integration follows the same trend, approaching $\overline{F} = 0.79734$ as the integration order increases from 0 to 8 (a 4.97 % change), though it requires 7388 s (44 times the hemicube solution time) to do so.

Across the broader view-factor literature, ray-shooting and Monte Carlo methods are consistently reported as the most computationally expensive of the standard techniques, while hemicube evaluation is rated substantially faster [50]. GPU-accelerated hemicube implementations have been reported to outperform ray-tracing-based view-factor calculations by up to three orders of magnitude [50], consistent with the present comparison. The higher cost of direct integration reflects its direct evaluation of mutual irradiation between every pair of facing surface elements, in contrast to the z-buffered hemicube projection that resolves each element's irradiation from all surrounding elements in a single pass [78].

The ECSS-based estimate used in the lumped-parameter model gives

$$F_{2S,\mathrm{ECSS}} = 0.78596,$$

consistently 1.4–1.5 % lower than the three converged numerical values (1.43 %, 1.46 %, and 1.46 % below the hemicube, direct integration, and ray shooting results, respectively). This offset is attributable to the zero-thickness surface assumption underlying the ECSS analytical formulation: solar panel thickness is neglected in the analytical model, whereas the panel edge faces are resolved explicitly in the hemicube, ray shooting, and direct integration meshes.

Based on the comparison, the hemicube method is selected for the final thermal model. It provides the best compromise between solution time, convergence of the selected ambient view factor, and compatibility with the required shadowing treatment. The selected nominal setting is a hemicube resolution of 256. Although lower resolutions already provide a stable value for the selected view-factor quantity, the higher resolution is retained to improve robustness for the complete spacecraft model, particularly with respect to backside irradiation and topology-consistency checks. These checks verify that one-sided radiating boundaries receive irradiation only on their physically radiating side. After the radiation configuration has been verified, the consistency check can be disabled to reduce computational cost [78].

The comparison is performed without radiation grouping. Radiation grouping is subsequently applied where appropriate to further reduce computational cost by restricting radiative exchange to physically relevant boundary sets. The final radiation setup therefore uses the hemicube method with a nominal resolution of 256 together with radiation groups.

Together, the lumped-parameter network and the finite-element model developed in this chapter provide two complementary levels of fidelity for representing the thermal response of a VLEO CubeSat. The lumped network is used for rapid case screening and heat-path interpretation. The finite-element model, using the radiation setup selected above, provides the spatially resolved predictions used for the thermal performance analysis and passive thermal-control sizing in Chapter 4. In this way, the chapter completes the thermal analysis framework.

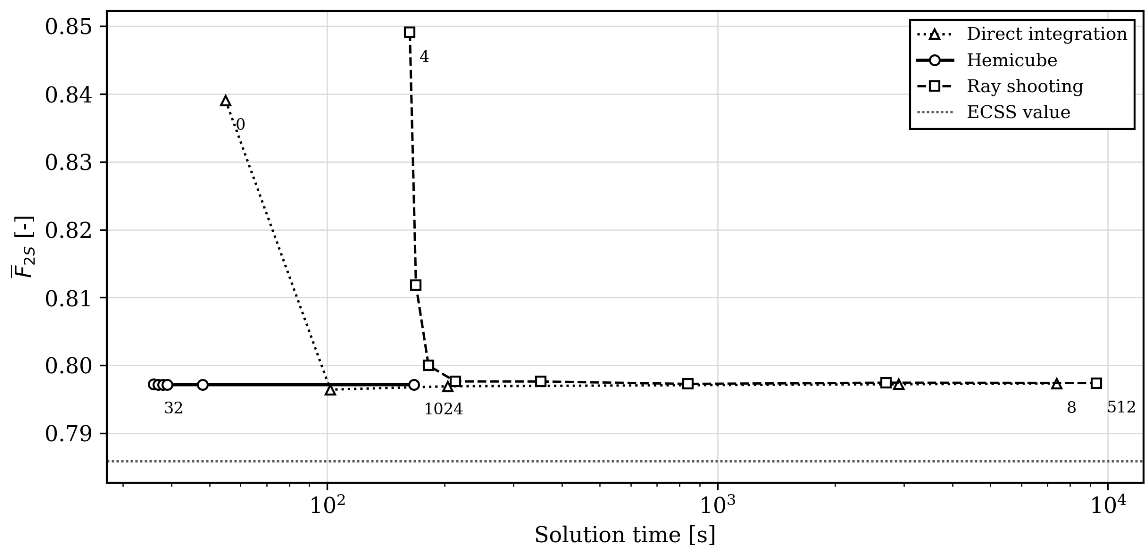

Figure 3.6: Radiation method comparison for the 16U CubeSat finite-element thermal model. The convergence quantity is the area-averaged ambient view factor $\overline{F}_{\mathrm{amb},2}$ of the body panel adjacent to the solar panels, evaluated at constant mesh resolution. Three view-factor methods are compared: hemicube (resolution 32-1024), ray shooting (resolution 4-512) with 1000 allowed reflections, and direct area integration (integration order 0-8). Solution time is shown on the secondary axis. The hemicube method converges to $\overline{F}_{\mathrm{amb},2} \approx 0.797$ at resolution 32 and retains this value through resolution 1024, at the lowest solution time of the three methods. The ECSS analytical estimate $F_{2S,\mathrm{ECSS}} = 0.786$ is shown for reference. The hemicube method at resolution 256 is selected for the final model.

4 Thermal Performance Analysis and Design Development

This chapter develops the passive bus thermal configuration from baseline assessment to uncertainty-based verification. Section 4.1 first evaluates the unmodified configuration in the bounding hot and cold cases and identifies the limiting thermal behaviour. Section 4.2 then applies Morris elementary-effects screening to rank the influential inputs and distinguish environmental drivers from actionable design levers. Based on this ranking, Section 4.3 sizes the retained design parameters through targeted sweeps and defines the selected passive configuration. Finally, Section 4.4 propagates the dominant uncertainties through this configuration and quantifies the residual probability of exceeding the structure-to-payload interface requirement.

4.1. Baseline Assessment

The baseline configuration corresponds to the reference platform defined in Section 1.3. It includes no dedicated thermal-control provisions beyond the heat shield already incorporated in that design. The remainder of the bus is assumed to be uncoated, with the aluminium side panels treated as polished aluminium and assigned an infrared emissivity of $\varepsilon_{sp} = 0.06$ [79]. This configuration therefore serves as the starting point against which the passive thermal-control modifications developed later in this chapter are evaluated.

This section assesses the thermal response of the baseline configuration for the minimum-perigee hot case and the cold reference orbit. It identifies the regions of the structure where heat accumulates and defines the design problem addressed in Section 4.3.

The first quantity evaluated is the structure sink temperature, computed with the lumped-parameter model. It is defined as the temperature of the spacecraft bus under external loads only, with internal dissipation $Q_{gen} = 0$. The sink temperature is independent of the operational duty cycle and therefore represents a lower-bound estimate of the achievable bus temperature in each orbital case. The converged structure sink temperatures for both the hot and cold cases are shown in Figure 4.1.

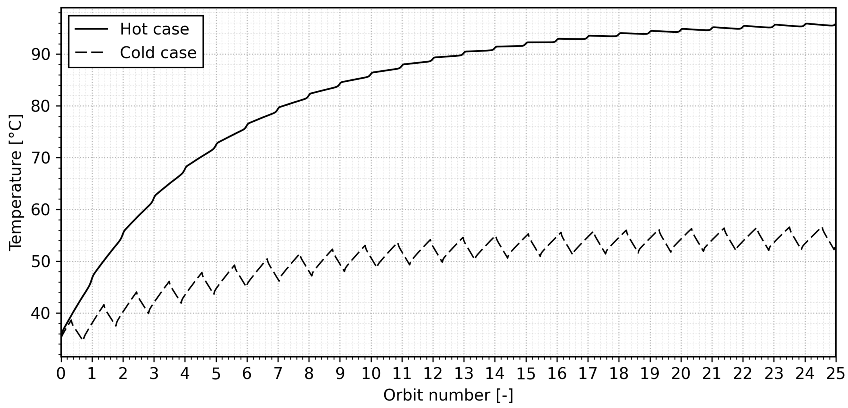

Figure 4.1: Baseline structure sink temperature exceeds the +40 °C hot-case design target for both orbital cases over 25 orbits, with zero internal dissipation (Q_{gen} = 0), computed with the lumped-parameter thermal network. The dawn-dusk hot case at h_p = 130 km reaches a steady-state value of 96 °C; the LTAN 12:00 cold case at h = 550 km reaches a periodic steady state with a cyclic minimum of 52 °C. The +40 °C target follows from the 50 °C structure-to-payload interface requirement with a 10 K uncertainty margin (Section 1.3.2).

For the hot transient case at h_p = 130 km, the structure sink temperature converges to a value of 96 °C, which lies 56 °C above the 40 °C hot-case design target. Since the exceedance occurs under zero internal dissipation, the baseline configuration cannot meet the hot-case requirement at the minimum perigee regardless of the operational duty cycle. For the cold transient case at the 550 km reference orbit, the minimum sink temperature over the orbital cycle converges to 52 °C, which lies 12 °C above the same 40 °C target. Because this is the cyclic minimum, the structure remains non-compliant at every point of the cold-case orbit.

The other components exposed to the orbital environment are not design-driving and are reported here for completeness. The hot case transient temperatures of the heat shield and deployed solar panels are shown in Figure 4.2. The heat shield remains below 100 °C throughout the hot case, well within its material capability. The deployed solar panels, assumed thermally isolated from the spacecraft bus at their mounting interface, fluctuate between 36 and 62 °C in the hot case and between −74.6 and 36 °C in the cold case, both within their qualified operating limits. Thus neither component drives the thermal design.

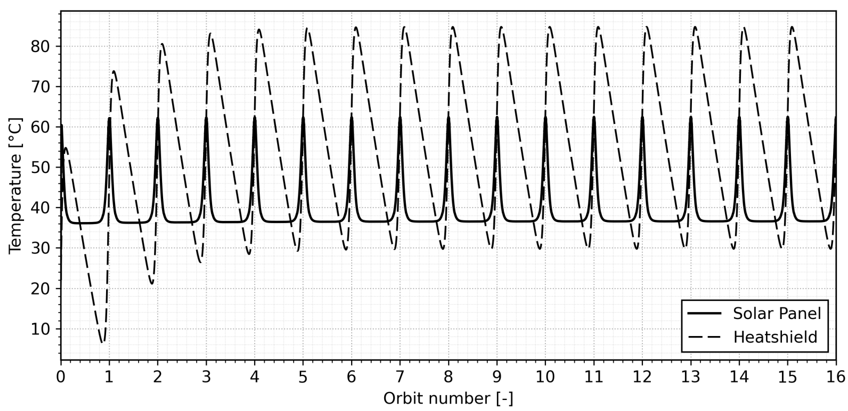

Figure 4.2: Baseline bus structure temperature distribution exceeds the +40 °C hot-case design target across the entire bus at the hot-case perigee (h_p = 130 km, Q_{gen} = 0, ε_{sp} = 0.06), computed with the COMSOL finite-element model; the heat shield is assumed insulated from the rest of the bus. Heat absorbed on the sun-facing surfaces conducts through the aluminium bus structure and is rejected exclusively via infrared emission from the side panels, with no dedicated radiator supplementing this path. Peak temperature reaches 99.8 °C on the body-mounted solar panel face; the opposing, self-shadowed face of the bus remains at 84.6 °C, reflecting the spatial variation in local radiative rejection capacity across the bus.

The lumped-parameter sink temperature establishes that the bus exceeds the hot-case design target at the structure level, but does not resolve where the excess heat accumulates spatially or which heat-rejection paths are available. To answer these questions, the COMSOL finite-element model is evaluated at the hot transient perigee under zero internal dissipation and baseline optical properties. The resulting steady-state structure temperature distribution is shown in Figure 4.3.

Figure 4.3: Steady-state spatial temperature distribution of the spacecraft bus structure at the hot-case perigee (h_p = 130 km, Q_{gen} = 0, ε_{sp} = 0.06), computed with the COMSOL finite-element model. Heat absorbed on the solar-illuminated surfaces conducts through the aluminium bus structure and is rejected exclusively via infrared emission from the side panels; no dedicated radiator supplements this path. Heat shield is assumed insulated from the rest of the bus. The resulting temperature gradient, 99.8 °C on the body-mounted solar panel face versus 84.6 °C on the opposing eclipse-facing side, reflects the spatial variation in local radiative rejection capacity across the bus.

Two features are visible in the temperature field. Peak temperatures of 99.8 °C occur on the side that carries the body-mounted solar panel, while the opposing side, which is in eclipse during the hot transient case, remains at 84.6 °C. The cross-bus temperature gradient reflects the dependence of the heat-rejection rate on the local surface emission, but at ε_{sp} = 0.06 neither side reaches the 40 °C design target. In the baseline configuration, side-panel radiation is the only available path by which the absorbed external load can be rejected to space, no dedicated radiator or thermally enhanced interface supplements this path. The result is that the bus temperature is governed by the balance between external load absorption and the radiative capacity of the side panels, independently of internal dissipation. The relative importance of individual conductive paths under operating duty cycles is examined in the subsequent design development, where the sensitivity of the peak structure temperature to each parameter is quantified directly.

It can be concluded that the baseline configuration is non-compliant in both the LTAN 12:00 cold case and the dawn-dusk hot case at h_p = 130 km. Both exceedances occur at Q_{gen} = 0, so reducing internal power dissipation cannot bring the bus into compliance in either case. The heat shield removes the single largest flux path, the ram-facing aerothermal load, but does not address the heat absorbed and conducted through the remaining bus surfaces, and an uncoated aluminium exterior lacks the radiative capacity to reject it at either bounding case. A heat shield alone is therefore not sufficient to ensure survivability at VLEO altitudes down to h_p = 130 km. The design problem is therefore to increase the bus heat-rejection capability beyond the heat shield alone, and the following sensitivity screening determines which design and uncertainty parameters govern this capability most strongly.

4.2. Sensitivity Screening

The baseline assessment identified the structure-to-component interface as the limiting hot-case quantity and the side-panel radiative path as the dominant heat-rejection mechanism. Before the design changes are sized, a sensitivity screening is performed to confirm which input parameters most strongly influence the maximum structure temperature $\max T_{\text{struc}}$, and to identify the secondary parameters that should be retained for the uncertainty propagation of Section 4.4.

4.2.1. Morris Elementary-Effects Method

The elementary-effects method of Morris [81], in the absolute-value form proposed by Campolongo and Cariboni [82], is used for the screening. For each input parameter x_i in the k-dimensional input vector $\mathbf{x} = (x_1, \ldots, x_k)$, the method computes a finite-difference sensitivity, the elementary effect:

$$EE_i(\mathbf{x}) \;=\; \frac{Y(x_1, \ldots, x_i + \Delta_i, \ldots, x_k) \;-\; Y(\mathbf{x})}{\Delta_i}, \tag{4.1}$$

where Δ_i is a finite perturbation applied to x_i and Y is the quantity of interest, here $\max T_{\text{struc}}$. The input space is sampled along r trajectories, each starting at a random base point and advancing through the parameters one at a time. The total computational cost is $r(k+1)$ thermal-model evaluations and is well below that of a full variance-based analysis.

Two statistics are extracted for each parameter from its r elementary effects. The mean of the absolute values,

$$\mu_i^* \;=\; \frac{1}{r}\sum_{j=1}^{r}\left|EE_i^{(j)}\right|, \tag{4.2}$$

ranks parameters by their overall influence on the output, with the absolute value preventing the cancellation of positive and negative effects that affects the original μ statistic of Morris. The standard deviation σ_i of the unsigned elementary effects measures how the effect of x_i varies across the input space and is large when the response is nonlinear in x_i or when x_i interacts with other parameters. The screening retains only the relative ordering of μ^* and σ, their absolute magnitudes are not interpreted as variance contributions.

4.2.2. Screened Parameter Set and Ranges

The screening parameters fall into three classes:

- **Design** parameters that can be modified by hardware or configuration choices prior to launch
- **Operational** parameters that can be adjusted in flight
- **Environmental** or modelling parameters that reflect the input uncertainty developed in Chapter 2

The complete set of input parameters and their ranges are listed in Table 4.1.

Table 4.1: Parameters and ranges used in the Morris elementary-effects screening. The quantity of interest is $\max T_{\text{struc}}$. Ranges are physical bounds, not probabilistic intervals.

Class	Parameter	Symbol	Units	Range
Design	Heat-shield-to-structure conductance	h_{HS}	$\mathrm{W\,m^{-2}\,K^{-1}}$	[0, 1000]
Design	Side-panel infrared emissivity	ε_{sp}	–	[0.06, 0.90]
Design	Side-panel solar absorptivity	α_{sp}	–	[0.16, 0.9]
Design	Bolted-interface contact factor	f_{c}	–	[0.5, 1.5]
Operational	Payload duty cycle	η_{pld}	–	[0, 1]
Environmental	Aerothermal heat-flux scale	k_{aero}	–	[0.5, 2.0]
Environmental	Yaw disturbance amplitude	θ_{dist}	deg	[0, 1.23]
Environmental	Yaw disturbance frequency	ω_{β}	Hz	[0, 0.05]
Environmental	Apogee altitude	h_{a}	km	[500, 550]

The design-parameter ranges span the realistic implementation space. h_{HS} varies from a fully isolated heat-shield interface to a bolted metallic contact. $\varepsilon_{\mathrm{sp}}$ spans from the bare aluminium structure value of 0.06 to a high-emissivity coating value of 0.90. The bolted-interface contact factor f_{c} scales the effective joint conductance relative to the nominal value determined from the mean tightening torque via the ECSS-E-HB-31-01 Part 4A look-up data [75]. It is varied over [0.5, 1.5] to reflect the uncertainty in the realised contact conductance about that nominal [55]. The operational and environmental ranges follow the hot-case definitions of Chapter 2, with θ_{dist} taken at its HWM-derived bound and ω_β taken at its altitude-dependent value, both from Section 2.4, and k_{aero} centred at unity to allow both an over- and an under-estimation of the aerothermal load relative to the nominal Caruso prediction.

4.2.3. Screening Results and Design-Lever Identification

The Morris elementary-effects statistics for the maximum structure temperature are shown in Figure 4.4. The full screening separates the input set into two clear groups. The environmental and modelling parameters carry the largest effects. The design parameters cluster closer to the origin. The hot-case structure temperature is therefore not only governed by the selected thermal-control configuration but also by the uncertain external forcing applied to it.

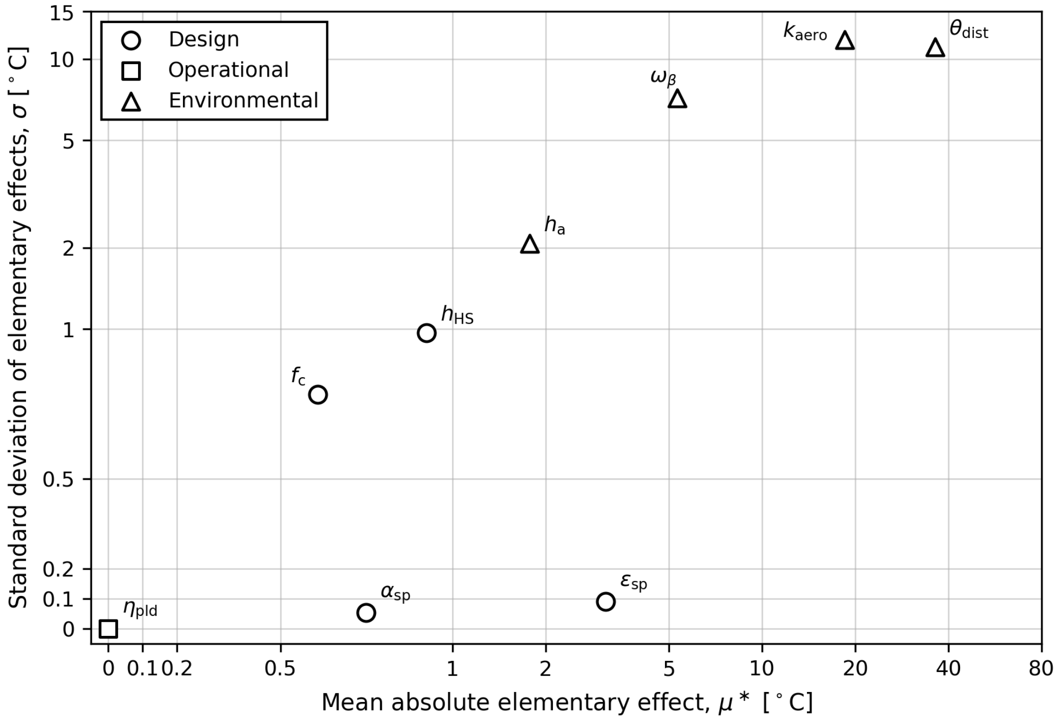

Figure 4.4: Morris elementary-effects screening for max T_{struc}, showing the mean absolute elementary effect μ^* against the standard deviation σ of the elementary effects for all nine screened parameters. Parameters with large μ^* exert a strong overall influence on the quantity of interest; large σ indicates nonlinear behaviour or interaction with other inputs. The environmental parameters dominate the response: the aerothermal scale factor k_{aero} and yaw disturbance amplitude θ_{dist} carry the largest mean effects, followed by the yaw disturbance frequency ω_β and apogee altitude h_{a}. Among the design and operational parameters, the side-panel emissivity $\varepsilon_{\mathrm{sp}}$ and heat-shield conductance h_{HS} are resolved as the dominant and second-ranked hardware levers, respectively; the remaining parameters (α_{sp}, f_{c}, η_{pld}) are compressed near the origin and cannot be ranked reliably from this screening alone.

