CubeSat Communication Architectures for the ESA RAMSES Mission

Paolo Concari

This thesis is submitted to the University of Hertfordshire in partial fulfilment of the requirements of the degree of MSc by Research

School of Physics, Engineering and Computer Science,
University of Hertfordshire

September 2026

Principal Supervisor: Kourtessis, Pandelis
Second Supervisor: Robinson, Matthew

Abstract

This thesis investigates communication architectures for CubeSats deployed as part of the ESA RAMSES mission to the near-Earth asteroid (99942) Apophis during its exceptional close encounter with Earth in April 2029. Two CubeSat configurations are defined: an orbiter (RCS1) and a lander (RCS2), both subject to severe constraints in power, mass, volume and antenna accommodation. The study focuses on scientific telemetry return and evaluates two candidate architectures: Direct-To-Earth (DTE) communication and relay via an Inter-Satellite-Link (ISL) through the RAMSES spacecraft.

The originality of the research lies in the mission-specific integration and application of established communication-analysis methods to the RAMSES CubeSat cases. A communication modelling and assessment framework is developed that combines time-varying Earth–Apophis and local spacecraft geometry, Apophis rotational behaviour, antenna radiation patterns, link feasibility, discrete communication modes and cumulative data return. This provides a consistent basis for comparing DTE and ISL for both the orbiting and landed CubeSat configurations. For RCS2, a statistical treatment of the landed population accounts for variations in landing location and initial conditions, allowing communication performance and robustness to be assessed beyond a single nominal landing case.

The results show that the two CubeSat cases lead to different communication-architecture implications. For RCS1, DTE and ISL provide approximately 39.6 Gbit and 62.1 Gbit, respectively, over the analysed mission interval. DTE is strongly dependent on the Earth–Apophis range and returns most of its data around close approach, whereas ISL provides a more sustained data-return profile with lower onboard resource demand. For RCS2, performance is strongly dependent on the realised surface geometry and asteroid rotational state. At 60 h, the median DTE and ISL returns are approximately 0.59 Gbit and 1.5 Gbit, respectively. DTE can achieve approximately 10–11 Gbit under favourable conditions but has weak median and lower-tail performance. ISL provides more consistent performance across the landed population, although its median return remains below the assumed 9.2 Gbit data allocation.

The thesis supports the adoption of ISL as the preferred combined architecture for the RAMSES CubeSats, using RAMSES as the Earth-return path. This conclusion is strongest for RCS1, while RCS2 requires further refinement, particularly to consolidate the data-volume capability of the landed ISL architecture. The work identifies the principal performance, robustness, resource and operational factors governing the architecture trade-off and provides a structured basis for early-phase CubeSat communication-architecture assessment for RAMSES and other relay-based small-body mission concepts.

Acknowledgment

I would like to express my deepest gratitude to my principal supervisor, Pandelis Kourtessis, whose support, vision and determination made this research possible. I am equally grateful to my second supervisor, Matthew Robinson, whose constant guidance throughout the entire journey has been invaluable. With his continuous availability and insight, I could not have asked for a better guide.

I would also like to thank my examiners, Dr Gavin Cox and Dr Olu Simpson, for the genuine interest they showed in my work and for the constructive and engaging discussion during the examination. I am grateful for the care they brought to the process and for their thoughtful comments, which helped me further strengthen the thesis and bring this work to its conclusion.

My sincere thanks go to my colleagues at ESA and in industry, many of whom supported and encouraged me throughout this research. In particular, I would like to acknowledge Ian as an outstanding source of inspiration and motivation, and Paolo for his continuous and immense support. I would also like to thank Bruno, Charlie, Ingo and Dani for the engaging, not only technical, discussions; and Victoria, my antenna guru, who first told me about the MSc by Research path. I would also like to thank Karim, whose support, guidance and trust over the years have helped me grow and reach the point where I could undertake this research. I am also deeply grateful to Sergio, a long-standing friend who has been there since my journey in space began many years ago, and who remains closely present today. Thank you.

I would also like to thank my family, without whom this journey would have been far more difficult, if not impossible. To my children, Tobia and Juno, thank you for your patience with my endless “papà has to study”, and for your genuine curiosity and interest in what I was doing, which at times led to new ideas or helped reveal problems I did not yet know I had. And for the challenge of writing “Donald Duck” somewhere in this thesis: here it is, Juno and Tobia — this is all for you. To my wife, for facing the weather, no matter what.

Finally, I would like to acknowledge the personal effort behind this work. This journey required deep commitment through demanding periods, long nights, and weekends. The perseverance I discovered during the days of Hera helped me stay dedicated throughout and allowed me to reach an objective I have pursued for a long time.

Table of Contents

Abstract I
Acknowledgment II
List of Acronyms/Abbreviations VI
Chapter 1 Introduction 1
1.1 Motivation 1
1.2 The ESA RAMSES mission and the Apophis opportunity 2
1.3 Research problem and context 3
1.4 Research aim and objectives 5
1.5 Thesis structure 6
Chapter 2 Technology review and RAMSES CubeSat communication design space 9
2.1 Introduction 9
2.2 RAMSES Mission Context and CubeSat Operational Roles 9
2.2.1 RCS1 Orbiting CubeSat 11
2.2.2 RCS2 Landed CubeSat 14
2.3 Spacecraft TT&C and Communication Architecture Fundamentals 15
2.3.1 Direct-To-Earth communication for deep-space CubeSats 16
2.3.2 Inter-Satellite-Link communication and relay architectures 19
2.4 Reference missions and technology heritage 21
2.5 Scope, Assumptions and Analysis Requirements 27
2.5.1 Scope of the Analysis 28
2.5.2 Baseline Assumptions and Input Data 28
2.5.3 Simplifications and Limitations 29
2.6 Chapter summary 31
Chapter 3 Communication modelling framework 32
3.1 Introduction 32
3.2 Mission Geometry and Reference Frames 32
3.3 Propagation assumptions and geometrical visibility 41