The yaw disturbance amplitude θ_{dist} and the aerothermal heat-flux scale factor k_{aero} produce the dominant effects. Both also show large σ, indicating that their influence is not purely linear over the screened range. This is consistent with the VLEO hot-case environment, where the applied heat input depends directly on the aerothermal load level and on the attitude-dependent exposure of the spacecraft surfaces. The yaw disturbance frequency ω_β produces a smaller but non-negligible effect. The apogee altitude h_a is the weakest environmental contributor but still exceeds most design parameters in μ^*, with the exception of the side panel emissivity ε_{sp}.

This ranking justifies the worst-case treatment of the environmental inputs adopted in Chapter 2. The hot-case definition fixes each environmental parameter at the bound that maximises the heat load on the structure. Because these parameters dominate the variability of $\max T_{struc}$, anchoring them at their worst-case bounds establishes a design envelope within which the configuration must remain compliant. The screening therefore confirms the appropriateness of the envelope approach.

This bounding treatment is not equally exhaustive for every environmental parameter, however. The apogee altitude h_a and the aerothermal heat-flux scale factor k_{aero} are fixed at bounds that already represent the true worst case: h_a is set by the mission profile itself, its screened range of 500-550 km spanning the hot- and cold-case envelope endpoints of Chapter 2, and k_{aero} is anchored at the conservative bound of the aerothermal model. Neither retains meaningful residual uncertainty once that bound is adopted, so both are held fixed and excluded from the propagation in Section 4.4. The yaw disturbance amplitude θ_{dist} and frequency ω_β each retain residual uncertainty after the hot-case bound is applied, for distinct physical reasons. $\theta_{dist,max} = 1.23°$ is the maximum wind-induced disturbance angle over the evaluated hot- and cold-case atmospheric datasets, adopted as a single altitude-independent bound. It is a finite-sample maximum under the HWM wind model, not a closed bound like h_a or k_{aero}, so residual uncertainty in the amplitude remains.

ω_β carries uncertainty of a different kind. By (2.33), $\omega_{\beta,0} = \sqrt{q k_\beta / I_{zz}}$, where the dynamic pressure q is environmental, the aerodynamic stiffness k_β depends on both the environment and the design, and the moment of inertia I_{zz} is set by the design mass distribution. ω_β is therefore a hybrid environmental-design quantity: its screened range of $[0, 0.05]$ Hz reflects altitude variation alone, with k_β and I_{zz} fixed at the configuration selected in Section 2.4. This is why it is retained alongside the design levers in Figure 4.5, in addition to the full screening. Because the screening identifies these same parameters as dominant drivers of $\max T_{struc}$, their residual uncertainty is propagated in Section 4.4. This treats θ_{dist} and ω_β, together with the four design levers identified below, as random variables, and propagates their combined effect on $\max T_{struc}$ through the selected configuration. Restricting the propagation to these six parameters also keeps the surrogate-based uncertainty propagation of Section 4.4 computationally tractable.

Among the design parameters, ε_{sp} has the largest mean elementary effect, identifying the radiative rejection capability of the side panels as the primary hardware lever for the bus heat balance. Its small σ relative to its μ^* indicates a largely monotonic response over the screened range. The heat-shield-to-structure conductance h_{HS} is the second resolved design contribution, with a standard deviation comparable to its mean elementary effect, indicating a coupled response, as the conductive heat leak from the ram face becomes more significant when the side panels cannot reject heat radiatively and less significant when they can. The remaining design parameters α_{sp}, f_c, and η_{pld} are compressed near the origin and their relative ranking cannot be established from the full screening, because their individual effects are small compared to the environmental dominators that drive the spread of the trajectories. A design-only screening is therefore performed to isolate the relative importance of the hardware levers from the environmental forcing and to identify which design parameters should be prioritised in the subsequent sizing. The result is shown in Figure 4.5.

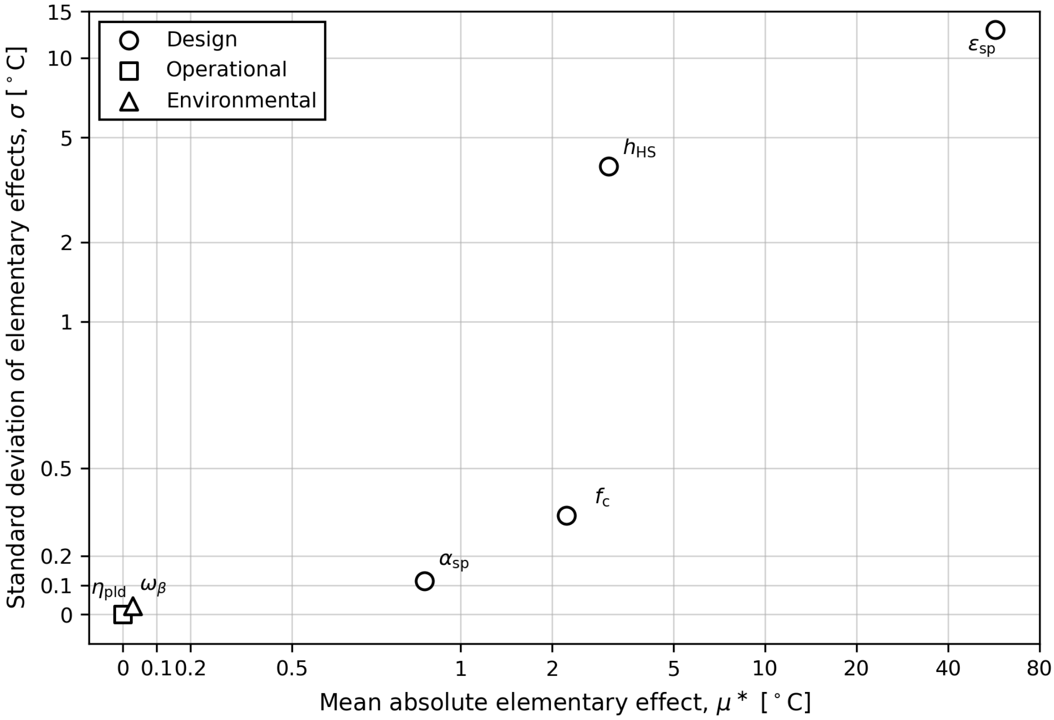

Figure 4.5: Morris elementary-effects screening restricted to the four design parameters for max T_{struc}, with the payload duty cycle η_{pld} (operational) and yaw disturbance frequency ω_β (environmental). Removing the dominant environmental drivers of the full screening resolves the relative ranking of the design levers compressed near the origin in that result. The side-panel emissivity ε_{sp} emerges as the primary design lever, followed by the heat-shield-to-structure conductance h_{HS}, which exhibits a large σ relative to its μ^*, indicative of coupling with the surrounding configuration. The bolted-interface contact factor f_c ranks third. The solar absorptivity α_{sp} produces a negligible effect over its admissible range, and both η_{pld} and ω_β collapse to near-zero effect once the dominant environmental parameters are no longer co-varied.

The design-only screening confirms ε_{sp} as the dominant design lever and supports increasing the side-panel emissivity to improve radiative heat rejection from the bus. The heat-shield-to-structure conductance h_{HS} is resolved as the second design contribution, with a standard deviation comparable to its mean elementary effect, consistent with the coupled behaviour already observed in the full screening. The bolted-interface contact factor f_c ranks third. Because f_c lumps several interfaces into one variable, the screening does not attribute the effect to a specific joint. The dominant interface is identified at the configuration level in Section 4.3. The solar absorptivity α_{sp} and payload duty cycle η_{pld} produce negligible effects over their admissible ranges. The yaw disturbance frequency ω_β collapses to near zero effect when the other environmental parameters are not co-varied. This confirms that its apparent contribution in the full screening originates entirely from interaction with θ_{dist} and k_{aero}.

The combined interpretation of both screenings yields two outcomes. The environmental parameters, led by θ_{dist} and k_{aero}, dominate the variability of max T_{struc} and validate the worst-case envelope adopted in the environmental definition. The main design levers are the side-panel emissivity, the aggregated bolted-interface conductance, and the heat-shield-to-structure conductance. Together with the baseline assessment of Section 4.1, this screening establishes how the VLEO environment and the spacecraft configuration determine the hot-case thermal response. The sizing and verification that follow address which passive configuration keeps that response within the interface requirement.

4.3. Thermal Design Development

The sensitivity screening identified the side-panel emissivity, the heat-shield-to-structure conductance, and the body-mounted solar panel interface conductance as the relevant hardware levers. Each is examined through a targeted parametric sweep at the nominal hot-case operating point. The resulting configuration is verified against the cold-case requirement and passed to the reliability assessment of Section 4.4.

4.3.1. Effect of Heat-Shield Interface Conductance

The screening of Section 4.2 identified the heat-shield-to-structure conductance h_{HS} as a relevant design parameter. Its mean elementary effect on max T_{struc} was measurable, and its standard deviation indicated coupling with the surrounding thermal configuration. A parametric sweep is performed to characterise the contribution of h_{HS} in isolation. The sweep is run at the nominal hot-case condition with zero internal dissipation ($Q_{\mathrm{gen}} = 0$). The contact conductance is varied from 0 W m^{-2} K^{-1} to 1000 W m^{-2} K^{-1}. The lower bound represents a fully insulated heat-shield interface and the upper bound a metallic bolted interface. The reported sweep is taken at a fixed side-panel emissivity of $\varepsilon_{\mathrm{sp}} = 0.40$, which differs from the baseline value of $\varepsilon_{\mathrm{sp}} = 0.06$ established in Section 4.1.

The choice of an elevated emissivity for this sweep avoids a competing radiative mechanism that becomes active at the baseline value. When the side panels cannot reject the absorbed heat load efficiently, as is the case at $\varepsilon_{\mathrm{sp}} = 0.06$, the heat shield itself begins to behave as a partial radiative surface. Heat returns into the shield through the interface and is emitted from its front face. In the baseline configuration the heat shield remains cooler than the structure for most of the orbit (30-85 °C, Figure 4.2, versus a 96 °C structure sink temperature, Section 4.1), so an unconstrained interface would conduct heat from the hotter structure into the cooler shield, the reverse of the direction exercised below. The elevated emissivity of $\varepsilon_{\mathrm{sp}} = 0.40$ suppresses this reversal and exercises h_{HS} against a single, unambiguous heat-transfer direction.

The maximum structure temperature from the sweep is shown in Figure 4.6.

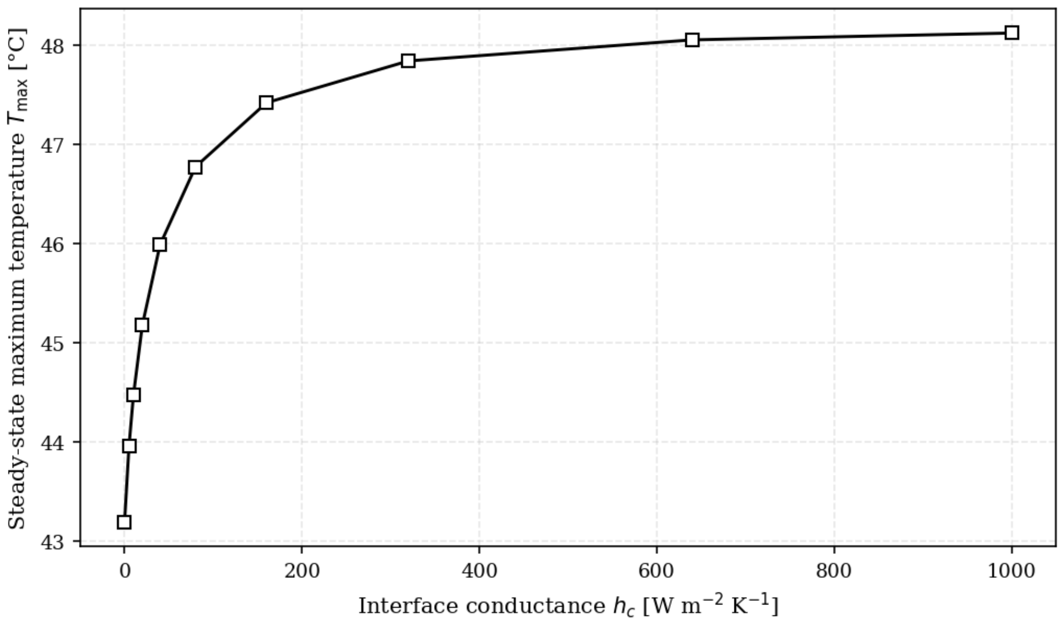

Figure 4.6: Maximum structure temperature in the steady-state portion of the hot transient case as a function of the heat-shield interface conductance h_{HS} (h_c in the plot), at zero internal dissipation ($Q_{\mathrm{gen}} = 0$) and side-panel emissivity $\varepsilon_{\mathrm{sp}} = 0.40$. The response is essentially binary: temperature rises by 4.65 °C (from 43.2 to 47.8 °C) between the insulated case and $h_{\mathrm{HS}} \approx 320$ W m^{-2} K^{-1}, then changes by only 0.28 °C up to $h_{\mathrm{HS}} = 1000$ W m^{-2} K^{-1}.

The response is essentially binary. Between $h_{\mathrm{HS}} = 0$ and $h_{\mathrm{HS}} = 320$ W m^{-2} K^{-1} the maximum structure temperature rises by 4.65 °C, from 43.2 to 47.8 °C. Between $h_{\mathrm{HS}} = 320$ and $h_{\mathrm{HS}} = 1000$ W m^{-2} K^{-1} the maximum temperature changes by 0.28 °C and is essentially flat. The contact conductance therefore acts as an effective switch. Once a meaningful conductive contact is established between the heat shield and the structure, the conductive resistance ceases to be the limiting step in the heat path from the shield into the bus. Further increases in h_{HS} produce no significant effect on the bus temperature.

The design implication is that insulating the heat-shield interface is the relevant design action. The precise value of the residual contact conductance is not critical as long as the interface is kept outside the upper plateau. Heat-shield isolation is particularly relevant for operation at the lower end of the perigee altitude range. The aerothermal load on the shield is largest there, and the conductive leak into the bus is correspondingly higher. The same isolation that produces a few kelvin of margin at the present minimum-perigee operating point produces a larger margin at altitudes below it. Heat-shield isolation is therefore one of the practical levers for extending the operational altitude range downward.

The heat-shield interface geometry has not yet been fixed at this design stage. The implementation through stand-offs, low-conductivity washer stacks, or a discrete-point mounting concept would each yield a different effective h_{HS}. The interface conductance is therefore retained as an uncertain input, propagated through the reliability assessment of Section 4.4 over the same $[0, 1000]$ W m^{-2} K^{-1} interval swept above. For the deterministic sizing of the remaining design parameters, the conservative bound $h_{HS} = 1000$ W m^{-2} K^{-1} is adopted, except explicitly stated otherwise. Given the 0.28 °C plateau established above, this bound overstates the conductive heat leak by no more than 0.28 °C relative to any conductance above 320 W m^{-2} K^{-1}.

4.3.2. Effect of Side-Panel Emissivity

The side-panel infrared emissivity ε_{sp} is the dominant design parameter identified by the screening. It controls the radiative heat-rejection capability of the spacecraft bus. The sweep is performed with the body-mounted solar panel interface conductance h_{MSP} included as a second parameter. The body-mounted solar panels and the heat shield form the main external surfaces through which absorbed environmental heat can be conducted into the bus. The hot case is evaluated at zero internal dissipation, yielding the structure sink temperature. The cold case is evaluated at the safe-mode dissipation of Section 1.3.3, since the continuous safe-mode heat input is not negligible against the cold-case Earth-radiation and eclipse-cooling budget.

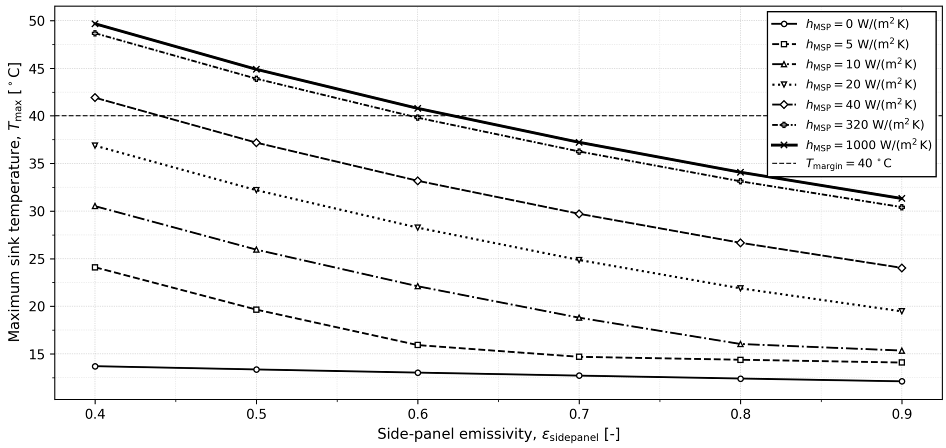

Figure 4.7: Maximum structure sink temperature in the hot transient case as a function of the side-panel emissivity ε_{sp}, for seven values of the body-mounted solar panel interface conductance h_{MSP}, at zero internal dissipation ($Q_{gen} = 0$). The interface conductance is the dominant control variable: with the interface insulated ($h_{MSP} = 0$) the sink temperature remains near 13 °C across the full emissivity range, while for $h_{MSP} \geq 320$ W m^{-2} K^{-1} the curves collapse to within 1 °C of each other and cross the +40 °C hot-case design target only at $\varepsilon_{sp} \approx 0.60$-0.62. Side-panel emissivity acts as a secondary lever, reducing sink temperature progressively along each h_{MSP} curve.

The hot-case result is shown in Figure 4.7. The body-mounted solar panel interface conductance is the dominant control variable. With the interface fully insulated ($h_{\mathrm{MSP}} = 0$) the maximum sink temperature stays close to 13 °C across the entire emissivity range. For $h_{\mathrm{MSP}} \geq 320\ \mathrm{W\,m^{-2}\,K^{-1}}$ the curves are essentially collapsed, differing by less than 1 °C across the entire emissivity range, and cross the 40 °C design target at $\varepsilon_{\mathrm{sp}} \approx 0.60$–$0.62$. At $\varepsilon_{\mathrm{sp}} = 0.5$ the maximum sink temperature rises from 13 °C at $h_{\mathrm{MSP}} = 0$ to 32 °C at $h_{\mathrm{MSP}} = 20\ \mathrm{W\,m^{-2}\,K^{-1}}$ and 37 °C at $h_{\mathrm{MSP}} = 40\ \mathrm{W\,m^{-2}\,K^{-1}}$, before saturating to 44 °C for $h_{\mathrm{MSP}} \geq 320\ \mathrm{W\,m^{-2}\,K^{-1}}$.

The response in h_{MSP} is binary in the same practical sense observed for the heat-shield interface conductance in Section 4.3.1. The body-mounted solar panel is attached through a bolted COTS interface. For this type of interface, ECSS-E-HB-31-01 Part 4A gives a torque-dependent joint conductance of approximately h_{MSP} = 800-1000 $\mathrm{W, m^{-2}, K^{-1}}$ [75]. This range lies entirely within the saturated plateau shown in Figure 4.7, where the resulting hot-case temperature is insensitive to the exact realized conductance, with variations below 1 °C. As currently configured, the interface is therefore effectively fixed in the high-temperature regime regardless of the specific tightening torque applied within the bolted joint's normal range. Only a deliberate insulating redesign, of the same kind identified for the heat-shield interface, would move the design into the qualitatively different low-h_{MSP} regime. Whether this insulation is achievable depends on the mounting hardware, and it is not a given at the present level of design. The side-panel emissivity itself acts as a secondary control variable that reduces the hot-case temperature progressively across the screened range. The emissivity is selected at $\varepsilon_{\mathrm{sp}} = 0.90$. This is the highest value in the screened range and produces the lowest hot-case sink temperature regardless of the regime in which h_{MSP} ends up.

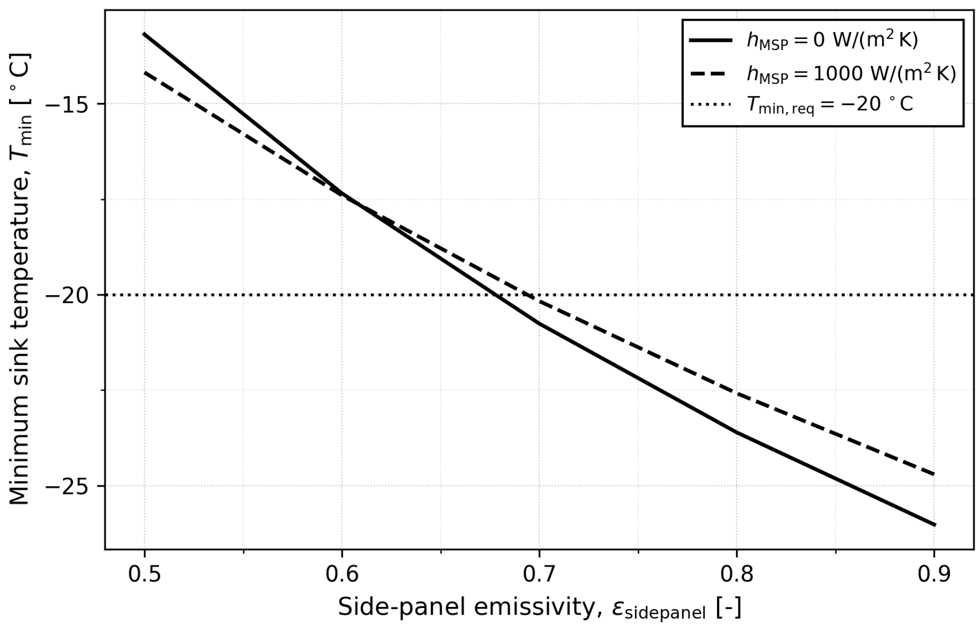

Figure 4.8: Minimum structure temperature in the cold transient case at safe-mode dissipation as a function of the side-panel emissivity, for the two bounding values of h_{MSP} (0 and 1000 $\mathrm{W\,m^{-2}\,K^{-1}}$). The −20 °C cold-case bus-level reference is shown as a visual guide. The actual cold-case compliance check is performed at subsystem level in Section 4.3.4.

The cold-case result is shown in Figure 4.8. The two h_{MSP} bounding curves nearly collapse, indicating that h_{MSP} has only a marginal effect in the cold case. The −20 °C bus-level line is shown as a reference, as the cold-case requirement is defined at subsystem level and is verified directly against the subsystem operating limits in Section 4.3.4. The minimum bus structure temperature at the selected $\varepsilon_{\mathrm{sp}} = 0.90$ lies a few kelvin below the bus-level reference. The selection is driven by the hot case because the hot-case margin is the limiting one, and because heating a subsystem locally is in general easier than cooling the structure globally.

4.3.3. Maximum Allowable Payload Duty Cycle

The screening of Section 4.2 did not identify the payload duty cycle η_{pld} as a strong contributor to the hot-case structure temperature. The duty cycle is nevertheless the only thermal lever available in flight once the bus configuration is fixed. The question addressed here is what maximum duty cycle the hot-case design target permits for a given level of passive thermal design.

To place this result in the context of the broader design trade-off, the duty cycle sweep is first performed for three heat-shield design configurations spanning different levels of thermal-design intervention. Design D1 uses the PEEK heat shield with a heat-shield-to-structure contact conductance of $h_{\text{HS}} = 1000\ \text{W}\,\text{m}^{-2}\,\text{K}^{-1}$, representative of a bolted interface. Design D2 retains the PEEK heat shield but assumes the interface to be fully isolated from the primary structure ($h_{\text{HS}} = 0\ \text{W}\,\text{m}^{-2}\,\text{K}^{-1}$), representing the idealised upper bound on achievable thermal isolation. Design D3 replaces the PEEK heat shield with an aluminium plate on a plain bolted interface, representing a fully commercial-off-the-shelf (COTS) platform without dedicated thermal design at the ram face. The three configurations therefore span from the most extensive thermal design and most complex interface (D2), through a more realistic intermediate configuration (D1), to the simplest COTS configuration (D3). All three are evaluated at the selected side-panel emissivity $\varepsilon_{\text{sp}} = 0.90$, with the payload duty cycle varied between 0 % and 50 % of the payload acquisition-mode dissipation. The sweep is run at the minimum-perigee operating point $h_{\text{p}} = 130$ km, the limiting case of the mission envelope.

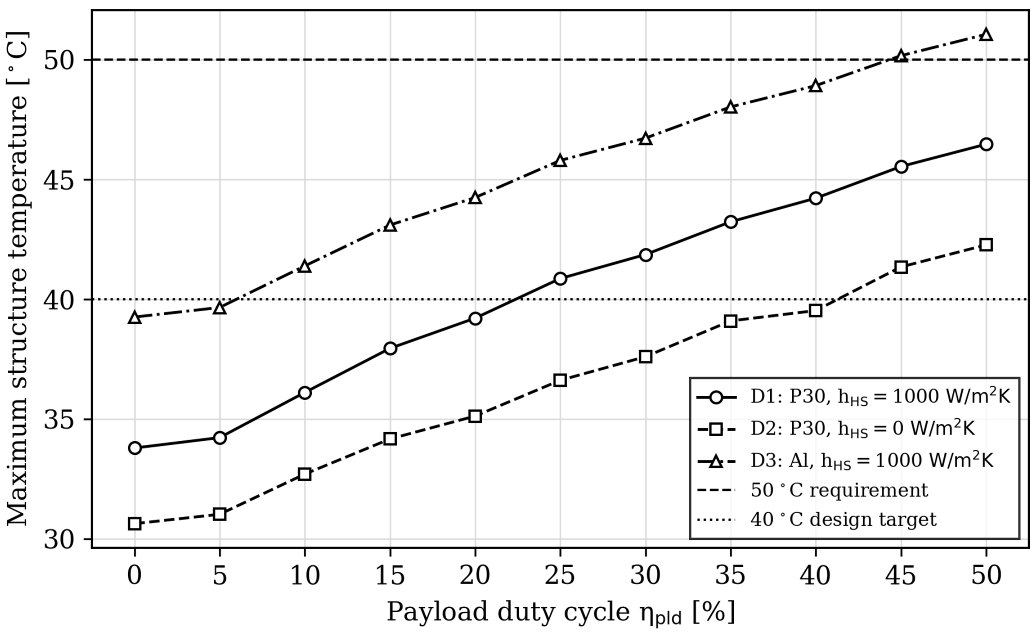

Figure 4.9: Maximum structure temperature of the hot-transient case at $h_{\text{p}} = 130$ km as a function of the payload duty cycle, for three heat-shield design configurations at $\varepsilon_{\text{sp}} = 0.90$: D1 (PEEK heat shield, $h_{\text{HS}} = 1000\ \text{W}\,\text{m}^{-2}\,\text{K}^{-1}$), D2 (PEEK heat shield, fully isolated, $h_{\text{HS}} = 0\ \text{W}\,\text{m}^{-2}\,\text{K}^{-1}$), and D3 (aluminium panel replacing the heat shield, bolted interface, $h_{\text{HS}} = 1000\ \text{W}\,\text{m}^{-2}\,\text{K}^{-1}$). The 50 °C structure-to-payload interface requirement and the 40 °C design target, including the 10 K uncertainty margin, are shown for reference. D2 performs best, permitting a duty cycle up to approximately 40 % at 130 km perigee. D3, the fully COTS configuration, performs worst, limited to approximately 5 %.

The maximum structure temperature increases monotonically with the duty cycle for all three configurations, with an approximately linear slope in each case. The 40 °C hot-case design target is reached at approximately $\eta_{\text{pld}} = 40\,\%$ for the fully isolated interface D2 and at $\eta_{\text{pld}} = 5\,\%$ for the COTS configuration D3. The intermediate configuration D1 lies between both results, at $\eta_{\text{pld}} \approx 20\,\%$.

The large difference in allowable duty cycle between the aluminium-plate configuration D3 and the PEEK-based configurations D1 and D2 can be traced to the difference in thermal conductivity between the two materials. CF30 PEEK has a thermal conductivity of approximately 0.3 $W\,m^{-1}\,K^{-1}$, compared to approximately 130 $W\,m^{-1}\,K^{-1}$ for the aluminium side panels, while the two materials have comparable density-specific-heat-capacity products. The low conductivity of PEEK limits the rate at which aerothermal heat absorbed at the heat shield reaches the structure, whereas the aluminium plate conducts this heat into the structure with little resistance. Limited operation with a commercial aluminium panel remains possible at higher perigee altitudes, but more extensive payload operation and operation at lower perigee altitudes require a heat-shield material combining low thermal conductivity with high thermal capacity, such as PEEK. This is particularly relevant for operations at or below the 130 km minimum-perigee altitude, where the aerothermal heat flux increases sharply.

A fully isolated heat-shield interface, as assumed in D2, is not physically realisable. Any mechanical mounting of the heat shield to the primary structure introduces a finite contact conductance. Configuration D1, with a contact conductance of $h_{HS} = 1000\ W\,m^{-2}\,K^{-1}$, is therefore selected as the baseline for the remainder of the analysis, representing an achievable value for the heat-shield interface. The maximum payload duty cycle compatible with the deterministic hot-case design target at minimum perigee for this configuration is $\eta_{pld} \approx 20\,\%$.

The same sweep is repeated for configuration D1 at perigee altitudes between 140 km and 230 km in 10 km increments, holding ε_{sp} and h_{HS} fixed at the D1 values, to quantify how the allowable duty cycle evolves across the mission envelope.

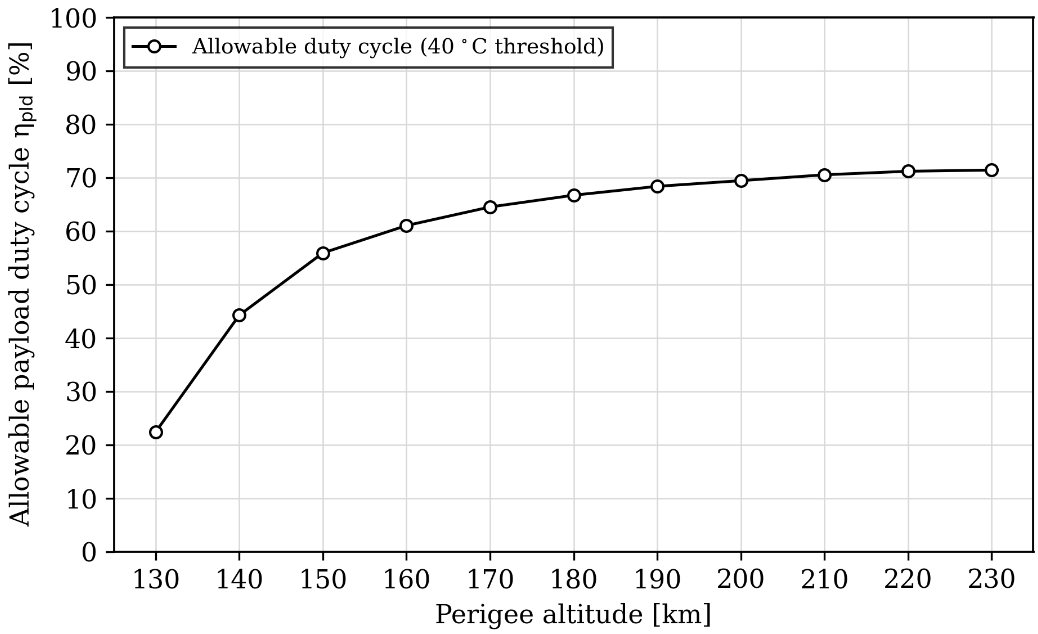

Figure 4.10: Maximum allowable payload duty cycle as a function of perigee altitude for the selected configuration D1, defined as the duty cycle at which the hot-case structure temperature reaches the 40 °C design target, at $\varepsilon_{sp} = 0.90$ and $h_{HS} = 1000\ W\,m^{-2}\,K^{-1}$. Each point is obtained by linear interpolation of a duty-cycle sweep performed at that altitude. The 130 km point corresponds to the D1 curve of Figure 4.9.

The constraint relaxes sharply with increasing perigee altitude, from $\eta_{pld} = 20.0\,\%$ at 130 km to $\eta_{pld} = 70.5\,\%$ at 210 km (Figure 4.10). The resulting thermal limit remains far below the approximately 97 % payload-acquisition duty cycle allowed by the dawn-dusk power budget of Section 1.3. The allowable duty cycle at minimum perigee is therefore set by the thermal constraint. The altitude-to-altitude increments shrink monotonically, from 22 percentage points between 130 and 140 km to 1.1 percentage points between 200 and 210 km, with the curve flattening from approximately 190-200 km onward. This behaviour is consistent with the reduction in atmospheric density, and hence in free-molecular aerothermal heat flux, with increasing altitude. As the aerothermal contribution to the hot-case load diminishes, the allowable duty cycle approaches an asymptotic value set primarily by the solar and internal heat loads, and further altitude increases yield rapidly diminishing thermal margin.

The single-scalar duty cycle adopted here treats the payload group as either fully active or fully inactive at a fixed orbit fraction. The actual operational profile need not follow this approximation. Two refinements are available without changing the hardware. The first is staggered scheduling, in which individual payload subsystems are operated at independent duty cycles such that the orbit-averaged heat load on the bus is reduced for the same scientific output. The second is temperature-aware operation, in which each payload is switched on or off as a function of its own measured temperature, so that the duty cycle adapts to the instantaneous thermal margin instead of following a fixed open-loop schedule. Both refinements would relax the $\eta_{\mathrm{pld}} \approx 20\,\%$ constraint of configuration D1 at minimum perigee. Their development lies outside the scope of the present analysis and is recommended as future work.

4.3.4. Selected Configuration

The selected passive thermal configuration combines the three parameter sweeps of the preceding subsections. The configuration is summarised in Table 4.2 and is the deterministic reference point for the hot- and cold-case verification below and the input to the reliability assessment of Section 4.4.

Table 4.2: Selected passive thermal configuration of the spacecraft bus.

Parameter	Value	Source
Side-panel infrared emissivity, $\varepsilon_{\mathrm{sp}}$	0.90	Section 4.3.2
Candidate side-panel coating	S13GP:6N/LO-1 (IITRI)	This subsection
Heat-shield interface conductance, h_{HS}	$1000\ \mathrm{W\,m^{-2}\,K^{-1}}$	Section 4.3.1, worst-case bound
Body-mounted solar panel conductance, h_{MSP}	$1000\ \mathrm{W\,m^{-2}\,K^{-1}}$	Section 4.3.2, worst-case bound
Maximum payload duty cycle at 130 km, η_{pld}	20 %	Section 4.3.3

The side-panel emissivity is committed deterministically at $\varepsilon_{\mathrm{sp}} = 0.90$, the highest value in the screened range. A candidate radiator coating that realises this emissivity with substantial space heritage is the S13GP:6N/LO-1 silicate-bonded matt-white paint produced by IITRI [83]. It has a hemispherical emittance of 0.90 and a solar absorptance of 0.18 measured per ECSS-Q-70-09. The coating satisfies the ECSS-Q-70-02 outgassing limits (TML = 0.5 %, RML = 0.15 %, CVCM between 0.01 % and 0.04 %) and the ECSS-Q-70-04 thermal-cycling test. The qualification of the coating against the VLEO atomic-oxygen environment is outside the scope of the present analysis. The time-invariant surface-properties standing assumption of Section 3.4 applies for the six-month mission window considered here. The longer-term atomic-oxygen erosion of the silicate binder is recommended as forward work before the coating selection is finalised for an extended-duration mission.

The two interface conductances h_{HS} and h_{MSP} are not committed to deterministic values at the present design stage. Both depend on mounting-hardware implementations that are not fixed at this level of design. For the deterministic sizing, both are taken at the upper bound of the screened interval $[0,\ 1000]\ \mathrm{W\,m^{-2}\,K^{-1}}$. This is the worst-case choice in the hot case, since higher conductance admits more external heat into the structure. Both conductances are retained as epistemic uncertain inputs in the reliability assessment of Section 4.4, where they are relaxed back to their full screened intervals.

Hot-Case Heat Distribution at the Selected Configuration

The hot-case maximum bus structure temperature at the selected configuration reaches the 40 °C design target at the $\eta_{\mathrm{pld}} = 20\,\%$ duty cycle established in Section 4.3.3. The 10 °C margin to the 50 °C structure-to-payload interface requirement is reserved for the dispersion produced by the uncertain inputs. The components mounted on the bus remain in the range of 21 to 57 °C across the configuration. The cooler end of this range corresponds to components on the self-shadowed side of the bus. The warmer end corresponds to components near the body-mounted solar panel and the front face, with each component's internal dissipation in the payload acquisition mode superimposed on the structural gradient. The COMSOL temperature distribution at the selected configuration can be visualised in Figure 4.11. The component temperatures are obtained from the simplified subsystem representations of the present thermal model. Refinement with detailed subsystem thermal models is recommended as future work.

Figure 4.11: Final-design hot-case component temperature distribution at $h_{\mathrm{p}} = 130$ km and $\eta_{\mathrm{pld}} = 20\,\%$, reaching the +40 °C bus structure design target with the 10 °C margin to the 50 °C structure-to-payload interface requirement reserved for input dispersion. Component temperatures range from 21 to 57 °C: the cooler end corresponds to components on the self-shadowed side of the bus, and the warmer end to components near the body-mounted solar panel and front face, where internal dissipation in payload acquisition mode is superimposed on the structural gradient.

Cold-Case Subsystem-Level Verification

The cold-case minimum bus structure temperature at the selected configuration lies approximately 5 °C below the −20 °C bus-level design target. The cold-relevant subsystems themselves remain within their individual operational limits. They are coupled to the structure with finite thermal resistance and dissipate a small but non-zero power in the safe mode. The most temperature-sensitive subsystems are the on-board computer and the ADCS controller, both with a chip-level lower limit of −30 °C. They reach a minimum of approximately −19 °C in the cold transient case. The remaining cold-relevant subsystems reach a minimum of approximately −26 °C against more permissive operational limits. The cold-case subsystem-level requirement is therefore satisfied by the selected configuration. The subsystem temperatures reported here are obtained from the simplified subsystem representations of the present thermal model, in which each subsystem is treated as a lumped capacity with a single conductive coupling to its mounting interface. The COMSOL results can be visualized in Figure 4.12.