3.4 Link Budget, Data-volume Models for Direct-To-Earth and Inter-Satellite-Links 48
3.5 Implementation Framework and Computational Architecture 54
3.5.1 Verification of the Implementation 58
3.6 Design Space Definition and Baseline Configuration 59
3.6.1 Platform Resource constraints 59
3.6.2 Communication Subsystem Sizing and Antenna Architecture 64
3.7 Model Applicability and Limitations 76
3.8 Chapter summary 77
Chapter 4 Orbiting CubeSat (RCS1) Communication analysis 79
4.1 Introduction 79
4.2 RCS1 Nominal Performance over the proximity-operations 79
4.2.1 Direct-To-Earth Communication Performance 79
4.2.2 Inter-Satellite-Link Communication Performance 86
4.3 RCS1 sensitivity analysis 93
4.4 Chapter summary 102
Chapter 5 Landed CubeSat (RCS2) Communication analysis 103
5.1 Introduction 103
5.2 RCS2 Nominal Performance across landed population 103
5.2.1 Geometric visibility characterisation 103
5.2.2 Direct-To-Earth Communication Performance 108
5.2.3 Inter-Satellite-Link Communication Performance 122
5.2.4 Consolidated nominal RCS2 communication behaviour 132
5.3 RCS2 Sensitivity analysis 133
5.3.1 Direct-To-Earth sensitivity and robustness 133
5.3.2 Inter-Satellite-Link sensitivity and robustness 142
5.4 Chapter summary 154
Chapter 6 Communication architecture evaluation and recommendations 156
6.1 Introduction 156

6.2 Evaluation framework and criteria 156
6.3 RCS1 architecture evaluation 158
6.3.1 RCS1 architecture observations 158
6.3.2 RCS1 DTE and ISL evaluation 160
6.4 RCS2 architecture evaluation 162
6.4.1 RCS2 architecture observations 162
6.4.2 RCS2 DTE vs ISL evaluation 164
6.5 Combined assessment and recommendations 166
6.6 Chapter summary 167
Chapter 7 Conclusion and Future Work 168
7.1 Research summary 168
7.2 Main findings 168
7.3 Achievement of research objectives 170
7.4 Limitations and future work 171
7.5 Final conclusion 172
References IX
Appendix XIII
Appendix I Baseline assumptions and analysis requirements XIII
Appendix II RCS2 landing-site population XVII
Appendix III Monte Carlo convergence assessment for RCS2 populations XXII
Appendix IV DTE TM rates and ISL data-mode definitions XXVI
Appendix V Detailed link budget plots XXVIII

List of Acronyms/Abbreviations

Acronym / abbreviation	Meaning
AU	Astronomical Unit
BW	Bandwidth
CA	Closest Approach
CAD	Computer-Aided Design
CCSDS	Consultative Committee for Space Data Systems
CDF	Cumulative Distribution Function
C/No	Carrier-to-noise-density ratio
C-DST	CubeSat Deep Space Transponder
CReMA	Consolidated Report on Mission Analysis
CUCU	COTS UHF Communication Unit
DART	Double Asteroid Redirection Test
DC	Direct Current
DLR	German Aerospace Center
DSN	Deep Space Network
DTE	Direct-To-Earth
DV	Data-volume
Eb/No	Energy-per-bit to noise-density ratio
ECSS	European Cooperation for Space Standardization
EIRP	Equivalent Isotropically Radiated Power
ESA	European Space Agency
ESTEC	European Space Research and Technology Centre
ESTRACK	ESA Tracking Station Network
FEC	Forward Error Correction
FHSS	Frequency Hopping Spread Spectrum
FSPL	Free-Space Path Loss
G/S	Ground Station

Acronym / abbreviation	Meaning
GSA	Global Sensitivity Analysis
G/T	Antenna gain-to-noise-temperature ratio
GNSS	Global Navigation Satellite System
HGA	High-Gain Antenna
HPA	High-Power Amplifier
ISD	Inter-Satellite Distance
ISL	Inter-Satellite-Link
ISS	International Space Station
JAXA	Japan Aerospace Exploration Agency
JPL	Jet Propulsion Laboratory
LEO	Low Earth Orbit
LGA	Low-Gain Antenna
LNA	Low-Noise Amplifier
MASCOT	Mobile Asteroid Surface Scout
MGA	Medium-Gain Antenna
MUPA	Multiple Uplink Per Aperture
NASA	National Aeronautics and Space Administration
NMS	Nested Monte Carlo
NPA	Non-Principal-Axis rotation
PLL	Phase-Locked Loop
QoS	Quality of Service
RAMSES	Rapid Apophis Mission for Space Safety
RC	Residual Carrier
RCS	RAMSES CubeSat
RCS1	RAMSES CubeSat 1
RCS2	RAMSES CubeSat 2
RF	Radio Frequency

Acronym / abbreviation	Meaning
RNG	Ranging
RT	Ranging Tone
RX	Receiver
S/C	Spacecraft
SDR	Software-Defined Radio
SNR	Signal-to-Noise Ratio
SVD	Singular Value Decomposition
TM	Telemetry
TS	Tilt-Spin
TT&C	Telemetry, Tracking and Command
TWTA	Travelling-Wave Tube Amplifier
TX	Transmitter
UTC	Coordinated Universal Time
VSWR	Voltage Standing Wave Ratio
XL	Extended-length CubeSat form factor

Chapter 1 Introduction

1.1 Motivation

The close approach of asteroid 99942 Apophis to Earth in April 2029 provides a rare opportunity to observe a small Solar System body before, during and after a strong planetary encounter. Such observations are relevant both to small-body science and to planetary defence, since the encounter may alter the asteroid's physical state, rotational behaviour, surface environment or dynamical properties [1]. The European Space Agency (ESA) Rapid Apophis Mission for Space Safety (RAMSES) mission is being developed in this context as a time-critical mission to characterise Apophis around the encounter epoch. The fixed encounter date constrains the mission not only by performance requirements, but also by schedule, maturity and implementation risks [2], [3].

A distinctive feature of the RAMSES concept is the use of a multi-spacecraft architecture, in which a main spacecraft operates with two deployed CubeSat-class miniature spacecraft in the vicinity of Apophis [4], [5]. These CubeSats can increase the scientific value of the mission by providing additional observation locations, viewing geometries and, with a landing element, surface-based measurements. However, their contribution depends not only on acquiring scientific data, but on returning a sufficient data-volume within the timescales imposed by their operational profiles.

Communication is therefore a mission-enabling function for the RAMSES CubeSats. This is challenging because deep-space CubeSats are strongly constrained in volume, mass, power and antenna accommodation [6], [7]. Their communication links are affected by large propagation losses, varying geometry, limited contact opportunities and degraded or uncertain operational conditions [8], [9]. The problem is not only whether communication is possible under favourable conditions, but whether scientific data and spacecraft-status information can be returned reliably across the range of operating conditions expected for the CubeSat missions.