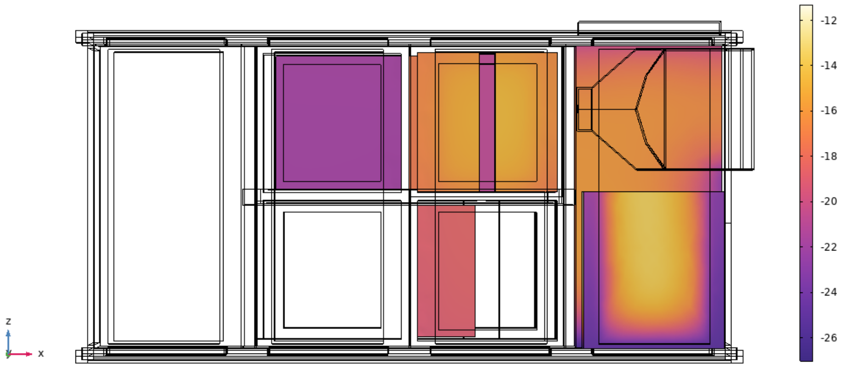

Figure 4.12: Final-design cold-case component temperature distribution at the 550 km LTAN 12:00 reference orbit in safe mode, satisfying the cold-case subsystem-level requirement. The minimum bus structure temperature lies approximately 5 °C below the −20 °C bus-level design target. The most temperature-sensitive subsystems, the on-board computer and ADCS controller (chip-level limit −30 °C), reach a minimum of approximately −19 °C; the remaining cold-relevant subsystems reach approximately −26 °C against more permissive operational limits.

Propulsion System

The propulsion system is omitted from the present thermal analysis. The thruster is assumed to be thermally insulated from the bus structure. This is consistent with the approach taken by Suresha et al. [84] in their thermal sensitivity analysis of a 22 N bipropellant thruster, where multi-layer insulation was used to protect the spacecraft structure from the heat generated during firing. A similar analysis is required for the thruster used in the present mission. The thruster hardware is not specified in sufficient detail at the present design stage to perform that analysis. The propulsion-thermal interaction is therefore identified as an open point of the thermal design and is recommended for future work once the thruster selection and its interface to the structure are fixed.

4.4. Reliability Assessment

The deterministic analysis of Section 4.3 established that the selected passive configuration meets the hot-case structure-to-payload interface requirement with the design margin of Subsection 1.3.2, conditional on point estimates of every input. The reliability assessment relaxes that assumption by propagating the dominant uncertain inputs through the selected configuration and quantifying the resulting distribution of the maximum structure temperature, from which the probability of meeting the requirement is read directly. It thereby converts the deterministic margin allocation of Subsection 1.3.2 into a quantitative compliance statement on the selected design.

4.4.1. Quantity of Interest and Compliance Criterion

The quantity of interest is the maximum structure-to-payload interface temperature over the hot transient case,

$$Q(\mathbf{x}) \;=\; \max_{t\in[0,T]} T_{\mathrm{struc}}(t;\mathbf{x})\,, \tag{4.3}$$

where $\mathbf{x}$ is the vector of uncertain inputs and T the simulated mission interval. Compliance is assessed against the interface temperature requirement $T_{\mathrm{req}} = 323.15$ K (50 °C) of Subsection 1.3.2. The margin-adjusted design target $T_{\mathrm{tgt}} = T_{\mathrm{req}} - \Delta T_{\mathrm{margin}} = 313.15$ K (40 °C), used as the deterministic sizing point, is retained as a reference. The exceedance probability is

$$p_{\mathrm{f}} \;=\; P\big[Q(\mathbf{x}) > T_{\mathrm{req}}\big]\,, \tag{4.4}$$

estimated directly from the propagated distribution of Q obtained in Section 4.4.4, where the configuration is compliant when p_{f} remains below the accepted target probability.

4.4.2. Uncertain Input Parameters

The uncertain inputs are selected from the Morris screening of Section 4.2. Six parameters are retained (Table 4.3): the heat-shield-to-structure conductance, the side-panel infrared emissivity and solar absorptivity, the bolted-interface contact factor, and the amplitude and frequency of the yaw disturbance. Together they span the dominant design, interface, and environmental contributions to the maximum structure temperature while keeping the propagation tractable. Internal power dissipation is not retained, because the screening showed a small parameter-level effect over its admissible range and because the payload duty cycle is fixed at the operating point selected in Section 4.3.

Inputs for which a characterised distribution is available are assigned that distribution: the side-panel emissivity is represented by a truncated normal with its datasheet mean and standard deviation, reflecting the documented unit-to-unit dispersion of production radiator coatings. Inputs known only by bounds are treated as uniform over their admissible range, the least-informative choice consistent with the available knowledge [85].

Table 4.3: Uncertain input parameters retained for the reliability assessment of the selected passive thermal configuration.

Parameter	Symbol	Representation	Source
Bolted-interface contact factor	f_c	$\mathcal{U}[0.58, 0.86]$	Refs. [86, 87]
Yaw disturbance frequency [Hz]	f_β	$\mathcal{U}[0, 0.05]$	Analysis (Sec. 2.4)
Yaw disturbance amplitude [deg]	θ_{dist}	$\mathcal{U}[0, 1.23]$	Analysis (Sec. 2.4)
Side-panel emissivity	ε_{sp}	$\mathcal{N}_{\text{trunc}}(0.9, 0.03)$	Datasheet
Side-panel solar absorptivity	α_{sp}	$\mathcal{U}[0.16, 0.20]$	Datasheet
Heat-shield contact conductance [$\mathrm{W\,m^{-2}\,K^{-1}}$]	h_{HS}	$\mathcal{U}[0, 1000]$	Analysis

4.4.3. Surrogate-Based Uncertainty Propagation

Evaluating the transient COMSOL thermal model at the sample sizes required to resolve the QoI distribution is computationally expensive. The propagation is therefore performed with the COMSOL Uncertainty Quantification module [88] on a surrogate model fitted to a limited set of COMSOL runs. A Gaussian process is used as the surrogate [89]. Its training points are placed by Latin hypercube sampling [90], and the surrogate is refined adaptively until a relative tolerance on the QoI statistics is met. The converged surrogate is then evaluated by Monte Carlo over the input distributions of Table 4.3 to approximate the probability density function of Q, which is estimated from the samples by kernel density estimation (KDE), a smoothing of the sample histogram into a continuous density [91]. The module additionally reports a confidence-interval table giving the mean, standard deviation, extrema, and the values of Q at prescribed cumulative distribution function (CDF) levels. These CDF levels provide a first-order estimate of the probability that Q exceeds a given threshold [88], which is how the exceedance probability of Eq. (4.4) is obtained.

A dedicated reliability analysis, such as efficient global reliability analysis (EGRA), would refine the exceedance estimate by concentrating additional model evaluations near the limit state $Q = T_{\text{req}}$ and is the appropriate tool when p_{f} is small [92, 93]. It is not pursued here because, as shown in Section 4.4.4, the propagated distribution lies entirely well below the requirement, leaving no probability mass near the limit state for the adaptive scheme to resolve. Therefore the uncertainty-propagation distribution alone establishes compliance with margin.

4.4.4. Results

The kernel density estimate of the maximum bus structure temperature is shown in Figure 4.13. The response is unimodal and approximately Gaussian, centred at 312.8 K (39.7 °C) with a standard deviation of 0.85 K and a sampled support of 309.3 to 315.6 K. Across the adaptive refinement iterations the predicted mean varies by less than 0.02 K and the standard deviation by less than 0.04 K. The surrogate has therefore converged.

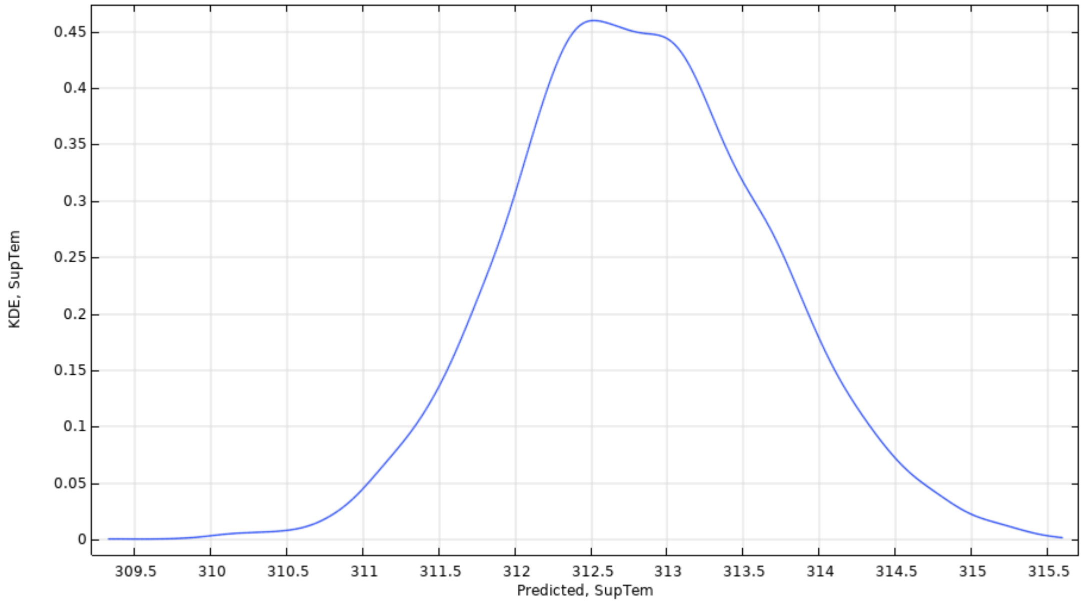

Figure 4.13: Kernel density estimate of the maximum bus structure temperature predicted by the surrogate for the selected configuration, centred at 312.8 K (39.7 °C) with a standard deviation of 0.85 K. The horizontal axis is in Kelvin; the 40 °C margin-adjusted design target corresponds to 313.15 K.

The entire sampled support lies more than 7 °C below the 50 °C (323.15 K) interface requirement, so the exceedance probability is below the resolution of the assessment, $p_f \approx 0$. The 10 °C margin reserved in the deterministic sizing of Section 4.3.3 is preserved almost in full, the response standard deviation amounting to only about 8.5 % of it. The propagated mean coincides with the 40 °C margin-adjusted design target (313.15 K) to within 0.4 °C, and that target is met with a probability of approximately 65 % (standardised distance $z \approx 0.4$). The interface requirement is therefore satisfied with substantial margin, and the deterministic design target is corroborated probabilistically as the centre of the propagated response.

The residual dispersion is the expected consequence of propagating the input uncertainties through the surrogate and, at 0.85 °C, is small relative to the 10 °C margin, so it does not threaten compliance. It reflects the combined effect of the retained inputs, dominated by those known only as bounds, whose admissible ranges are wide. Reducing it is therefore a matter of improving the state of knowledge rather than redesigning the configuration. Experimental correlation of the interface conductances and a higher-fidelity, interface-resolved analysis would tighten these ranges, and a more conservative design would fix the bolted interfaces at their worst-case rather than nominal contact factor. A quantitative attribution of the dispersion to individual inputs, and a differentiated interface-level torque selection that follows from it, are left as forward work.

5 Conclusion

This thesis set out to establish whether a predominantly commercial-off-the-shelf (COTS) 16U CubeSat can be operated thermally in very low Earth orbit (VLEO) down to a minimum perigee altitude of 130 km, guided by two questions: how the environmental conditions specific to VLEO and the spacecraft's own configuration determine its thermal response at low perigee, and how these effects can be captured in a physics-consistent thermal-analysis framework suitable for design use; and what passive thermal-control configuration allows the platform to meet its structure-to-payload interface temperature requirement at the minimum-perigee operating point, together with the probability of violating that requirement once the dominant environmental and modelling uncertainties are propagated through the selected design. The assessment was simulation-based. The VLEO environment was characterised and suitable modelling approaches were selected for the atmospheric, aerothermal, aerodynamic, and Earth-radiation inputs, and these inputs were propagated through a reduced-order LTspice thermal network and a finite-element COMSOL model. Sensitivity screening was then used to identify the relevant thermal design levers, after which the selected passive configuration was verified by uncertainty propagation.

The environmental modelling showed that analytical free-molecular methods are suitable for rapid design iteration in the considered VLEO regime. The Caruso aerothermal model and the Sentman aerodynamic model reproduced the DSMC reference trends with sufficient accuracy for spacecraft-level thermal design, while reducing the computational cost from days to milliseconds per evaluation. For the Earth-radiation environment, the comparison showed that CERES-based albedo and infrared models are preferable when the temperature uncertainty associated with planetary radiation should remain below the order of 10 °C. The Sasaki zonal-harmonics model was therefore selected for the final thermal cases.

From the thermal-design perspective, the analysis demonstrated that passive operation at a 130 km perigee is feasible for the considered platform, provided that the ram-facing aerothermal load is intercepted by a heat shield and the bus has sufficient radiative heat-rejection capability. The baseline configuration was not compliant in the hot case, even without internal dissipation, which identified bus heat rejection rather than payload power as the limiting mechanism. The selected design therefore combines high-emissivity side panels with an infrared emissivity of $\varepsilon_{\mathrm{sp}} = 0.90$ with a PEEK heat-shield. With this configuration, the deterministic hot-case design permits a payload duty cycle of 20 % at the minimum perigee, with larger duty cycles feasible at higher perigee altitudes.

The uncertainty propagation confirmed the deterministic sizing. The propagated maximum structure temperature was centred at 39.7 °C with a standard deviation of 0.85 °C, and the sampled response remained more than 7 °C below the 50 °C structure-to-payload interface requirement. No requirement exceedance was observed within the propagated uncertainty set, so the exceedance probability is below the resolution of the assessment. The residual dispersion is small compared with the reserved 10 °C margin and is governed mainly by interface and environmental parameters that are currently known only through bounds rather than characterised distributions.

Future work should focus on reducing these remaining uncertainties and refining the subsystem-level representation. Subsystem temperatures are currently obtained from a simplified representation in which each subsystem is treated as a lumped capacity with a single conductive coupling to its mounting interface. A more detailed thermal model of the individual subsystems is recommended, particularly for the propulsion system, whose thruster is assumed thermally insulated from the bus by analogy with a comparable bipropellant thruster in the literature but has not yet been specified in enough detail to confirm this. The heat-shield and body-mounted-solar-panel conductances are likewise only bounded rather than characterised. Representative thermal-vacuum testing or correlation, once a mounting concept is fixed, would narrow these bounds and reduce their leading contribution to the 0.85 °C standard deviation of the propagated response. More broadly, the aerothermal and aerodynamic models are validated against DSMC data at 140 and 180 km, whereas the 130 km minimum perigee lies close to the transitional-flow boundary where the largest deviation is expected; dedicated thermal-vacuum correlation, or ultimately flight telemetry, would replace this bounds-only treatment with validated design margins at the lowest operating altitude.

Two further effects were identified but not carried through the present model. Atomic-oxygen-driven drift in the optical properties of exposed surfaces was excluded by assuming time-invariant properties over a six-month mission window. As the selected configuration is sized directly on the side-panel emissivity, end-of-life properties should be incorporated before the coating is finalised for a longer mission. Separately, the stability analysis assumes perfect geometric and mass symmetry, so that small manufacturing offsets between the centre of gravity and the panel geometry, and more general tumbling beyond this small-perturbation regime, are not reflected in the assumed thermal loading footprint. A dedicated analysis is recommended to confirm that these effects do not erode the 7 °C margin to the interface requirement.

Bibliography

[1] Crisp, N., Roberts, P., Livadiotti, S., Oiko, V., Edmondson, S., Haigh, S., Huyton, C., Sinpetru, L., Smith, K., Worrall, S., Becedas, J., Domínguez, R., González, D., Hanessian, V., Mølgaard, A., Nielsen, J., Bisgaard, M., Chan, Y.-A., Fasoulas, S., Herdrich, G., Romano, F., Traub, C., García-Almiñana, D., Rodríguez-Donaire, S., Sureda, M., Kataria, D., Outlaw, R., Belkouchi, B., Conte, A., Perez, J., Villain, R., Heißerer, B., and Schwalber, A., "The benefits of very low earth orbit for earth observation missions," *Progress in Aerospace Sciences*, Vol. 117, 2020, p. 100619. https://doi.org/10.1016/j.paerosci.2020.100619.

[2] Huang, J., Chang, L., Dong, B., Wang, J., and Huang, H., "Variable Switching System for Heat Protection and Dissipation of Ultra-LEO Satellites Based on LHP Coupled with TEC," *Aerospace*, Vol. 11, No. 7, 2024, p. 539. https://doi.org/10.3390/aerospace11070539.

[3] Llop, J. V., "Spacecraft Flight in the Atmosphere," Ph.D. thesis, School of Engineering, Cranfield University, 2014. Supervisor: Dr. Peter C.E. Roberts. © Cranfield University, 2014. All rights reserved.

[4] European Space Agency, "Earth Explorer 10 Candidate Mission Daedalus: Report for Assessment," Report for Assessment ESA-EOPSM-DAED-RP-3793, European Space Agency, Noordwijk, The Netherlands, Nov. 2020. Issue/Revision 1.0, issued 13 November 2020.

[5] Corpino, S., Caldera, M., Nichele, F., Masoero, M., and Viola, N., "Thermal design and analysis of a nanosatellite in low earth orbit," *Acta Astronautica*, Vol. 115, 2015, pp. 247–261. https://doi.org/10.1016/j.actaastro.2015.05.012.

[6] Floberghagen, R., Fehringer, M., Lamarre, D., Muzi, D., Frommknecht, B., Steiger, C., Piñeiro, J., and da Costa, A., "Mission Design, Operation and Exploitation of the Gravity Field and Steady-State Ocean Circulation Explorer Mission," *Journal of Geodesy*, Vol. 85, 2011, pp. 749–758. https://doi.org/10.1007/s00190-011-0498-3.

[7] Kimoto, Y., Tsuchiya, Y., Miyazaki, E., Goto, A., Yukumatsu, K., and Imamura, S., "Analysis of upper atmospheric effects on material per onboard atomic oxygen monitor system of SLATS," *Frontiers in Space Technologies*, Vol. 3, 2022, p. 891753. https://doi.org/10.3389/frspt.2022.891753.

[8] Caruso, P. S., Jr., and Naegeli, C. R., "Low-Perigee Aerodynamic Heating During Orbital Flight of an Atmosphere Explorer," NASA Technical Note NASA-TN-D-8308, NASA Goddard Space Flight Center, Greenbelt, MD, United States, Sep. 1976. Public domain. Accession Number: 76N32505.

[9] Adam, F., Villanueva Aguado, M., Cravero, M., Dekker, J., Iannone, A., Jorissen, M., Leto, G., Rucita, Y., Venth, S., and van den Wijngaart, J., "RAMSES: RAM Satellite for Earth Sensing," Final Report DSE08 Spring 2019, FR-1.1, Delft University of Technology, Delft, The Netherlands, Jul. 2019. Project duration: 23 April 2019-4 July 2019.

[10] Canosa Ybarra, D., Janisch, K.-I., Kalis, N., Lentschig, D., Lopez Rivera, A., Manieri, M., Regnery, K., and van de Wal, T., "Sigma Sat Final Review: TU Delft AI Satellite," Final review report, Delft University of Technology, Department of Space Systems Engineering, 2025. DSE 01 Student Project.

[11] Paugnat, H., Gruet, M., Bizalion, H., Courson, R., Darsel, V., Douet, B., Ferrer, H., Fis, Y., Girault, F., and Hart-Davis, L., "IonSat: a Student Nanosat With An Iodine Thruster in Very Low Earth Orbit," *Proceedings of 71st International Astronautical Congress, The CyberSpace Edition*, 2020, pp. 12–14.

[12] Selvadurai, S., "Thermal Control Techniques for Infrared Imaging on Nanosatellites," Ph.D. thesis, Nanyang Technological University, School of Electrical & Electronic Engineering, 2023.

[13] Marca, T. A. L., Graziano, M. D., Grassi, M., Iannascoli, L., Cavicchiolo, L., Capanni, A., Duzzi, M., Capuano, G. M., Fortunato, V., Cardenio, C., Leccese, G., Rinaldi, M., Clemente, M. D., Luciani, R., Natalucci, S., and Pulcino, V., "Mission and System Design for the EarthNext Cubesat VLEO Mission," *CEAS Space Journal*, 2025. https://doi.org/10.1007/s12567-025-00663-2.