Two broad communication approaches are considered in this thesis. A Direct-To-Earth (DTE) architecture gives the CubeSat an independent path to ground stations, but requires it to close the long-range deep-space link using its own limited resources. An Inter-Satellite-Link (ISL) architecture instead allows the CubeSat to communicate locally with the

RAMSES spacecraft, which provides the onward Earth-return path. This reduces the long-range communication demand on the CubeSat but introduces dependence on the RAMSES spacecraft as the communication counterpart.

The candidate communication architectures therefore need to be assessed not only under nominal conditions, but across the geometric and operational variations relevant to each RAMSES CubeSat case. This includes the Earth–Apophis range evolution, the local spacecraft or surface geometry, and the operational constraints affecting each link.

This thesis addresses that problem by comparing DTE and ISL communication architectures in terms of link feasibility, data-return capability and robustness, while also considering operational constraints and implementation aspects.

1.2The ESA RAMSES mission and the Apophis opportunity

The close approach of asteroid 99942 Apophis to Earth on 13 April 2029 allows a near-Earth asteroid to be studied before, during and after a strong planetary encounter. At closest approach, Apophis will pass at a distance comparable with the geostationary-orbit region, creating an exceptional opportunity to observe a small Solar System body during a natural planetary flyby [1], [10]. The encounter is scientifically important because Earth's gravitational perturbation is expected to modify aspects of Apophis' dynamical and physical states or surface environment. Current knowledge of Apophis' shape and non-principal-axis rotational state is based on light-curve, radar and shape-modelling studies [11]-[13]. Observations before, during and after the encounter can therefore characterise the asteroid prior to the perturbation, monitor the encounter phase and assess the resulting changes afterwards [1], [10].

RAMSES is being developed to exploit this fixed encounter opportunity by accompanying Apophis around the 2029 flyby [3], [5]. The mission concept includes a main spacecraft and two CubeSat elements, referred to in this thesis as RAMSES CubeSat 1 (RCS1) and RAMSES CubeSat 2 (RCS2), which are carried to Apophis by the main RAMSES spacecraft and deployed after arrival in the asteroid environment [4], [5]. The CubeSats are intended to provide additional observation geometries and local measurement opportunities through two distinct operational profiles: an orbiting CubeSat, RCS1, associated with extended proximity operations around Apophis, and a landed CubeSat, RCS2, associated with a short surface-operation mission around the close encounter period [5], [10].

These different mission profiles create different communication conditions that cannot be addressed by a single generic communication approach. A more detailed description of the RAMSES mission context, CubeSat roles and operational timelines is provided in Chapter 2.

1.3Research problem and context

The communication architecture for the RAMSES CubeSats is not a generic choice between two radio-link concepts. It is a mission-specific architectural problem driven by the different roles of the two CubeSat elements, the time-varying Earth–Apophis geometry, the proximity geometry with the RAMSES spacecraft, and the limited resources available on CubeSat-class spacecraft. The same communication option may therefore have different implications for RCS1 and RCS2, and the most suitable solution for one CubeSat is not necessarily the most suitable solution for the other.

The core research problem is therefore to assess how Direct-To-Earth and Inter-Satellite-Link architectures perform for the two RAMSES CubeSat elements under their respective mission constraints. In this comparison, DTE provides architectural independence but requires the CubeSat to close the long-range deep-space link using its own resources, whereas ISL shifts the CubeSat communication task to a potentially less demanding local link with the RAMSES spacecraft. The resulting trade-off involves data-return capability, onboard resource allocation, dependence on the mother spacecraft, relay geometry and robustness to uncertain or degraded conditions.

Figure 1 illustrates the two candidate communication paths considered in this thesis.

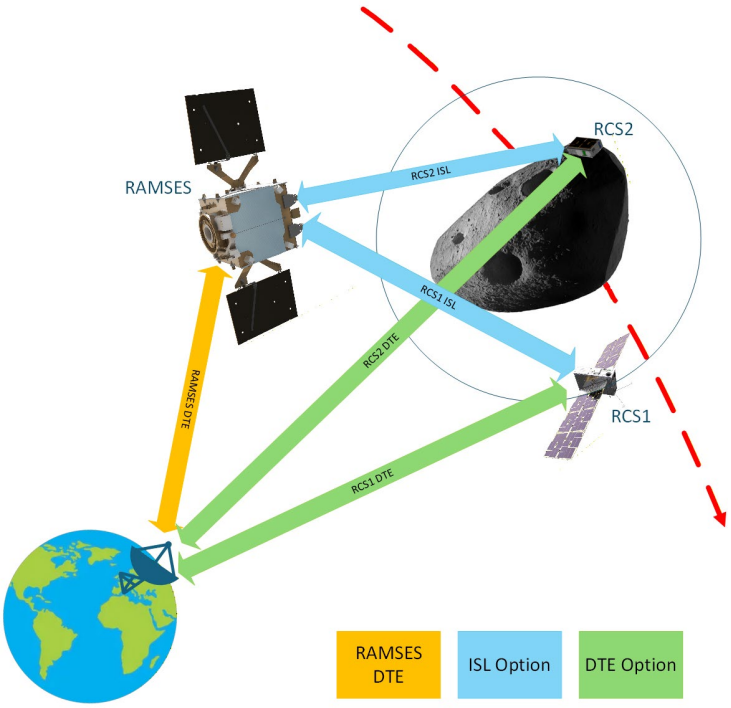


Figure 1: Candidate communication architectures for the RAMSES CubeSats

In Figure 1, the amber link represents the RAMSES Direct-To-Earth path, which is present in both architecture options. The green links represent the CubeSat Direct-To-Earth option, in which RCS1 and RCS2 return telemetry directly to Earth ground stations. The light-blue links represent the Inter-Satellite-Link option, in which the CubeSats transmit locally to RAMSES and the onward Earth return is provided by the RAMSES Direct-To-Earth link.

Communication performance in the RAMSES CubeSat cases is expected to depend strongly on where and when each link is established. For RCS1, the extended mission timeline means that DTE performance must be assessed across the changing Earth–Apophis distance, while the ISL option must be assessed against the local inter-spacecraft geometry during proximity operations. For RCS2, the communication problem is expected to be shaped by the short landed mission and by the rapidly varying geometry around closest approach. In the DTE case, the relevant conditions include the Earth–Apophis range evolution and Earth visibility from the realised landing location. In the ISL case, they include the local geometry between the lander and the RAMSES spacecraft.