[14] Bailet, G., Sakraker, I., Scholz, T., and Muylaert, J., "Qubesat for Aerothermodynamic Research and Measurement on AblatioN," *9th International Planetary Probe Workshop*, Toulouse, France, 2012, pp. 1–9. HAL Id: hal-02081892.

[15] Martella, D., Marugán Rubio, G., and Garcia, R., "EntrySat Thermal Control Subsystem," Final report, Institut Supérieur de l'Aéronautique et de l'Espace (ISAE-SUPAERO), Université de Toulouse, Toulouse, France, 2017.

[16] Galla, D., Klinkner, S., Herdrich, G., Kaiser, C., Kuhm, H., and Fischer, H., "The Detailed Design of the SOURCE Satellite for Demise Investigation," *33rd International Symposium on Space Technology and Science*, Beppu, Oita, Japan, 2022, p. 1. Presented 1 March 2022.

[17] Bailet, B., Asma, C. O., Muylaert, J., and Magin, T., "Feasibility Analysis and Preliminary Design of an Atmospheric Re-Entry CubeSat Demonstrator," *Proceedings of the 7th European Symposium on Aerothermodynamics*, ESA, Brugge, Belgium, 2011, p. 109.

[18] Pikus, A., Berger, A., Bolliger, M., Parkos, D., and Alexeenko, A., "DSMC Aerothermal Study for 3U CubeSat Probes in LEO," *47th AIAA Thermophysics Conference*, Denver, Colorado, 2017, pp. 1–11. https://doi.org/10.2514/6.2017-4027, aIAA Paper 2017-4027.

[19] Picone, J. M., Hedin, A. E., Drob, D. P., and Aikin, A. C., "NRLMSISE-00 empirical model of the atmosphere: Statistical comparisons and scientific issues," *Journal of Geophysical Research: Space Physics*, Vol. 107, No. A12, 2002, p. 1468. https://doi.org/10.1029/2002JA009430.

[20] International Organization for Standardization (ISO), "ISO 14222 space environment (natural and artificial)." , 2013.

[21] Bruinsma, S., "The DTM-2013 Thermosphere Model," *Journal of Space Weather and Space Climate*, Vol. 5, 2015, p. A1. https://doi.org/10.1051/swsc/2015001.

[22] Bowman, B., Tobiska, W. K., Marcos, F., Huang, C., Lin, C., and Burke, W., "A New Empirical Thermospheric Density Model JB2008 Using New Solar and Geomagnetic Indices," *AIAA/AAS Astrodynamics Specialist Conference and Exhibit*, American Institute of Aeronautics and Astronautics, 2008, pp. AIAA 2008–6438. https://doi.org/10.2514/6.2008-6438.

[23] Sentman, L. H., "Free-Molecule Flow Theory and Its Application to the Determination of Aerodynamic Forces," *ARS Journal*, Vol. 31, No. 12, 1961, pp. 1717–1726.

[24] Schaaf, S. A., and Chambre, P. L., *Flow of Rarefied Gases*, Princeton University Press, Princeton, NJ, 1961.

[25] Vallado, D. A., and Finkleman, D., "A Critical Assessment of Satellite Drag and Atmospheric Density Modeling," *Acta Astronautica*, Vol. 95, 2014, pp. 141–165. https://doi.org/10.1016/j.actaastro.2013.10.005.

[26] Sethuraman, P., Ram, V. R. P., Yang, Y., You, H., and Kim, J. G., "An improved analytical method for rarefied aerothermodynamics on a complex geometry in free-molecular flows," *Aerospace Science and Technology*, Vol. 155, 2024, p. 109644. https://doi.org/10.1016/j.ast.2024.109644.

[27] Bird, G. A., *Molecular Gas Dynamics and the Direct Simulation of Gas Flows*, Oxford Science Publications, Oxford, UK, 1994.

[28] Boyd, I. D., "Computation of Hypersonic Flows Using the Direct Simulation Monte Carlo Method," *Journal of Spacecraft and Rockets*, Vol. 52, No. 1, 2015, pp. 38–53. https://doi.org/10.2514/1.A32999.

[29] Hsieh, C., Pan, C., and Lo, M., "Determining Aerothermodynamic Effects on Very-Low-Earth-Orbit Satellites Using Direct Simulation Monte Carlo Method," *International Journal of Computational Fluid Dynamics*, Vol. 35, No. 5, 2021, pp. 388–411.

[30] Kaiser, C. F., "Numerical Simulation of the Rarefied, High-Enthalpy Gas Flow During Atmospheric Entry for SOURCE, a 3U CubeSat," Master's thesis, University of Stuttgart, 2020.

[31] Jin, X., Huang, F., Cheng, X., Wang, Q., and Wang, B., "Monte Carlo Simulation for Aerodynamic Coefficients of Satellites in Low-Earth Orbit," *Acta Astronautica*, Vol. 160, 2019, pp. 222–229. https://doi.org/10.1016/j.actaastro.2019.04.012.

[32] Walsh, J., Berthoud, L., and Allen, C., "Drag Reduction through Shape Optimisation for Satellites in Very Low Earth Orbit," *Acta Astronautica*, Vol. 179, 2021, pp. 105–121. https://doi.org/10.1016/j.actaastro.2020.09.018.

[33] Mehta, P. M., Minisci, E., Vasile, M., Walker, A. C., and Brown, M., "An Open-Source Hypersonic Aerodynamic and Aerothermodynamic Modeling Tool," *20th AIAA International Space Planes and Hypersonic Systems and Technologies Conference*, 2015, pp. 1–8. https://doi.org/10.2514/6.2015-3089, URL https://doi.org/10.2514/6.2015-3089.

[34] Jiang, Y., Zhang, J., Tian, P., Liang, T., Li, Z., and Wen, D., "Aerodynamic drag analysis and reduction strategy for satellites in Very Low Earth Orbit," *Aerospace Science and Technology*, 2022. https://doi.org/10.1016/j.ast.2022.108077.

[35] Gilmore, D., and Donabedian, M., *Spacecraft Thermal Control Handbook*, Vol. 1, Aerospace Press, El Segundo, CA, 2002.

[36] Knocke, P. C., Ries, J. C., and Tapley, B. D., "Earth Radiation Pressure Effects on Satellites," *AIAA/AAS Astrodynamics Conference*, AIAA, Minneapolis, MN, 1988, pp. 577–587. https://doi.org/10.2514/6.1988-4292.

[37] ECSS Secretariat, *ECSS-E-30 Part 1A: Space Engineering – Mechanical – Part 1: Thermal Control*, European Cooperation for Space Standardization, Noordwijk, The Netherlands, Apr. 2000.

[38] Sasaki, K., "Simplified Earth Infrared and Albedo Coefficient Models for Spacecraft Thermal Analysis Based on the CERES Data Products," *Advances in Space Research*, Vol. 75, 2025, pp. 1597–1615. https://doi.org/10.1016/j.asr.2024.11.017.

[39] Oshima, K., and Oshima, Y., "Analytical approach to the thermal design of spacecraft," ISAS Report 33, Institute of Space and Astronautical Science (ISAS), 1968.

[40] Gadalla, M., "Prediction of temperature variation in a rotating spacecraft in space environment," *Applied Thermal Engineering*, Vol. 25, No. 14, 2005, pp. 2379–2397.

[41] Pérez-Grande, I., Sanz-Andrés, A., Guerra, C., et al., "Analytical study of the thermal behaviour and stability of a small satellite," *Applied Thermal Engineering*, Vol. 29, No. 11, 2009, pp. 2567–2573.

[42] Özişik, M. N., *Finite Difference Methods in Heat Transfer*, 1st ed., CRC Press, 1994. https://doi.org/10.1201/9781315168784.

[43] Malde, C. K., Sharma, G., and Aneesh, A., "Development and Testing of Computational Model for the Thermal Analysis of a CubeSat Nanosatellite," *2021 IEEE Aerospace Conference (50100)*, IEEE, 2021, pp. 1–9. https://doi.org/10.1109/AERO50100.2021.9438469.

[44] Sundu, H., and Doner, N., "Detailed Thermal Design and Control of an Observation Satellite in Low Earth Orbit," *European Mechanical Science*, 2020. https://doi.org/10.26701/ems.730201.

[45] Diaz-Aguado, M. F., Greenbaum, J., Fowler, W. T., et al., "Small Satellite Thermal Design, Test, and Analysis," *Modeling, Simulation, and Verification of Space-based Systems III*, SPIE Proceedings, Vol. 6221, 2006, pp. 74–85. https://doi.org/10.1117/12.667369.

[46] Morsch Filho, E., Seman, L. O., and de Paulo Nicolau, V., "Simulation of a CubeSat with Internal Heat Transfer Using Finite Volume Method," *Applied Thermal Engineering*, Vol. 193, 2021, p. 117039. https://doi.org/10.1016/j.applthermaleng.2021.117039.

[47] Miller, S. W., and Walker, W. Q., "Improvements to a Response Surface Thermal Model for Orion," *Thermal and Fluids Analysis Workshop (TFAWS)*, NASA, Houston, TX, USA, 2011, pp. 23–28.

[48] Shi, R. H., Liu, L., Long, T., et al., "Surrogate Assisted Multidisciplinary Design Optimization for an All-electric GEO Satellite," *Acta Astronautica*, Vol. 138, 2017, pp. 301–317. https://doi.org/10.1016/j.actaastro.2017.05.009.

[49] Xiong, Y., Guo, L., Zhang, Y., et al., "Surrogate Modeling for Spacecraft Thermophysical Models Using Deep Learning," *Neural Computing and Applications*, Vol. 34, No. 19, 2022, pp. 16577–16603. https://doi.org/10.1007/s00521-022-07221-2.

[50] Shen, M., Zhang, L., Sun, Z., Kong, L., Liu, Y., and Xue, Z., "Models, Simulations, and Applications of Small Satellite Thermal Analysis," *Advances in Space Research*, Vol. 74, No. 2, 2024, pp. 836–860. https://doi.org/10.1016/j.asr.2024.04.026.

[51] Sun, D., You, E., Zhang, T., Xu, J., Wang, X., Ren, X., and Tao, W., "A review of thermal contact conductance research of conforming contact surfaces," *International Communications in Heat and Mass Transfer*, Vol. 159, 2024, p. 108065. https://doi.org/10.1016/j.icheatmasstransfer.2024.108065.

[52] Blanchete, N., and Bah, A., "CubeSat thermal analysis: evaluating models for thermal contact conductance," *Heat and Mass Transfer*, Vol. 60, 2024, pp. 1221–1233. https://doi.org/10.1007/s00231-024-03482-z.

[53] Gomez-Sanjuan, A., Pérez-Grande, I., and Sanz-Andrés, A., "Uncertainty calculation for spacecraft thermal models using a generalized SEA method," *Acta Astronautica*, 2018. https://doi.org/10.1016/j.actaastro.2018.05.045.

[54] Sola Chicote, Y., "Design and Performance Assessment of Very Low Earth Orbit Missions within the Framework of Responsive Space." Master's thesis, Universidad Carlos III de Madrid, September 2025. URL https://elib.dlr.de/217845/.

[55] European Cooperation for Space Standardization, "Space Engineering: Thermal Analysis Handbook," Handbook ECSS-E-HB-31-03A, European Cooperation for Space Standardization, Noordwijk, The Netherlands, Nov. 2016. Issued 15 November 2016.

[56] European Cooperation for Space Standardization, "Space Engineering: Space Environment," Standard ECSS-E-ST-10-04C Rev.1, European Cooperation for Space Standardization, Noordwijk, The Netherlands, Jun. 2020. Issued 15 June 2020.

[57] NASA Marshall Space Flight Center, "Solar Cycle Forecast," https://www.nasa.gov/msfcsolar, 2025. Accessed February 2025.

[58] Suggs, R., et al., "The MSFC Solar Activity Future Estimation (MSAFE) Model," Tech. rep., NASA, 2017.

[59] Moe, K., and Moe, M. M., "Gas–Surface Interactions and Satellite Drag Coefficients," *Planetary and Space Science*, Vol. 53, No. 8, 2005, pp. 793–801.

[60] Nizenkov, P., Noeding, P., Konopka, M., and Fasoulas, S., "Verification and Validation of a Parallel 3D Direct Simulation Monte Carlo Solver for Atmospheric Entry Applications," *CEAS Space Journal*, Vol. 9, No. 1, 2017, pp. 127–137. https://doi.org/10.1007/s12567-016-0143-8.

[61] Nizenkov, P., "Numerical Simulation of Rarefied, High-Enthalpy Gas Flows Around Complex Three-Dimensional Bodies During Atmospheric Entry," Ph.D. thesis, University of Stuttgart, Stuttgart, Germany, 2018.

[62] Fujita, K., and Noda, A., "Rarefied Aerodynamics of a Super Low Altitude Test Satellite," *41st AIAA Thermophysics Conference*, 2009, pp. 22–25. https://doi.org/10.2514/6.2009-3606, URL https://doi.org/10.2514/6.2009-3606.

[63] Sutton, E. K., "Normalized Force Coefficients for Satellites with Elongated Shapes," *Journal of Spacecraft and Rockets*, Vol. 46, No. 1, 2009, pp. 112–116. https://doi.org/10.2514/1.40940.

[64] Ramesh, S. S., Elhadidi, B., Khoo, B. C., and Chan, W. L., "Aerodynamic Force and Moment Computations for Very Low Earth Orbit Satellites Using a TPMC Model," *Acta Astronautica*, Vol. 228, 2025, pp. 88–100. https://doi.org/10.1016/j.actaastro.2024.11.034.

[65] Drob, D. P., Emmert, J. T., Meriwether, J. W., Makela, J. J., Doornbos, E. N., Conde, M., Hernandez, G., Noto, J., Zawdie, K. A., McDonald, S. E., Huba, J. D., and Klenzing, J. H., "An Update to the Horizontal Wind Model (HWM): The Quiet Time Thermosphere," *Earth and Space Science*, Vol. 2, No. 7, 2015, pp. 301–319. https://doi.org/10.1002/2014EA000089.

[66] Fu, X., Liang, L., Ma, W., Cui, H., and Zhao, Y., "EHFE Refinement and Extension for Variable Physical Earth Radiation in Spacecraft Uncertainty Thermal Analysis," *Journal of Physics: Conference Series*, Vol. 2746, No. 1, 2024, p. 012018. https://doi.org/10.1088/1742-6596/2746/1/012018.

[67] Wielicki, B. A., Barkstrom, B. R., Harrison, E. F., Lee, R. B., Smith, G. L., and Cooper, J. E., "Clouds and the Earth's Radiant Energy System (CERES): An Earth Observing System Experiment," *Bulletin of the American Meteorological Society*, Vol. 77, No. 5, 1996, pp. 853–868. https://doi.org/10.1175/1520-0477(1996)077<0853:CATERE>2.0.CO;2.

[68] Loeb, N. G., Su, W., Doelling, D. R., Wong, T., Minnis, P., Thomas, S., and Miller, W. F., "Earth's Top-of-Atmosphere Radiation Budget," *Comprehensive Remote Sensing*, edited by S. Liang, Elsevier, Oxford, 2018, pp. 67–84. https://doi.org/10.1016/B978-0-12-409548-9.10367-7.

[69] Shankar, M., Loeb, N. G., Smith, N., Kato, S., Thorsen, T., and Doelling, D. R., "Evaluating the Radiometric Performance of the Clouds and the Earth's Radiant Energy System (CERES) Instruments on Terra and Aqua over 20 Years," *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 61, 2023, pp. 1–11. https://doi.org/10.1109/TGRS.2023.3330398.

[70] Su, W., Corbett, J., Eitzen, Z., and Liang, L., "Next-Generation Angular Distribution Models for Top-of-Atmosphere Radiative Flux Calculation from CERES Instruments: Methodology," *Atmospheric Measurement Techniques*, Vol. 8, No. 2, 2015, pp. 611–632. https://doi.org/10.5194/amt-8-611-2015.

[71] Su, W., Corbett, J., Eitzen, Z., and Liang, L., "Next-Generation Angular Distribution Models for Top-of-Atmosphere Radiative Flux Calculation from CERES Instruments: Validation," *Atmospheric Measurement Techniques*, Vol. 8, No. 8, 2015, pp. 3297–3313. https://doi.org/10.5194/amt-8-3297-2015.

[72] Doelling, D. R., Loeb, N. G., Keyes, D. F., Nordeen, M. L., Morstad, D., and Wielicki, B. A., "Geostationary Enhanced Temporal Interpolation for CERES Flux Products," *Journal of Atmospheric and Oceanic Technology*, Vol. 30, No. 6, 2013, pp. 1072–1090. https://doi.org/10.1175/JTECH-D-12-00136.1.

[73] Doelling, D. R., Sun, M., Nguyen, L. T., Nordeen, M. L., Haney, C. O., Keyes, D. F., and Mlynczak, M. G., "Advances in Geostationary-Derived Longwave Fluxes for the CERES Synoptic (SYN1deg) Product," *Journal of Atmospheric and Oceanic Technology*, Vol. 33, No. 3, 2016, pp. 503–521. https://doi.org/10.1175/JTECH-D-15-0147.1.

[74] NASA Langley Research Center, "CERES and GEO-Enhanced TOA, Within-Atmosphere and Surface Fluxes, Clouds and Aerosols Daily Terra-Aqua Edition 4A," https://ceres.larc.nasa.gov/data/, 2017. https://doi.org/10.5067/Terra+Aqua/CERES/SYN1degDayL3.004A.

[75] ESA Requirements and Standards Division, "Space Engineering: Thermal Design Handbook – Part 4: Conductive Heat Transfer," Handbook ECSS-E-HB-31-01 Part 4A, ECSS Secretariat, ESA-ESTEC, Requirements & Standards Division, Noordwijk, The Netherlands, Dec. 2011.

[76] Walker, T., Xue, S., and Barton, G., "Numerical Determination of Radiative View Factors Using Ray Tracing," *Journal of Heat Transfer-transactions of The Asme*, Vol. 132, 2010, p. 072702. https://doi.org/10.1115/1.4000974.

[77] ESA Requirements and Standards Division, "Space Engineering: Thermal Design Handbook – Part 1: View Factors," Handbook ECSS-E-HB-31-01 Part 1A, ECSS Secretariat, ESA-ESTEC, Requirements & Standards Division, Noordwijk, The Netherlands, Dec. 2011.

[78] COMSOL AB, "Heat Transfer Module User's Guide," , 2023. COMSOL Multiphysics v.6.2, Stockholm, Sweden.

[79] Finckenor, M. M., and Coker, R. F., "Optical Properties of Nanosatellite Hardware," NASA Technical Memorandum NASA/TM-2014-218195, NASA Marshall Space Flight Center, Huntsville, Alabama, June 2014. URL https://ntrs.nasa.gov/citations/20140012786.

[80] Cohen, M. F., and Greenberg, D. P., "The hemi-cube: A radiosity solution for complex environments," *ACM SIGGRAPH Computer Graphics*, Vol. 19, 1985, pp. 31–40.

[81] Morris, M. D., "Factorial Sampling Plans for Preliminary Computational Experiments," *Technometrics*, Vol. 33, No. 2, 1991, pp. 161–174. https://doi.org/10.1080/00401706.1991.10484804.

[82] Campolongo, F., Cariboni, J., and Saltelli, A., "An Effective Screening Design for Sensitivity Analysis of Large Models," *Environmental Modelling & Software*, Vol. 22, No. 10, 2007, pp. 1509–1518. https://doi.org/10.1016/j.envsoft.2006.10.004.

[83] IITRI, "S13GP:6N/LO-1 Silicate Matt-white Paint Datasheet," Illinois Institute of Technology Research Institute, product datasheet, 2024. Measured per ECSS-Q-70-09; outgassing per ECSS-Q-70-02; thermal cycling per ECSS-Q-70-04.

[84] Suresha, S., and Katti, R. A., "Thermal sensitivity analysis of a 22 N bipropellant thruster in a spacecraft," *Proceedings of the Institution of Mechanical Engineers, Part G: Journal of Aerospace Engineering*, Vol. 213, No. 5, 1999, pp. 321–335. https://doi.org/10.1243/0954410991533052.