The assessment distinguishes nominal capability, sensitivity to degraded or uncertain conditions, and variation across representative scenario populations. Link feasibility, telemetry rate, coverage, and cumulative data return are evaluated alongside percentile and lower-tail behaviour, since robustness is often driven by adverse rather than favourable cases.

Existing deep-space CubeSat communication studies and mission heritage establish the feasibility of both DTE and relay communication, but do not provide a common methodology for comparing the two architectures on a consistent mission-performance basis. In particular, the effects of time-varying geometry, antenna performance and link availability are not generally combined into a comparative assessment of cumulative data return. This gap is particularly relevant to RAMSES, for which no mission-specific comparison of DTE and ISL has been developed for the distinct orbiting RCS1 and landed RCS2 cases.

The contribution of this research is the development and application of a communication-assessment framework that provides a common methodological basis for comparing DTE and ISL architectures across the distinct RAMSES CubeSat cases. The comparison extends beyond nominal communication performance through sensitivity and statistical analyses, allowing robustness and engineering trade-offs to be assessed. Architecture selection is based on the combined assessment of communication performance, robustness, resource and implementation implications, and operational dependence, providing a structured basis for the resulting selection of ISL as the preferred common communication architecture for the RAMSES CubeSats.

The problem is also shaped by the current maturity of the RAMSES mission and CubeSat concepts. The operational timeline, subsystem implementation and geometric conditions are not yet defined with the certainty of a final flight design. The assessment must therefore account for uncertainty in the communication environment and in the assumptions used to represent it, rather than relying only on a single deterministic configuration. In this thesis, robustness is assessed quantitatively as the extent to which communication performance is retained under representative geometric, operational and design variations. It is evaluated through changes in cumulative data return and communication availability or achievable rate relative to the corresponding baseline, together with the distribution of performance across statistical cases where applicable.

1.4Research aim and objectives

The research aim and objectives translate the RAMSES CubeSat communication-architecture problem into the modelling, analysis and evaluation tasks addressed in this thesis. They are defined as follows.

Research Aim

To analyse and compare candidate communication architectures for the RAMSES CubeSats under mission-specific geometric, operational and design constraints, assessing their performance and robustness in support of a risk-aware architecture evaluation for the time-critical deep-space mission to asteroid Apophis.

Research Objectives

The research aim is addressed through the following objectives:

Objective 1: Communication-architecture modelling framework

To develop a consistent modelling framework for evaluating DTE and ISL communication architectures, linking RAMSES mission geometry, communication feasibility and data-return performance.

Objective 2: Nominal performance assessment

To evaluate the nominal DTE and ISL communication performance of RCS1 and RCS2 using mission-specific assumptions and metrics such as link feasibility, achievable telemetry rate, coverage and cumulative data return.

Objective 3: Sensitivity and robustness assessment

To quantify how communication performance changes under representative geometric, operational and design variations, and to identify the main robustness drivers for each CubeSat case and communication architecture.

Objective 4: Statistical characterisation of performance

To assess how communication performance varies across representative populations of geometric and operational cases, identifying the spread between favourable, central and lower-performing cases for the orbiting and landed CubeSat configurations.

Objective 5: Architecture-level evaluation and recommendations

To synthesise the nominal, statistical and sensitivity results into an architecture-level evaluation, identifying the trade-offs, engineering implications and conditions under which DTE and ISL options are favourable for the RAMSES CubeSats.

1.5Thesis structure

The thesis is structured to progress from research framing, through mission and communication background, to modelling, CubeSat-specific performance assessment,

architecture-level evaluation and conclusions. The technical analysis is organised by CubeSat rather than by architecture, with the DTE and ISL options assessed first within the distinct RCS1 and RCS2 mission contexts.

This structure is summarised in Figure 2, which shows how the common modelling framework is applied separately to the RCS1 and RCS2 cases before the results are brought together in the final architecture evaluation.

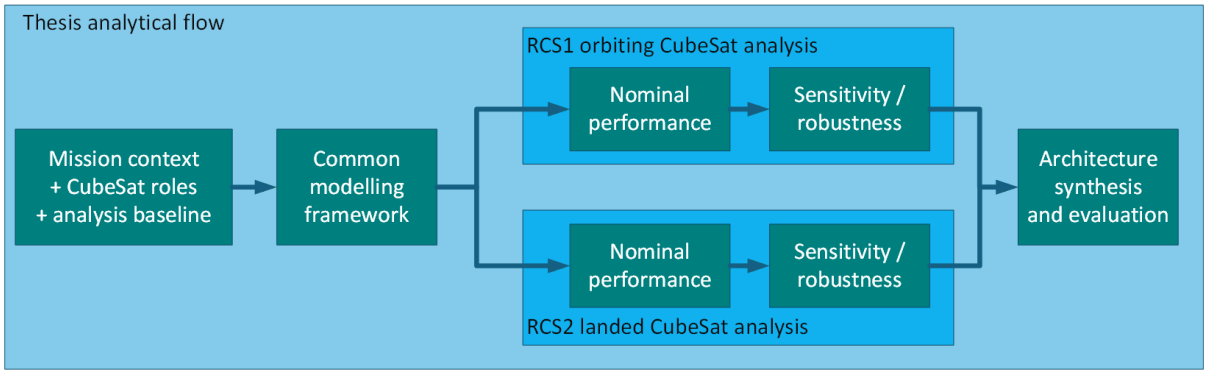


Figure 2: Thesis analytical flow

Chapter 1 introduced the research context, motivation, research problem, aim and objectives, and framed the high-level scope of the study. *Chapter 2* provides the mission, communication and technology background needed to position the work, covering the RAMSES CubeSat operational context, deep-space telemetry, tracking and command (TT&C) functions, DTE and ISL architectures, relevant mission and technology heritage, and the scope, assumptions, limitations and analysis requirements that define the study baseline.

Chapter 3 develops the common modelling framework for the technical analysis. It defines how the RAMSES mission geometry, visibility conditions, communication-link performance and data-return metrics are represented, and establishes the model boundary conditions, applicability limits and computational implementation.

Chapters 4 and *5* apply this framework to the two RAMSES CubeSat cases. Chapter 4 addresses the orbiting CubeSat, RCS1, whose analysis is driven by an extended proximity-operations timeline. Chapter 5 addresses the landed CubeSat, RCS2, whose analysis is driven by the short surface-operation timeline and by the dependence of communication performance on landed geometry. In both chapters, the DTE and ISL options are assessed under nominal conditions and selected sensitivity cases.

Chapter 6 synthesises the RCS1 and RCS2 results into an architecture-level evaluation. It considers the performance, robustness and operational implications of the candidate

communication architectures, including the implications of adopting either a common or differentiated solution for the two CubeSat elements.

Chapter 7 concludes the thesis by summarising the main findings, assessing the achievement of the research objectives, identifying the main limitations of the study and outlining possible future work.

Personal academic work and source basis

This thesis was carried out as personal academic work for the MSc by Research degree. It is based on publicly available sources and on assumptions defined within the thesis. The analysis, interpretations and conclusions presented here are those of the author and do not represent an official position of ESA or the RAMSES project.

Chapter 2 Technology review and RAMSES CubeSat communication design space

2.1Introduction

This chapter provides the mission, communication and technology context for the RAMSES CubeSat communication analysis. It summarises the operational roles of RCS1 and RCS2 and introduces the two communication architectures considered in the thesis: Direct-To-Earth (DTE) communication and Inter-Satellite-Link (ISL) communication through the RAMSES spacecraft.

Relevant mission and technology heritage is then reviewed to support the definition of the communication design space. The chapter closes by defining the reference basis used for the subsequent performance assessment.

2.2RAMSES Mission Context and CubeSat Operational Roles

The RAMSES mission architecture includes the main spacecraft and two deployed CubeSat elements, RAMSES CubeSat 1 (RCS1) and RAMSES CubeSat 2 (RCS2), which are carried during launch and interplanetary cruise and deployed after arrival in the Apophis environment [4], [5]. RCS1 is the orbiting CubeSat and is associated with extended proximity operations around Apophis. RCS2 is the landed CubeSat and is associated with a short surface-operation phase around closest approach. These different roles define the communication timelines, visibility conditions and architecture trade-offs analysed in later chapters.

Figure 3 provides a mission-level view of the RAMSES transfer trajectory towards Apophis and the 2029 Earth close-encounter geometry taken from the RAMSES Consolidated Report on Mission Analysis (CReMA) [10].

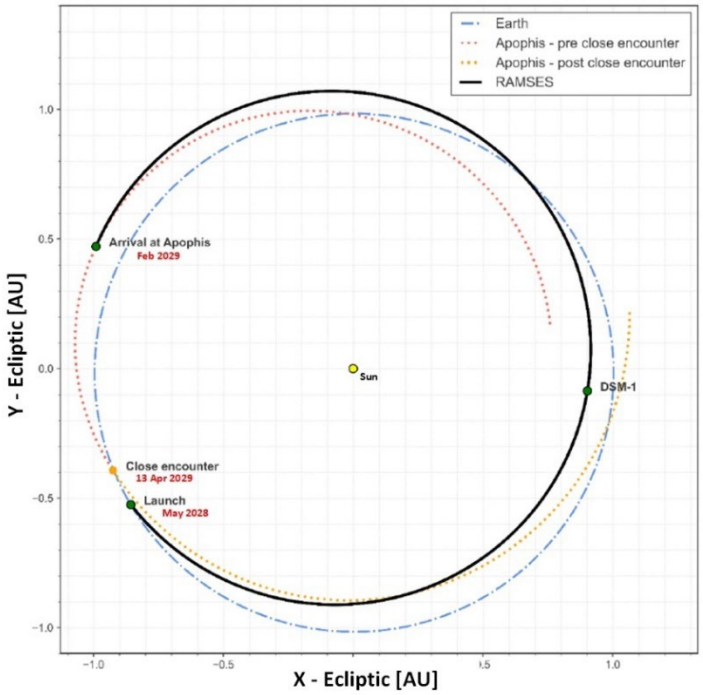


Figure 3: RAMSES spacecraft interplanetary trajectory to Apophis. Source: [10]

Figure 3 shows the RAMSES transfer and the Apophis close encounter in their heliocentric context. The figure identifies the sequence from launch, interplanetary transfer and arrival at Apophis to the close-encounter epoch.

The mission baseline defines a launch period from late April to mid-May 2028, with transfer durations of approximately 269–297 days depending on launch date. Arrival at Apophis is constrained to occur no later than two months before closest approach, with proximity operations starting on 3 March 2029. Around closest approach, Apophis reaches a minimum Earth distance of approximately 32,000 km, whereas RAMSES arrival occurs at an Earth distance of approximately 0.22–0.26 astronomical units (AU), or approximately 33–39 million km depending on launch date. The post-encounter phase occurs as the Earth–Apophis range increases again to several tens of millions of kilometres.

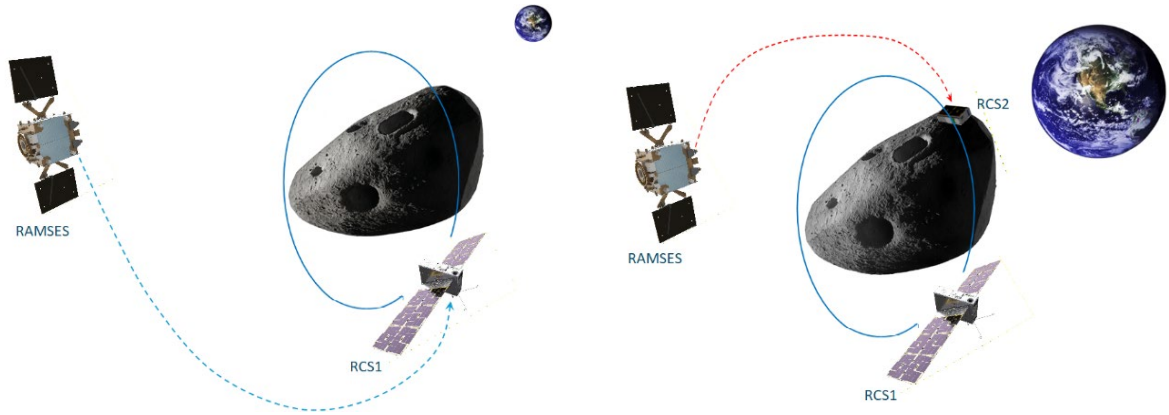
The mission timeline divides proximity operations into pre-encounter, close-encounter and post-encounter phases corresponding respectively to pre-flyby characterisation, encounter monitoring and post-flyby assessment. These phases provide the mission-level reference for the RCS1 and RCS2 communication cases.