[85] Der Kiureghian, A., and Ditlevsen, O., "Aleatory or epistemic? Does it matter?" *Struct. Saf.*, Vol. 31, No. 2, 2009, pp. 105–112. https://doi.org/10.1016/j.strusafe.2008.06.020.

[86] Elhady, A., "Design and analysis of a LEO micro-satellite thermal control including thermal contact conductance," *IEEE Aerospace Conference*, 2010, pp. 1–11. https://doi.org/10.1109/aero.2010.5446763.

[87] Ajul, E., and Chanda, S., "Estimation of thermal contact conductance of spacecraft heat sink bolted joints," *Applied Thermal Engineering*, 2023. https://doi.org/10.1016/j.applthermaleng.2023.120078.

[88] COMSOL AB, "Uncertainty Quantification Module User's Guide," , 2024. COMSOL Multiphysics v.6.4, Stockholm, Sweden.

[89] Rasmussen, C. E., and Williams, C. K. I., *Gaussian Processes for Machine Learning*, MIT Press, 2006.

[90] McKay, M. D., Beckman, R. J., and Conover, W. J., "A comparison of three methods for selecting values of input variables in the analysis of output from a computer code," *Technometrics*, Vol. 42, No. 1, 2000, pp. 55–61. https://doi.org/10.1080/00401706.2000.10485979.

[91] Sheather, S. J., "Density estimation," *Statist. Sci.*, Vol. 19, No. 4, 2004, pp. 588–597. https://doi.org/10.1214/088342304000000297.

[92] Bichon, B. J., Eldred, M. S., Swiler, L. P., Mahadevan, S., and McFarland, J. M., "Efficient global reliability analysis for nonlinear implicit performance functions," *AIAA J.*, Vol. 46, No. 10, 2008, pp. 2459–2468. https://doi.org/10.2514/1.34321.

[93] Echard, B., Gayton, N., and Lemaire, M., "AK-MCS: An Active Learning Reliability Method Combining Kriging and Monte Carlo Simulation," *Structural Safety*, Vol. 33, No. 2, 2011, pp. 145–154. https://doi.org/10.1016/j.strusafe.2011.01.002.

[94] Muccino, M., Dell'Agnello, S., Delle Monache, G., Denni, U., Montanari, M., Porcelli, L., Salvatori, L., Tibuzzi, M., Filomena, L., Mauro, L., Mondaini, C., Petrassi, M., Remujo Castro, A., Rubino, L., Traini, M., Chandler, J., Viloresi, P., Vallone, G., Capozziello, S., Vittori, R., Lognonné, P., and Jani, K., "Next-Generation Laser Retro-Reflectors and Application for Lunar Gravitational-Wave Detectors," Presentation slides, 2023. PowerPoint presentation.

A
Spacecraft Data

A.1. Spacecraft Architecture

The reference spacecraft considered in this thesis is a conceptual 16U CubeSat platform intended for operation in the VLEO environment. The platform is based on a modular commercial-off-the-shelf architecture and includes scientific payloads, propulsion hardware, deployable and body-mounted solar arrays, communication units, attitude-control hardware, command and data handling electronics, and thermal protection elements.

For confidentiality reasons, study-specific names, supplier names, institutional affiliations, and detailed component identifiers are not reported. Instead, the spacecraft is described using generic subsystem and component labels. The numerical properties used in the thermal model, such as power dissipation levels and temperature limits, are retained because they define the thermal boundary conditions and design constraints of the conceptual study.

The spacecraft uses a 16U aluminium structural frame as the primary mechanical backbone. A polymer-composite heat shield is mounted on the ram-facing side to reduce the exposure of sensitive front-end components to aerothermal heating and atomic oxygen. The internal configuration is arranged around a set of major equipment groups: payload units, avionics, communication equipment, power-system components, attitude-control hardware, propulsion hardware, and structural/thermal protection elements.

A.1.1. Payload Equipment

The payload equipment is represented by five generic payload units. These include instruments for in-situ environmental measurements, atmospheric composition analysis, optical imaging, navigation support, and flow-intake demonstration. In the thermal model, the payloads are treated primarily through their power dissipation, mounting location, and allowable temperature range.

- **Payload A:** In-situ plasma/environmental monitoring unit.
- **Payload B:** Compact mass-spectrometry or atmospheric-composition measurement unit.
- **Payload C:** Passive intake or flow-guiding demonstrator integrated near the ram-facing side.
- **Payload D:** Compact optical imaging unit.
- **Payload E:** Navigation and timing support unit.

A.1.2. Propulsion System

The spacecraft includes a compact propulsion system used for orbit-control manoeuvres and attitude-support functions. In the thermal model, the propulsion system is assumed to be thermally insulated from the rest of the spacecraft.

A.1.3. Electrical Power System

The electrical power system provides power generation, storage, conditioning, and distribution. It consists of deployable solar-array panels, a body-mounted solar panel, a power-management and distribution unit, and battery packs. The solar-array drive mechanisms are included as separate heat-dissipating elements where relevant.

A.1.4. Command and Data Handling

The command and data handling architecture is represented using a primary spacecraft computer and a payload data-handling unit. These components are responsible for command execution, housekeeping acquisition, payload data handling, and communication with the spacecraft communication units. In the thermal model, they are included as internal heat sources with mode-dependent power dissipation.

A.1.5. Communication Subsystem

The communication subsystem contains equipment for telemetry, telecommand, and payload-data downlink. It is represented by generic radio-frequency transceiver units, antenna units, and high-rate communication electronics. The communication units are activated according to the assumed operational mode and altitude-dependent duty-cycle logic.

A.1.6. Attitude Determination and Control System

The attitude determination and control system includes attitude sensors, angular-rate sensors, magnetic-field sensors, magnetic torquers, and reaction wheels. The thermal model distinguishes between low-power standby or sampling states and higher-power operating states where applicable.

A.1.7. Structure and Thermal Protection

The primary structure is represented as a 16U aluminium CubeSat frame. The ram-facing thermal protection element is modelled as a high-temperature polymer-composite heat shield. The structure provides the main conductive path between internal equipment, external panels, solar arrays, and the heat shield.

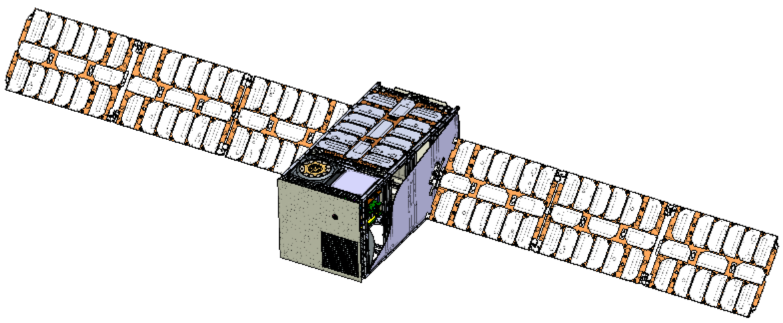

Figure A.1: External view of the conceptual 16U CubeSat platform.

A.2. Power Consumption Analysis

The power dissipation of the major spacecraft equipment groups under the defined operating modes is summarized in Table A.1. These values are derived from component-level specifications and subsystem estimates and are used as direct inputs to the thermal model.

Table A.1: Power consumption and operational status in cold and hot modes.

Equipment Group	Component	Mode	Power (W)	Cold	Hot
Attitude Control	Control electronics	Typical	0.17		
		Maximum	0.90		x
	Navigation receiver	On	1.32		x
	Reaction wheel assembly	Zero torque	1.20		
		Maximum torque	10.00		x
	Magnetic torquer x-axis	Standby	0.003		
		Average	0.50		
		Maximum	0.625		x
	Magnetic torquer y-axis	Standby	0.003		
		Average	0.50		
		Maximum	0.625		x
	Magnetic torquer z-axis	Standby	0.003		
		Average	0.72		
		Maximum	0.9		
	Angular-rate sensor	On	1.50		x
	Optical attitude sensor	Minimum	0.37		
		Typical	0.70		
		Maximum	1.00		x
	Magnetic-field sensor	Not sampling	0.00033		
		Sampling	0.00825		x
	Sun sensors x6	Not sampling	0.00016		
		Sampling	0.0693		x
Power System	Battery packs x2	Typical	0.256	x	x
	Power management unit	Typical	0.50	x	x
	Solar-array drive mechanisms x2	Minimum	0.65		
		Maximum	2.25		
Avionics	Spacecraft computer	Typical	0.17	x	
		Maximum	0.90		x
	Telecommand transceiver	RX	0.61	x	x
		TX (min)	1.20	x	
		TX (max)	3.60		x
Communication	Low-gain antenna unit	Maximum	5.00	x	
	High-rate communication unit	Standby	7.5		x
		On	10.2		
	High-rate antenna unit	Standby	1.00		x
		On (min)	12.70		
		On (max)	23.00		
	Duplex antenna unit	Standby	0.30		x
		TX (min)	4.10		
		TX (max)	16.00		
Propulsion	Compact propulsion system	Idle	4.4		x
		Heating	17.3		
		Firing	6.4		
		Attitude control	6.9		
Payloads	Payload D	Idle	0.38		
		On	0.80		x
	Payload A	On	12.00		x
		Standby	2		
	Payload B	On	17		x

Equipment Group	Component	Mode	Power (W)	Cold	Hot
	Payload C	On	—	x	x
	Payload E	On	2.30		x
	Payload E antenna	On	—		x
	Payload data unit	On	2.00		x
Total Power				**7.74 W**	**69.4 W**

A.3. Thermal Requirements

Thermal control is a critical design driver for the conceptual 16U CubeSat platform, particularly because of the demanding VLEO environment. In addition to conventional orbital heat loads, the spacecraft is exposed to aerothermal heating on the ram-facing side. The thermal design must ensure that all equipment remains within its allowable temperature range during both nominal and off-nominal mission phases.

The thermal requirements are defined across three categories:

- **Operational temperature range:** The temperature range within which a component must function reliably during normal operation.
- **Survival temperature range:** The temperature limits beyond which a component may suffer permanent damage, even if it is not operating.
- **Interface temperature requirements:** The allowable temperature range at selected mechanical or electrical interfaces, used to limit thermal stress, local overheating, or performance degradation.

The operational and survival temperature ranges for the major equipment groups are summarized in Table A.2 and Table A.3, respectively. Interface temperature limits are given in Table A.4. Following ECSS-E-30 Part 1A, a thermal uncertainty margin of $\pm 10\,\mathrm{K}$ is applied to the temperature requirements when evaluating deterministic thermal margins [37].

Table A.2: Operational temperature range of spacecraft equipment.

Equipment Group	Component	Operational Range (°C)
Attitude Control	Control electronics	-30 to +85
	Navigation receiver	TBD
	Reaction wheel assembly	-40 to +80
	Magnetic torquer assembly	-40 to +85
	Angular-rate sensor	-40 to +85
	Optical attitude sensor	-20 to +40
	Sun sensors	-40 to +100
	Magnetic-field sensor	-40 to +85
Power System	Battery pack	-10 to +50
	Power management unit	-40 to +85
	Deployable solar array	-101 to +101
	Body-mounted solar panel	-40 to +85
	Solar-array drive mechanism	-40 to +80
Avionics	Spacecraft computer	-30 to +85
	Telecommand transceiver	-40 to +75
Communication	Duplex antenna unit	-40 to +85
	High-rate antenna unit	-40 to +85
	Low-gain antenna unit	-40 to +85
	High-rate communication unit	-40 to +85
Payloads	Payload A	-40 to +75
	Payload B	TBD
	Payload D	0 to +60
	Payload C	-35 to +145
	Payload E	-20 to +70
	Payload E antenna	-65 to +125
	Payload data unit	-40 to +50
Propulsion	Compact propulsion system	-16 to +30
Structure	Aluminium spacecraft structure	-100 to +150
	Polymer-composite heat shield	-100 to +1600

Table A.3: Survival temperature range of spacecraft equipment.

Equipment Group	Component	Survival Range (°C)
Attitude Control	Control electronics	-30 to +85
	Navigation receiver	TBD
	Reaction wheel assembly	-40 to +85
	Magnetic torquer assembly	-40 to +85
	Angular-rate sensor	-40 to +90
	Optical attitude sensor	-20 to +40
	Sun sensors	-40 to +125
	Magnetic-field sensor	-40 to +85
Power System	Battery pack	-10 to +50
	Power management unit	-40 to +85
	Deployable solar array	-101 to +101
	Body-mounted solar panel	-40 to +85
	Solar-array drive mechanism	-40 to +80
Avionics	Spacecraft computer	-30 to +85
	Telecommand transceiver	-40 to +85
Communication	Duplex antenna unit	-40 to +85
	High-rate antenna unit	-40 to +85
	Low-gain antenna unit	-40 to +85
	High-rate communication unit	-40 to +85
Payloads	Payload A	-40 to +75
	Payload B	TBD
	Payload D	-40 to +85
	Payload C	-35 to +145
	Payload E	-40 to +80
	Payload E antenna	-70 to +125
	Payload data unit	-40 to +85
Propulsion	Compact propulsion system	-34 to +71
Structure	Aluminium spacecraft structure	-100 to +150
	Polymer-composite heat shield	-100 to +1600

Table A.4: Interface temperature requirements for selected critical components.

Interface	Maximum Allowable Temperature (°C)
Avionics unit to structure	+65
High-rate communication unit to structure	+50
High-rate antenna unit to structure	+57
Duplex antenna unit to structure	+57
Payload data unit to structure	+50

B
Albedo and OLR Plots

This appendix provides the supporting material for the Earth-radiation model comparison in Section 2.5.4: Section B presents the complete latitude-resolved accuracy comparison against CERES underlying the reported RMSE values, and Section B presents the corresponding altitude-resolved thermal-response sensitivity of the simplified test plate.

Quantitative comparison

This appendix provides the complete set of zonal-mean latitude profiles and signed percent-error plots underlying the RMSE values summarised in Section 2.5.4, evaluated at four reference dates spanning the 2022 seasonal cycle (both equinoxes and both solstices). Figs. B.1 to B.4 compare the constant, ECSS, and Sasaki predictions directly against the CERES zonal-mean reference; Figs. B.5 to B.8 express the same comparison as signed percent error, showing the latitude- and season-dependent bias structure summarised in Section 2.5.4.

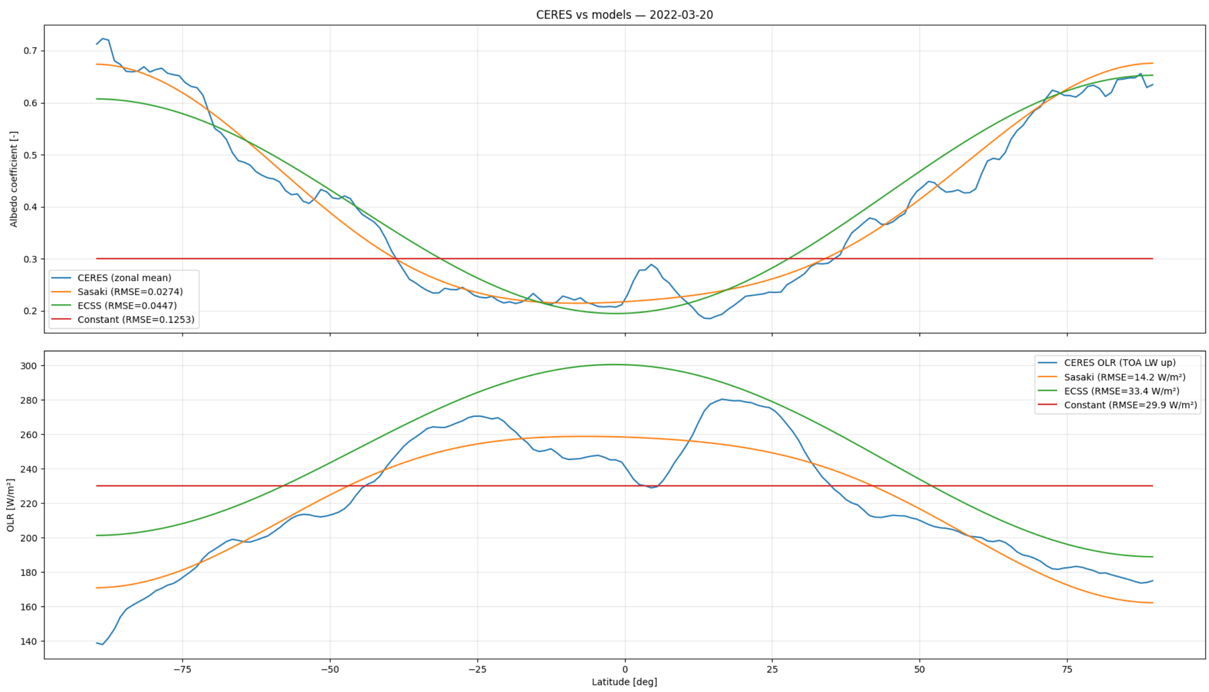

Figure B.1: CERES zonal-mean latitude profiles compared with the constant, ECSS and Sasaki models on 2022-03-20 (equinox). RMSE values are computed using cosine-latitude area weighting; albedo RMSE excludes polar-night latitudes using a validity mask.

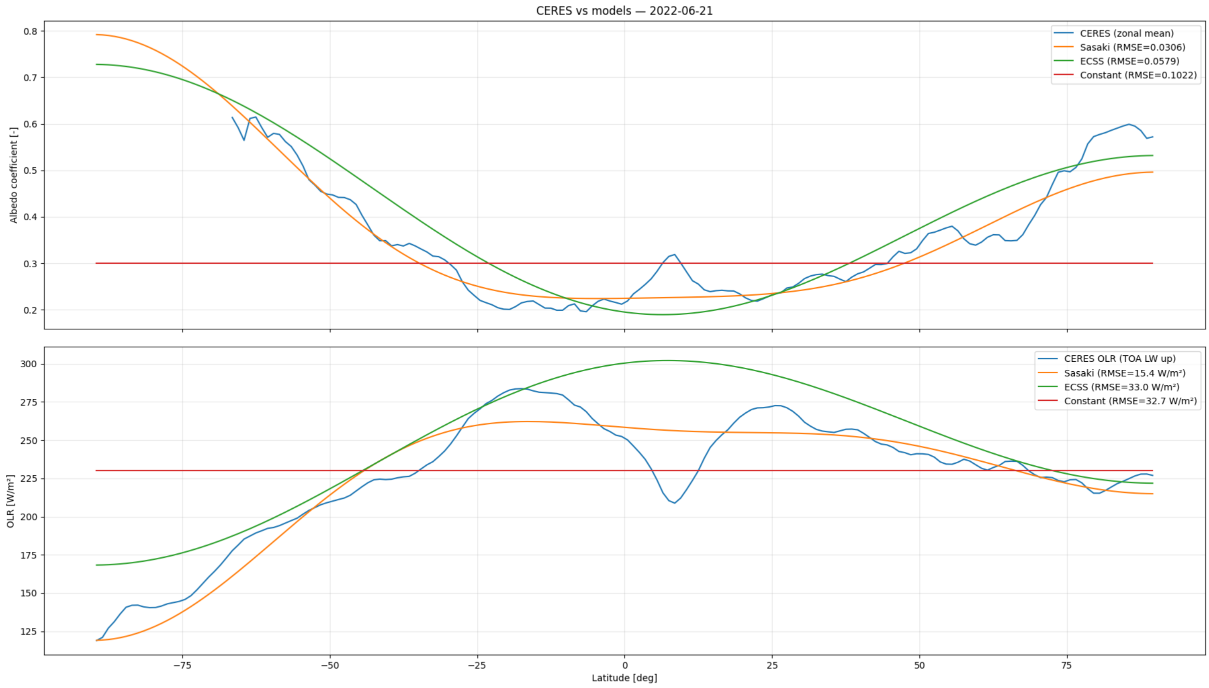

Figure B.2: CERES zonal-mean latitude profiles compared with the constant, ECSS and Sasaki models on 2022-06-21 (June solstice). RMSE definition as in Fig. B.1.