The orbiting RCS1 is characterised by extended proximity operations around Apophis. After deployment and commissioning, it operates around Apophis and contributes observations before, during and after the close encounter. The reference RCS1 operations window starts on 15 March 2029, approximately one month before closest approach, when the Earth–Apophis

range is approximately 15 million km. The range then decreases rapidly towards approximately 32,000 km at closest approach and increases again after the encounter, reaching above 40 million km by the end of the 100-day operations window.

The landed RCS2 case is characterised by surface operations around the close encounter. Its deployment and landing are associated with the encounter period, and its surface mission is concentrated over a short battery-supported interval. The reference RCS2 communication window starts 20 h before closest approach, when the Earth–Apophis range is of the order of 500,000 km. Representative 20 h, 40 h and 60 h surface-operation windows are considered, corresponding respectively to data return before closest approach, shortly after the encounter, and over an extended post-encounter surface interval. Across these windows, the Earth–Apophis range decreases to approximately 32,000 km around closest approach and then increases again to approximately 900,000 km by the end of the 60 h window.

Figure 4 provides a simplified visual summary of the two RAMSES CubeSat operational concepts after deployment.



(a) RCS1 (b) RCS2 and RCS1

Figure 4: RAMSES CubeSat simplified operational concepts after deployment

Figure 4 shows in panel (a) the RCS1 deployment and subsequent orbit insertion around Apophis, while in panel (b) the later close-encounter configuration, in which RCS2 reaches its landed configuration on the asteroid surface while RCS1 remains in proximity orbit.

The following subsections define the RCS1 and RCS2 operational roles used in the communication analysis.

2.2.1 RCS1 Orbiting CubeSat

After deployment from the RAMSES main spacecraft and commissioning, RCS1 transfers to an operational orbit around Apophis. For the communication analysis, this orbit is approximated as circular with a characteristic radius of approximately one kilometre. From

this orbit, RCS1 performs close-range observations that complement those of the main spacecraft. Its operations span the pre-encounter, close-encounter and post-encounter phases, and data are therefore generated over the full 100-day operations window.

Because a detailed public observation timeline is not available, data generation is represented as a continuous process over the mission duration using an equivalent constant average rate. This allows communication performance to be assessed independently of detailed operational scheduling assumptions while still capturing cumulative data return over the mission timeline.

From a communication perspective, two geometric scales are relevant for RCS1: the large-scale Earth–Apophis geometry and the local proximity geometry around the asteroid. For DTE, the Earth–Apophis range provides the main range-dependent driver because it determines the free-space path-loss variation over the 100-day window. For ISL, the relevant geometry is the local RCS1–RAMSES separation and mutual viewing geometry within the Apophis proximity environment.

The local proximity geometry is represented differently for the two architectures. For DTE, the RCS1 orbital radius is small compared with the Earth–Apophis range, meaning that the spacecraft-to-Earth distance can be approximated for the purpose of the communication analysis by the Earth–Apophis distance. The RCS1 orbit is nevertheless retained for visibility modelling, since Apophis can block the line of sight between RCS1 and Earth during occultation intervals. For ISL, the local RCS1–RAMSES geometry is represented directly, since the link is established within the Apophis proximity environment.

For the DTE case, the main operational constraints are represented through time-dependent CubeSat pointing strategy, ground-segment performance and contact-availability assumptions. During most of the RCS1 operational timeline, the DTE configuration is assumed to point its antenna toward Earth to improve the link performance. This creates an operational constraint because Earth-pointed DTE communication cannot be maintained simultaneously with the asteroid-pointed attitude that scientific observations may require. Such observations would therefore need to be scheduled outside the communication windows. Around closest approach, the DTE case instead includes a time interval with no dedicated pointing, relying instead on a wide-coverage configuration, which removes the dependence on dedicated Earth pointing in this mission-critical phase.

The ground segment is modelled using time-dependent receive-performance and contact-availability assumptions. Ground station access is represented by an equivalent contact duty

cycle rather than by a detailed contact schedule, avoiding the need to define the RAMSES CubeSat communication operations in detail.

For the RCS1 DTE case, ground station support is represented using the RAMSES mission-analysis, space-to-ground interface and system-requirements information available for this study [10], [14], [15]. The reference case assumes representative receive-performance, antenna-use, Earth-pointing and contact-availability profiles for CubeSat DTE support, with RAMSES/RCS operations compatibility adopted as a modelling assumption. The time evolution of these main operational inputs is summarised in Figure 5.

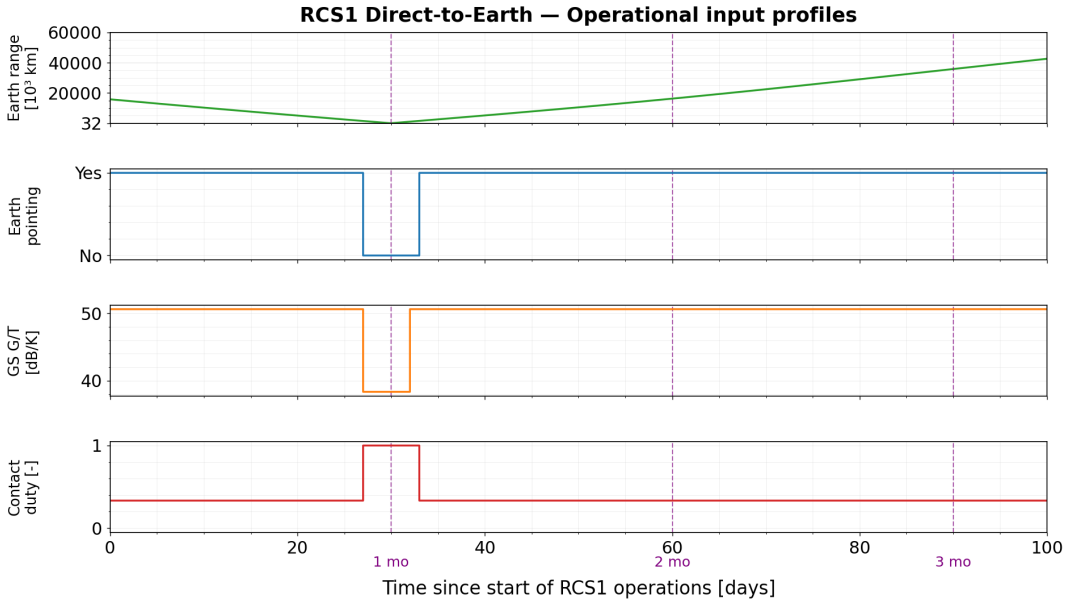


Figure 5: RCS1 Direct-To-Earth operational profile over mission time.

Figure 5 shows in the top panel the rapid decrease in Earth–Apophis range during the first month of RCS1 operations, followed by the post-encounter range increase over the remainder of the analysis window. In the other three panels, the profiles shown represent the adopted Earth-pointing capability, ground station receive-performance condition and equivalent contact duty cycle.

For the ISL case instead, both communicating spacecraft are located within the Apophis proximity environment. The link is therefore represented as a direct RF connection between RCS1 and the RAMSES spacecraft, available whenever the relative geometry allows it to be established. Unlike directive DTE communication, the ISL case is assumed not to require a dedicated pointing attitude, and is therefore compatible with simultaneous asteroid-pointing

science observations. The detailed ISL terminal and data-mode assumptions are introduced later as part of the baseline communication configuration.

These definitions establish the RCS1 DTE and ISL communication contexts adopted in the modelling framework of Chapter 3 and applied in the following performance analysis of Chapter 4.

2.2.2 RCS2 Landed CubeSat

RCS2 is the landed CubeSat of the RAMSES mission. After touchdown on Apophis, it provides local surface observations that complement those performed by the RAMSES main spacecraft and by RCS1. Its operations are concentrated around the close-encounter phase over a short battery-supported surface period, with only limited potential for extension through body-mounted solar panels. The timely return of telemetry is therefore a primary operational driver, since data acquired during the landed mission have value only if they can be returned within the available operational window.

The reference RCS2 surface-operation mission starts 20 h before closest approach, when the Earth–Apophis range is of the order of 500,000 km. Three representative surface-operation windows are considered: 20 h, 40 h and 60 h. These correspond respectively to telemetry return before closest approach, shortly after the encounter, and over an extended post-encounter surface interval. Over this period, the Earth–Apophis range decreases to approximately 32,000 km around closest approach and then increases again to approximately 900,000 km by the end of the 60 h window.

Once landed, RCS2 remains fixed at its touchdown location on the asteroid surface. It cannot manoeuvre to improve visibility or link geometry, while the rotation of Apophis continuously changes the directions of Earth and the RAMSES spacecraft as seen from the lander. The communication geometry is therefore defined by the realised landed conditions and by the asteroid rotational evolution during the short surface mission. For DTE, two conditions must be satisfied at the same time. The Earth–Apophis range must allow the link budget to close, and Earth must also be above the local horizon of the realised landing site. A favourable range interval around closest approach is therefore not sufficient by itself if the landed geometry blocks the Earth line of sight. For ISL, the relevant driver is instead the local RCS2–RAMSES geometry. The link is not governed by the Earth–Apophis range minimum, but by whether the RAMSES spacecraft is visible from the realised landing site and by the resulting local link distance and antenna geometry.

As for RCS1, a detailed public instrument timeline is not available for RCS2. The surface mission is therefore evaluated using the achievable telemetry rate and cumulative data return over the assumed surface-operation windows. This avoids introducing a detailed science schedule while retaining the short-duration character of the landed mission. Unlike RCS1, RCS2 does not have a conflict between communication and science pointing needs, because the lander attitude is fixed after touchdown.

These definitions establish the RCS2 communication cases used in the modelling framework of Chapter 3 and the performance analysis of Chapter 5.

2.3Spacecraft TT&C and Communication Architecture Fundamentals

Spacecraft communication is a mission-enabling function rather than only a data-return service. A spacecraft must transmit telemetry, receive telecommands and, where required, support tracking functions used to estimate its trajectory and operational state. These functions are commonly grouped under Telemetry, Tracking and Command (TT&C), and are essential in both nominal and contingency operations [6], [16].

For deep-space missions, TT&C becomes more demanding because the spacecraft operates outside Global Navigation Satellite Systems (GNSS) coverage and cannot rely on the near-Earth navigation support available to spacecraft in Earth orbit. Tracking must instead be supported through the TT&C radio link with ground, while the large distance to Earth introduces high free-space path loss, weak received signals, long propagation delay and Doppler effects [9], [17].

The need for robust communication affects both the operational use of the link and the communication-system architecture. At operational level, communication can be preserved or recovered by reconfiguring the link, for example by reducing the telemetry rate or increasing the transmitted power where possible. At architecture level, spacecraft commonly rely on antenna diversity and redundancy to maintain communication under degraded pointing conditions, when precise Earth pointing may not be achievable [6]. In conventional deep-space TT&C architectures, wide angular coverage can be achieved using two opposite low-gain antennas (LGAs), each connected in a one-to-one configuration to a dedicated transceiver chain. On CubeSat-class platforms, this separated architecture may not be feasible because of tight resource constraints. Multiple LGAs may instead be combined into a single RF path, with combining losses and adverse polarisation effects reducing the effective link performance.

Higher-gain antennas can improve the link performance, but they provide this improvement over a narrower angular region and usually require more demanding accommodation and pointing capabilities. The communication architecture is therefore a trade-off between robustness, achievable performance and onboard resources. For the RAMSES CubeSats, this trade-off is central: a DTE architecture provides an independent path to Earth, but requires the CubeSat to close a long-range deep-space link with limited power, constrained antenna accommodation and limited operational flexibility [6], [18].

An Inter-Satellite-Link (ISL) architecture changes the communication problem by replacing the CubeSat direct Earth link with a local space-to-space link to the main spacecraft. In this case, the CubeSat communicates over a comparatively short RF link, while the main spacecraft provides the onward communication path to Earth. This reduces the need for the CubeSat to close the long-range Earth link directly, but introduces dependence on the main spacecraft as relay, and on the local visibility and relative geometry between the spacecraft.