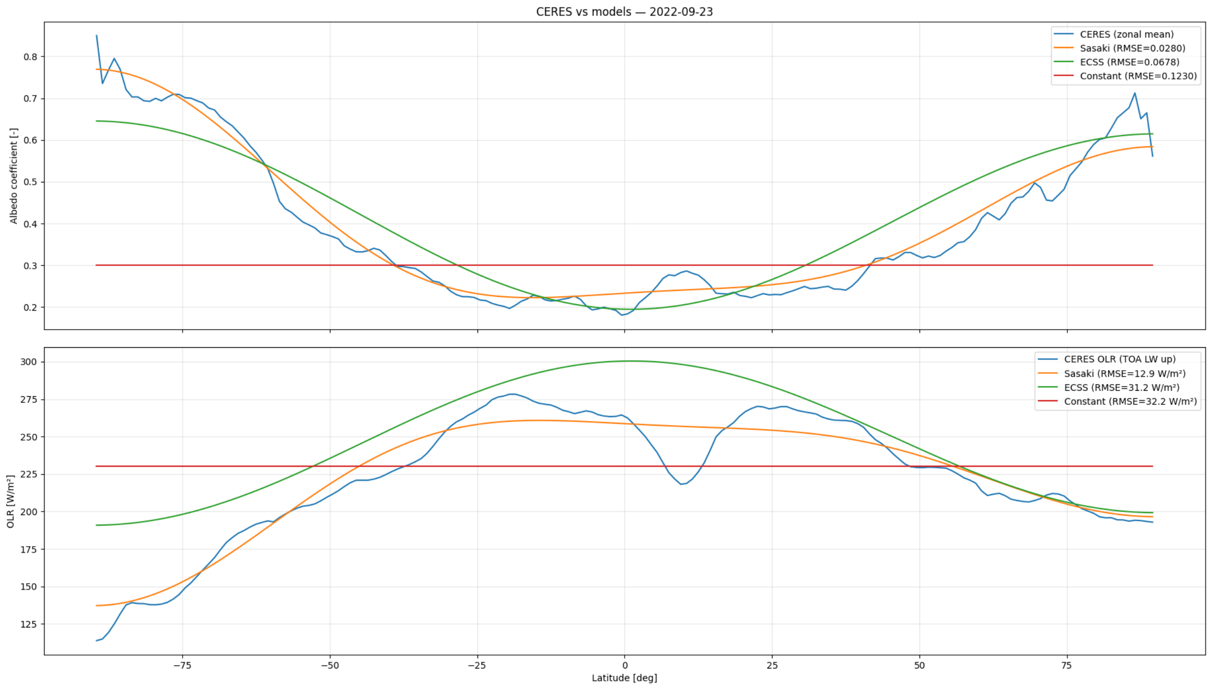

Figure B.3: CERES zonal-mean latitude profiles compared with the constant, ECSS and Sasaki models on 2022-09-23 (equinox). RMSE definition as in Fig. B.1.

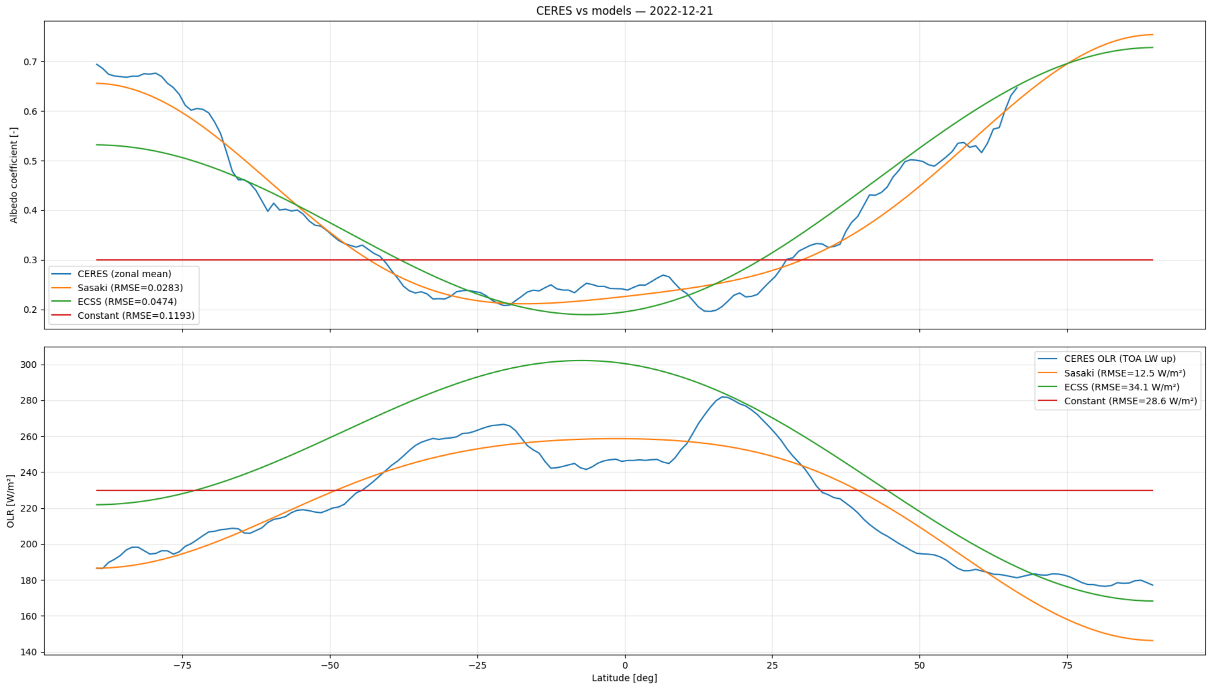

Figure B.4: CERES zonal-mean latitude profiles compared with the constant, ECSS and Sasaki models on 2022-12-21 (December solstice). RMSE definition as in Fig. B.1.

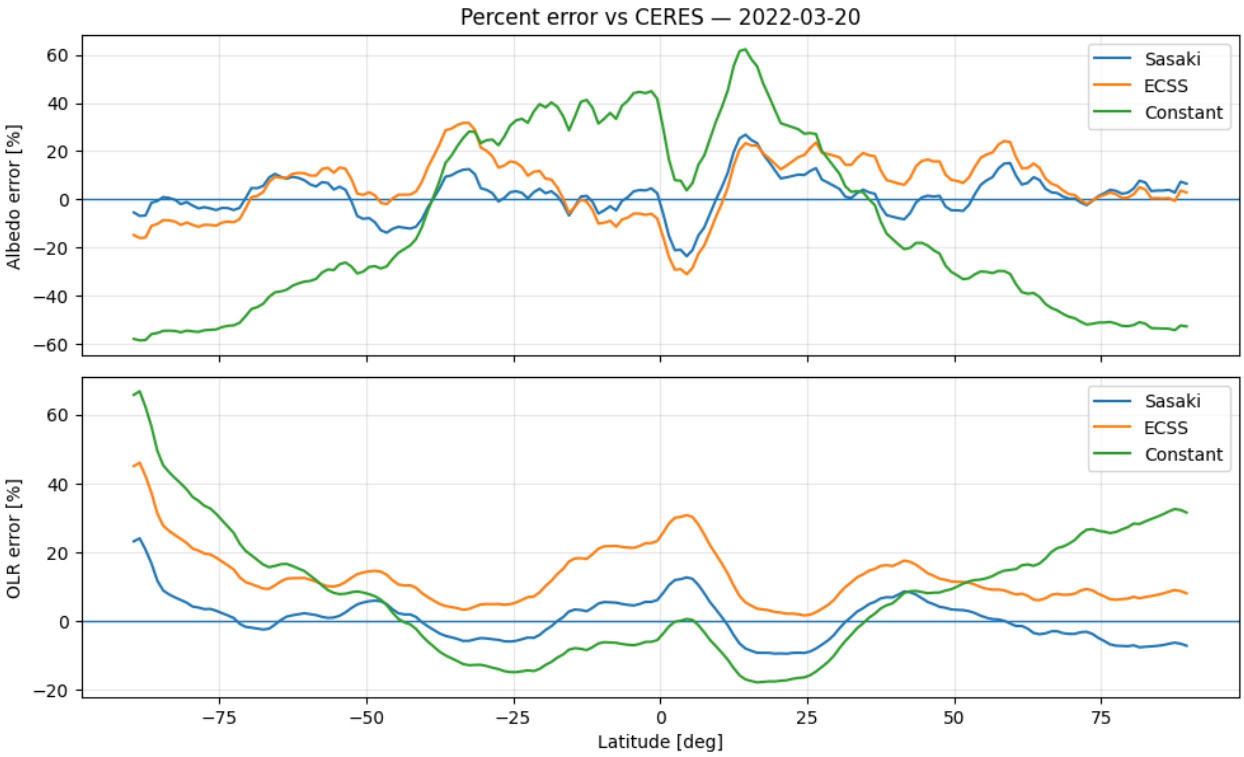

Figure B.5: Signed percent error, (model – CERES)/CERES, of the constant, ECSS, and Sasaki models relative to CERES on 2022-03-20 (equinox); positive values indicate overprediction.

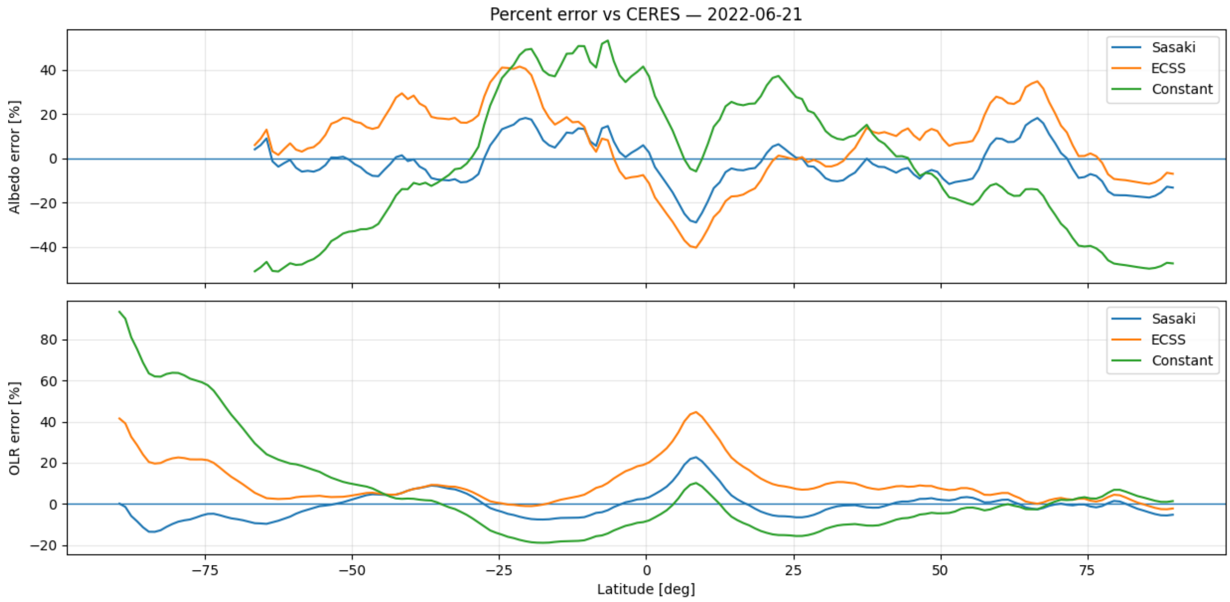

Figure B.6: Signed percent error, (model – CERES)/CERES, of the constant, ECSS, and Sasaki models relative to CERES on 2022-06-21 (June solstice); positive values indicate overprediction.

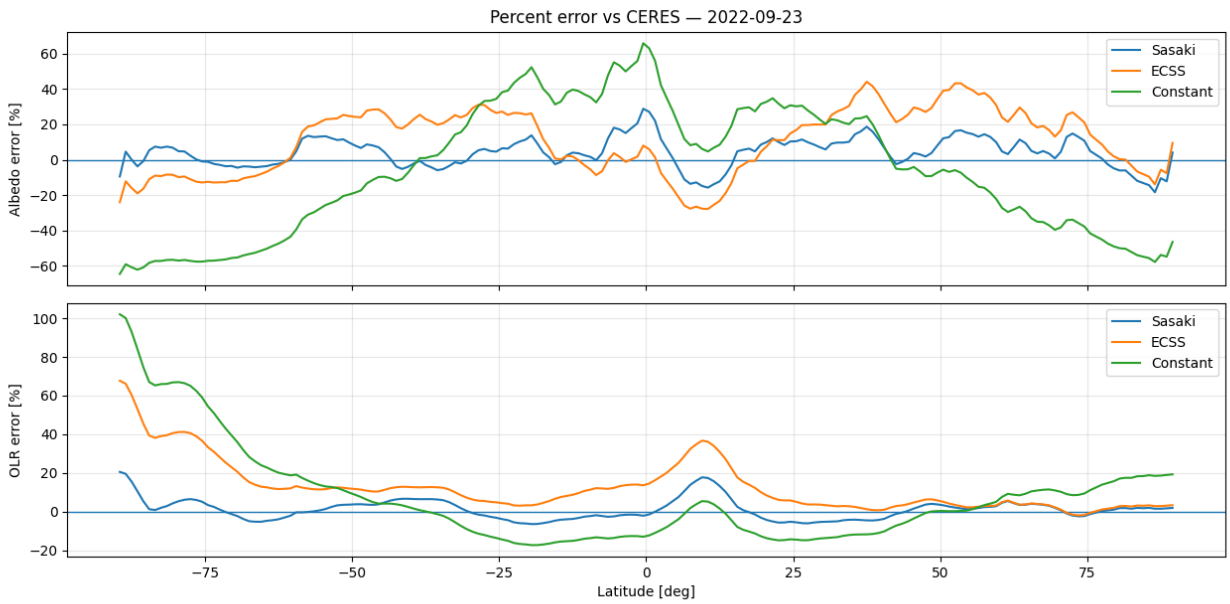

Figure B.7: Signed percent error, (model – CERES)/CERES, of the constant, ECSS, and Sasaki models relative to CERES on 2022-09-23 (equinox); positive values indicate overprediction.

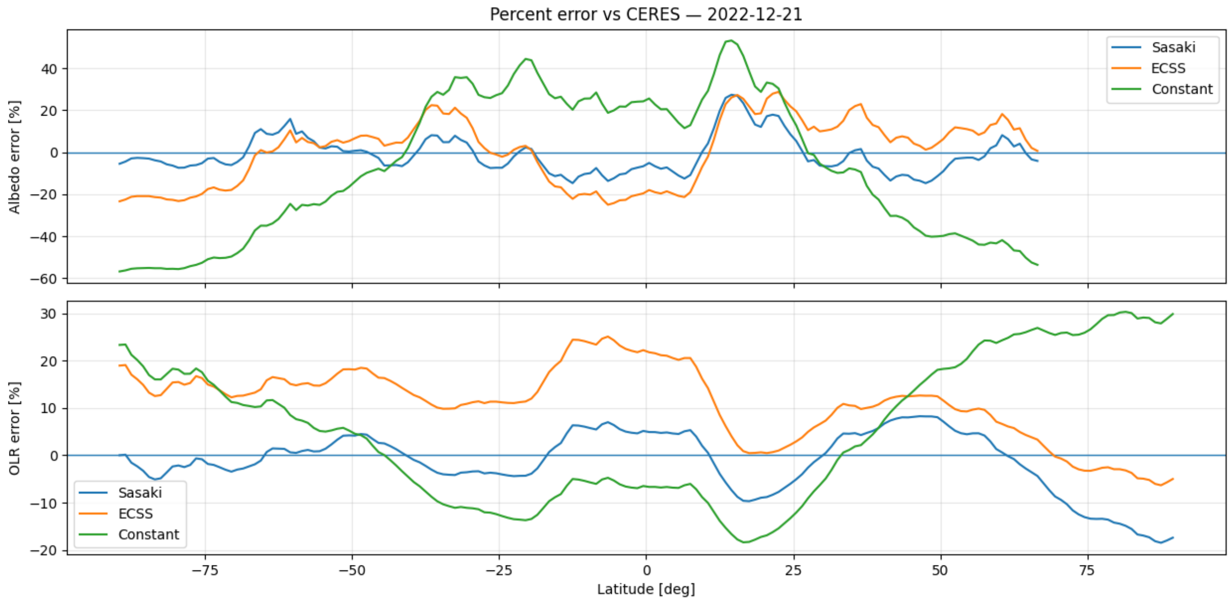

Figure B.8: Signed percent error, (model − CERES)/CERES, of the constant, ECSS, and Sasaki models relative to CERES on 2022-12-21 (December solstice); positive values indicate overprediction.

Temperature Effects

This appendix provides the complete altitude-resolved temperature and deviation curves summarised at 100 km in Section 2.5.4, for the full 100–500 km test sweep. The ΔT spread between the three Earth-radiation models does not decrease appreciably with altitude, since neither the OLR nor the albedo flux varies strongly with spacecraft altitude over this range. Aerothermal heating, by contrast, decreases by several orders of magnitude across the same range (Fig. 2.8). The practical significance of Earth-radiation-model choice therefore grows toward the upper end of the VLEO altitude range, where aerothermal heating no longer dominates the total heat load.

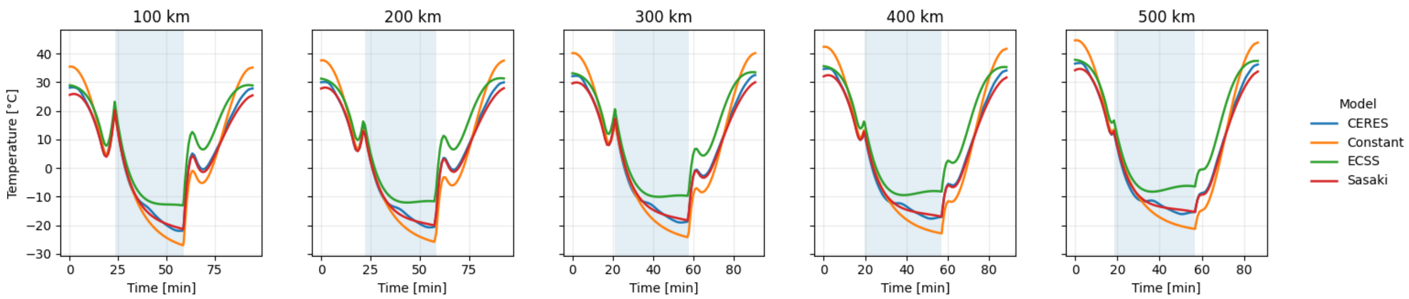

Figure B.9: Temperature evolution over one orbit on 21 December 2022, comparing the CERES-driven reference case with the constant, ECSS, and Sasaki Earth-radiation models, across circular polar orbit altitudes of 100 to 500 km (five panels). Shaded regions indicate eclipse intervals.

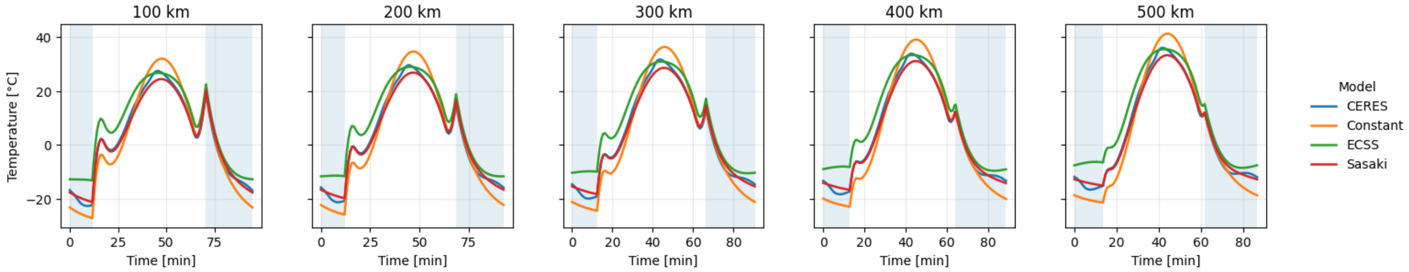

Figure B.10: Temperature evolution over one orbit on 21 June 2022, comparing the CERES-driven reference case with the constant, ECSS, and Sasaki Earth-radiation models, across circular polar orbit altitudes of 100 to 500 km (five panels). Shaded regions indicate eclipse intervals.

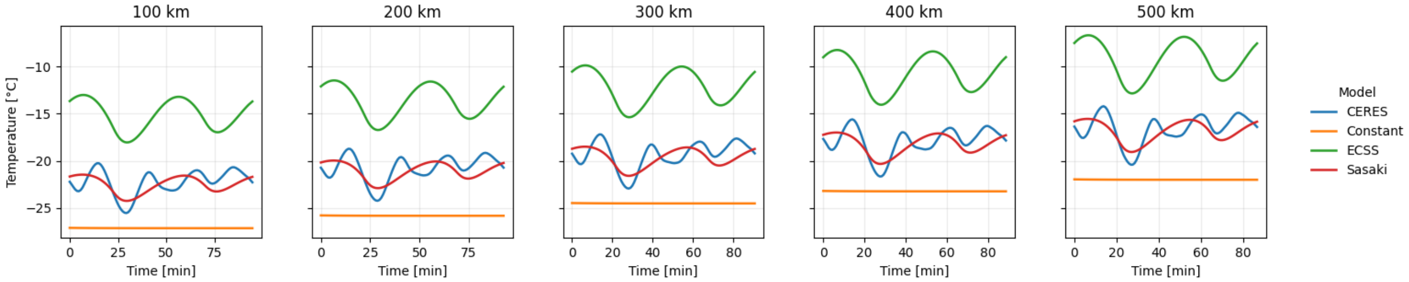

Figure B.11: Temperature evolution over one orbit on 6 March 2022, comparing the CERES-driven reference case with the constant, ECSS, and Sasaki Earth-radiation models, across circular polar orbit altitudes of 100 to 500 km (five panels). Shaded regions indicate eclipse intervals.

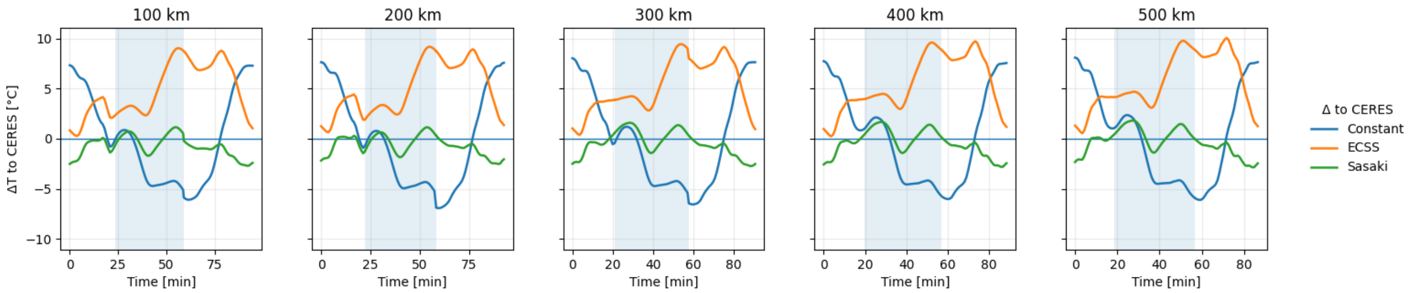

Figure B.12: Temperature deviation ΔT relative to the CERES-driven reference case on 21 December 2022, for the constant, ECSS, and Sasaki Earth-radiation models, across circular polar orbit altitudes of 100 to 500 km (five panels). Shaded regions indicate eclipse intervals.

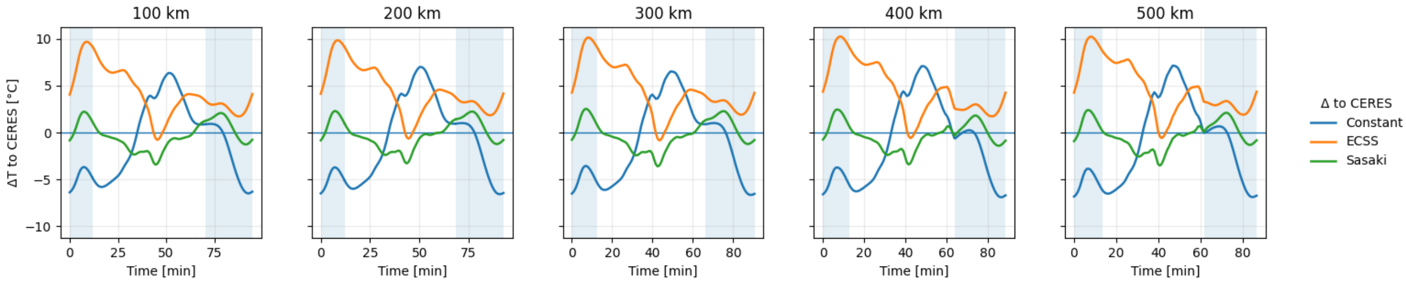

Figure B.13: Temperature deviation ΔT relative to the CERES-driven reference case on 21 June 2022, for the constant, ECSS, and Sasaki Earth-radiation models, across circular polar orbit altitudes of 100 to 500 km (five panels). Shaded regions indicate eclipse intervals.

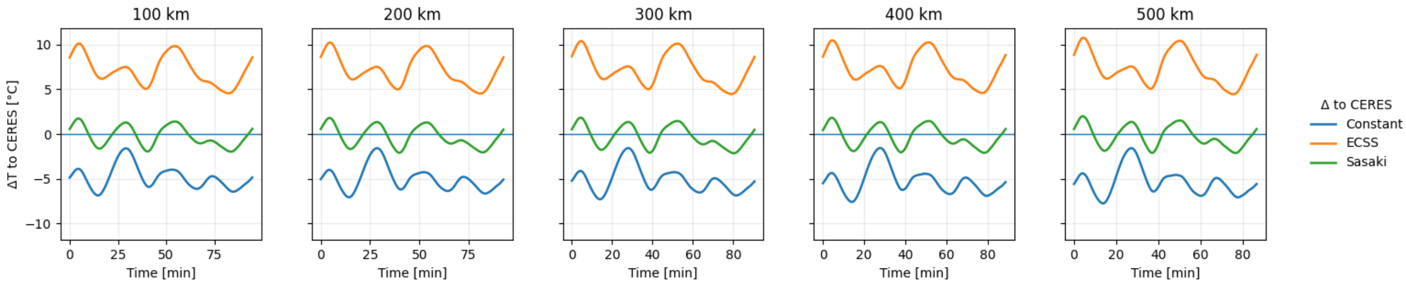

Figure B.14: Temperature deviation ΔT relative to the CERES-driven reference case on 6 March 2022, for the constant, ECSS, and Sasaki Earth-radiation models, across circular polar orbit altitudes of 100 to 500 km (five panels). Shaded regions indicate eclipse intervals.

C

Solar-Panel-to-Structure View-Factor and Gebhart-Factor Calculation

This appendix summarises the analytical calculation used to determine the radiative coupling between the solar-panel node and the structure node in the lumped-parameter thermal network. The calculation consists of three steps: first, the direct view factors are obtained from ECSS-E-HB-31-01 Part 1A [37]; second, the corresponding space-facing view factors are calculated from view-factor summation; third, the Gebhart factors are calculated to account for diffuse multiple reflections.

The solar panel and spacecraft side wall are idealised as rectangular surfaces connected along a common edge. The view factor between composite rectangular areas is evaluated using the ECSS relation

$$F_{3(2,4,6)} = \frac{A_{(1,3)}F_{(1,3)(2,4)} + A_{(3,5)}F_{(3,5)(4,6)} - A_1F_{12} - A_5F_{56}}{2A_3}. \tag{C.1}$$

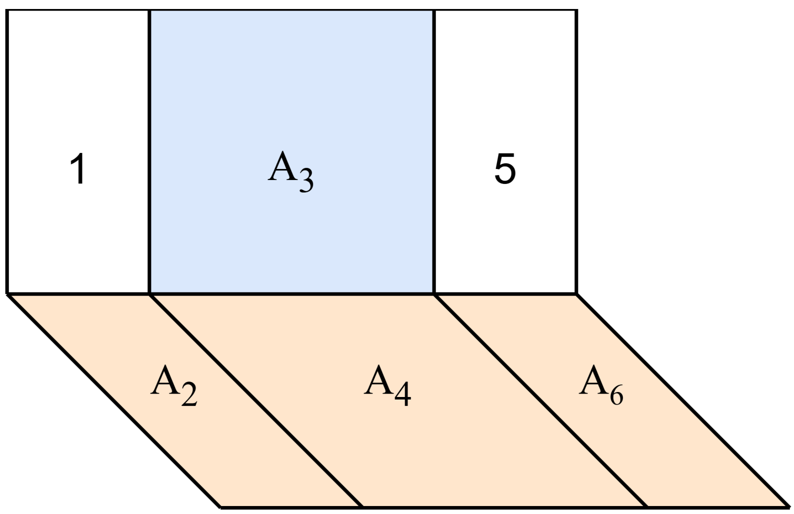

Figure C.1: ECSS view-factor configuration for two rectangular surfaces with one common edge, adapted from ECSS-E-HB-31-01 Part 1A. [37]

The individual rectangular view factors required in Equation C.1 are calculated using the ECSS expression for two perpendicular rectangles with a common edge

$$F_{ij} = \frac{1}{\pi L}\left\{L\tan^{-1}\left(\frac{1}{L}\right) + N\tan^{-1}\left(\frac{1}{N}\right) - \sqrt{N^2+L^2}\,\tan^{-1}\left(\frac{1}{\sqrt{N^2+L^2}}\right)\right.$$
$$+\frac{1}{4}\ln\left[\frac{(1+N^2)(1+L^2)}{1+N^2+L^2}\left(\frac{L^2(1+N^2+L^2)}{(1+L^2)(N^2+L^2)}\right)^{L^2}\right.$$
$$\left.\left.\times\left(\frac{N^2(1+N^2+L^2)}{(1+N^2)(N^2+L^2)}\right)^{N^2}\right]\right\}, \tag{C.2}$$

where L and N are the non-dimensional geometric parameters of the rectangular-surface configuration shown in Figure C.2.

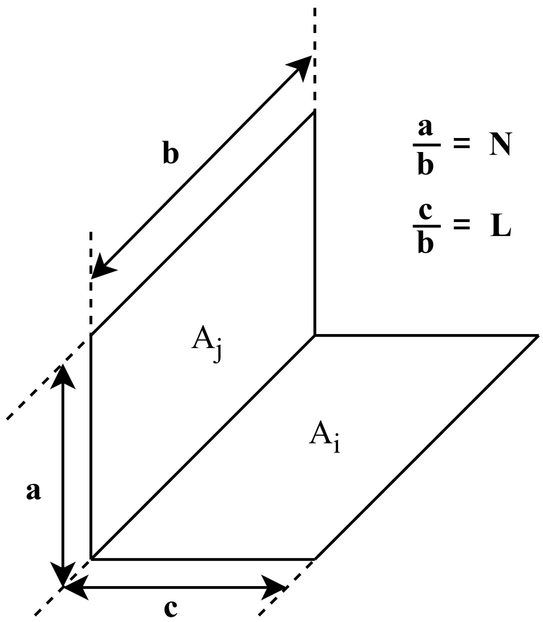

Figure C.2: ECSS view-factor configuration for two perpendicular rectangles with a common edge, adapted from ECSS-E-HB-31-01 Part 1A. [37]

The resulting solar-panel-to-structure view factor is denoted by F_{12} in the lumped-parameter network. The reverse view factor is obtained from reciprocity:

$$A_1F_{12} = A_2F_{21}, \tag{C.3}$$

such that

$$F_{21} = \frac{A_1}{A_2}F_{12}. \tag{C.4}$$

Once F_{12} and F_{21} are known, the corresponding space-facing view factors follow from view-factor summation. For the solar-panel and structure nodes,

$$F_{1s} = 1 - F_{12}, \tag{C.5}$$
$$F_{2s} = 1 - F_{21}. \tag{C.6}$$

The heat shield is assumed to radiate only to deep space, hence

$$F_{4s} = 1. \tag{C.7}$$

The Gebhart factors are calculated from the direct view factors and infrared emissivities. For diffuse-grey surfaces, the Gebhart factor is defined as

$$B_{ij} = F_{ij}\varepsilon_j + \sum_{k=1}^{N} (1-\varepsilon_k)\, F_{ik} B_{kj}. \tag{C.8}$$

For the two-surface exchange between the solar-panel node and the structure node, with deep space as the remaining radiative sink, the Gebhart equations are written for surfaces 1 and 2. Since deep space is treated as an absorbing sink, reflected radiation from space is not returned to the enclosure. The Gebhart factors between the two spacecraft surfaces are therefore

$$B_{11} = (1-\varepsilon_2)\, F_{12} B_{21}, \tag{C.9}$$
$$B_{12} = F_{12}\varepsilon_2 + (1-\varepsilon_2)\, F_{12} B_{22}, \tag{C.10}$$
$$B_{21} = F_{21}\varepsilon_1 + (1-\varepsilon_1)\, F_{21} B_{11}, \tag{C.11}$$
$$B_{22} = (1-\varepsilon_1)\, F_{21} B_{12}. \tag{C.12}$$

Equivalently, B_{12} may be obtained from the Gebhart reciprocity relation

$$\varepsilon_1 A_1 B_{12} = \varepsilon_2 A_2 B_{21}, \tag{C.13}$$

which gives

$$B_{12} = \frac{\varepsilon_2 A_2}{\varepsilon_1 A_1} B_{21}. \tag{C.14}$$

The space-facing Gebhart factors follow from conservation:

$$B_{1s} = 1 - B_{11} - B_{12}, \tag{C.15}$$
$$B_{2s} = 1 - B_{21} - B_{22}, \tag{C.16}$$
$$B_{4s} = 1. \tag{C.17}$$

The radiative coupling factors used in the lumped-parameter thermal network are then obtained from

$$R_{ij} = \varepsilon_i A_i B_{ij} = \varepsilon_j A_j B_{ji}. \tag{C.18}$$

For the four-node network, this gives the retained radiative couplings

$$R_{12} = \varepsilon_1 A_1 B_{12} = \varepsilon_2 A_2 B_{21}, \tag{C.19}$$
$$R_{1s} = \varepsilon_1 A_1 B_{1s}, \tag{C.20}$$
$$R_{2s} = \varepsilon_2 A_2 B_{2s}, \tag{C.21}$$
$$R_{4s} = \varepsilon_4 A_4 B_{4s}. \tag{C.22}$$

D

LTspice Implementation Details

Initial Conditions

The initial thermal state of the four-node network is prescribed using the `.ic` command. The same initial temperature is assigned to all nodes. The node voltage depends on the selected network potential: for a temperature-potential formulation $V_{\text{init}} = T_{\text{init}}$, while for the radiative-potential formulation used here, $V_{\text{init}} = \sigma T_{\text{init}}^4$. The initial conditions are implemented as

```
.ic V(n001)=Vinit V(n002)=Vinit V(n003)=Vinit V(n004)=Vinit
```

Heat Load Implementation

Constant internal dissipation terms are implemented as fixed current sources. Time-varying loads, such as orbit-dependent aerothermal heating, are prescribed as piecewise-linear (`PWL`) data consisting of time–heat-rate pairs, entered either inline,

```
PWL(t1 W1 t2 W2 t3 W3)
```

or from an external file for longer time series,

```
PWL FILE="Example.txt"
```

Transient Solver Command

The transient analysis is executed using

```
.tran 0 t_end 0 dt_max uic
```

where `t_end` is the simulated duration and `dt_max` is the maximum permitted time step. The maximum time step is specified explicitly to ensure adequate resolution of the orbital heat-load profiles. The `uic` flag enforces the prescribed initial conditions rather than deriving them from an operating-point solution.

Integration Method

LTspice offers trapezoidal, Gear, and modified trapezoidal integration. Trapezoidal integration is accurate but can produce numerical oscillations (trap ringing) at discontinuous load inputs. Gear integration suppresses these artefacts through added numerical damping, at the risk of attenuating genuine transient behaviour. Modified trapezoidal integration eliminates trap ringing while preserving the accuracy of the trapezoidal method and is therefore selected for all transient simulations in this work.

E

Material and Optical Property Tables

E.1. Representative Material Properties of Simplified Component Domains

Table E.1: Representative properties used for the simplified box-like component structures in the COMSOL thermal model. Values are effective bulk properties derived from subsystem mass budgets and material data.

Component	**Dimension [mm]**	**Mean Density** ($kg\,m^{-3}$)	**Mass** (g)	**Mean Thermal Cond.** ($W\,m^{-1}\,K^{-1}$)	**Mean Heat Capacity** ($J\,kg^{-1}\,K^{-1}$)
Payload E	44x100x95	1913.87	800	0.3	1369
Payload A	100x100x243.13	822.6	2000	0.3	1369
Payload B	96x94x107.149	1900	TBD	0.3	1369
Payload C	N.A	2700	TBD	238	900
Compact propulsion system	209.3x209.3x110	1257	6060	238	900
Deployable solar array	207.1x930x9.8	470.46	888	58.47	804.43
Body-mounted solar panel	326.5x209x3	2212.83	453	138.46	796.50
Battery pack	94.8x95x42	1271.6	482	13.40	1120.1
Power management unit	95x95x38.8	1030.3	360.8	112.67	800.81
Solar-array drive mechanism	98x54.1x10	1508.9	80	127.92	834.61
Spacecraft computer	92.1x86.5x16.3	668.87	87	61.03	671.0
High-rate communication unit	95x95x31.65	1060.7	303	122.87	815.73
Duplex antenna unit	98.0x98.0x20	567.4	109	160.9795	634.402
High-rate antenna unit	95.60x100.5x12.85	1223.1	151	144.960947	821.18
Telecommand transceiver	98x82.6x8.83	338.22	24	124.31	623.83
Optical attitude sensor	29x29x38.1	3495.40	112	160	900
Control electronics	92.59x89.46x16.44	929.47	126.57	62.39	661.11
Sun sensors	22x11x5	1818.18182	2.2	103.87	869.86
Magnetic-field sensor	23x20x8	2173.9	8	110.67	883.61
Magnetic torquer assembly	95x95x83.6	1760.129	1328	141.4998	632.7248
Aluminium spacecraft structure	N.A	2800	2295	130	862
Support structure	N.A	2800	445	130	862
Polymer-composite heat shield	N.A	1380	TBD	0.3	1250

E.2. Initial Optical Properties of Spacecraft External Surfaces

Table E.2: Optical properties assigned to the spacecraft external surfaces in the initial configuration of the thermal model. Entries marked unk are not specified at the modelling stage; these are treated as design parameters in Chapter 4.

Component	Surface Coating / Material	Emissivity [−]	Absorptivity [−]
Payload E	unk	-	-
Payload A	unk	-	-
Payload B [79]	PCB	0.89	0.46
Payload C	unk	-	-
Compact propulsion system	unk	-	-
Deployable solar array cells [79]	Triple-junction solar cell	0.85	0.93
Body-mounted solar panel cells [79]	Triple-junction solar cell	0.85	0.93
Deployable solar array PCB [79]	PCB	0.89	0.46
Body-mounted solar panel PCB [79]	PCB	0.89	0.46
Battery pack	unk	0.92	-
Power management unit	unk	0.92	-
Solar-array drive mechanism	unk	-	-
Spacecraft computer	unk	0.92	-
High-rate communication unit	unk	0.92	-
Duplex antenna unit [35]	Low-loss RF substrate	0.8	0.22
High-rate antenna unit [35]	Low-loss RF substrate	0.8	0.22
Telecommand transceiver [79]	PCB	0.89	0.46
Optical attitude sensor	unk	-	-
Control electronics	unk	0.92	-
Sun sensors	unk	0.92	-
Magnetic-field sensor	unk	0.92	-
Magnetic torquer assembly	unk	0.92	-
Aluminium spacecraft structure [94]	Bare-aluminium protective treatment	0.06	0.13
Rail structure [35]	Hard-anodised aluminium	0.76-0.84	0.27-0.53
Support structure [94]	Bare-aluminium protective treatment	0.06	0.13
Polymer-composite heat shield	Carbon-filled polymer composite	0.8	0.12