The following subsections introduce the two communication architectures considered in this thesis: Direct-To-Earth communication, in which the CubeSat communicates directly with Earth, and Inter-Satellite-Link communication, in which the CubeSat communicates through the RAMSES spacecraft.

2.3.1 Direct-To-Earth communication for deep-space CubeSats

In a Direct-To-Earth architecture, the CubeSat communicates directly with Earth ground stations using its own onboard radio-frequency subsystem. This provides an independent Earth-return path, but requires the CubeSat to close the full deep-space link with limited transmitter power, antenna gain, pointing capability and energy resources.

The dominant physical driver of a deep-space DTE link is propagation distance. Free-space path loss increases with range, following a square-law dependence in linear terms, making the received signal level strongly dependent on the Earth–spacecraft distance. Link feasibility depends on the complete link budget, including spacecraft transmitter power, antenna gain, coding performance and ground station receive sensitivity [8], [18]. In the RAMSES case, this range dependence is strongly time-varying because the Earth–Apophis distance changes by orders of magnitude around the 2029 close encounter, making range evolution a primary driver of DTE performance.

Deep-space DTE communication must support not only telemetry transfer. The radio link must also provide a stable RF signal for receiver acquisition and tracking, and support radiometric measurements used in flight dynamics for trajectory reconstruction. In coherent communications, the spacecraft downlink RF frequency is generated from the received uplink using a defined turnaround ratio. This allows Doppler measurements on the returned signal to be related to spacecraft motion, with the downlink frequency tied to the ground station oscillator greatly improving the performance [9].

For this reason, deep-space TT&C links commonly use standardised signal structures in which the transmitted RF power is shared between telemetry modulation, a residual carrier component and, where required, a ranging component [17], [19]. Figure 6 illustrates a representative signal for this downlink structure.

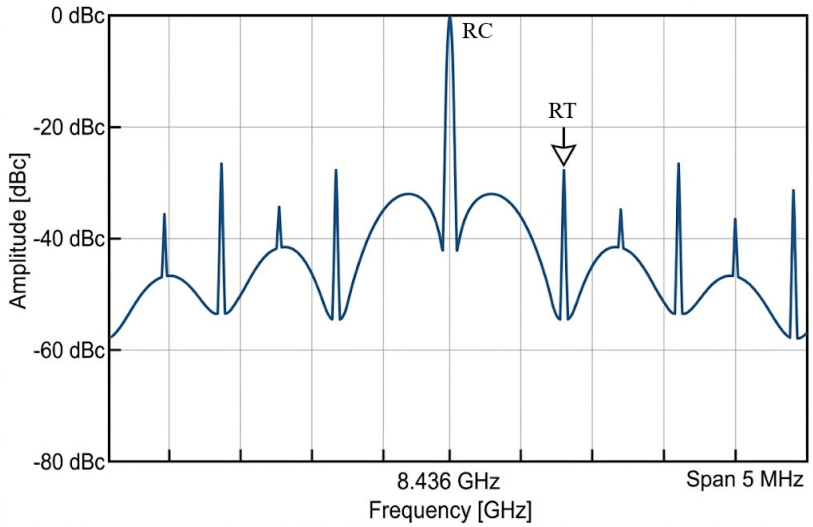


Figure 6: Representative deep-space downlink spectrum.

Figure 6 shows the RF spectrum of a representative deep-space TT&C downlink, illustrating the coexistence of residual carrier, telemetry and ranging components. The residual carrier is located at the centre of the spectrum and marked as RC, while the telemetry modulation forms the main sideband structure around it. The ranging tones, one of which is marked as RT, are placed near spectral nulls of the telemetry modulation to reduce overlap and thereby mutual interference. This signal structure allows the same downlink to support carrier acquisition and tracking, telemetry transfer and radiometric measurements for Doppler and range estimates.

Because the available RF power is shared between the residual carrier, telemetry modulation and, where present, ranging components, the signal allocation affects data recovery, carrier

tracking and radiometric support. Increasing the power assigned to telemetry can improve the data link, but reducing the residual-carrier component too far can compromise acquisition and tracking. Similarly, including a ranging component reduces the power available to the other signal components, but may be required when range measurements are part of the tracking concept. In the DTE link budget model, this power decomposition is represented through telemetry modulation-loss terms, with link feasibility assessed through separate telemetry and residual-carrier checks.

At subsystem level, these functions are implemented by the spacecraft TT&C radio system, including the transponder, RF amplification, RF distribution section and antenna interfaces [8]. The transponder is a highly specialised radio unit that supports uplink reception, downlink generation with telemetry modulation and, where required, coherent turnaround and ranging functions. Figure 7 shows a simplified representative DTE RF-chain architecture, used here to illustrate the main subsystem functions rather than to define a specific RAMSES implementation.

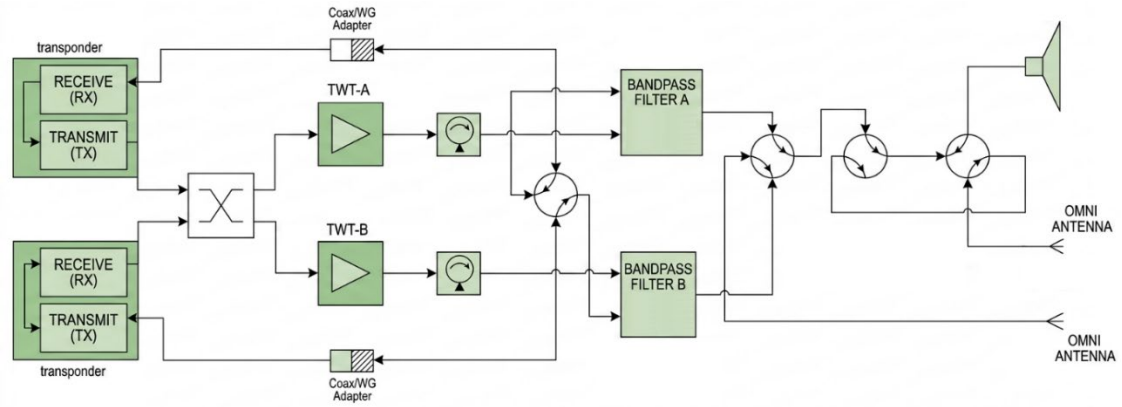


Figure 7: Representative deep-space TT&C transponder and RF-chain architecture.

Figure 7 illustrates that even a relatively simple deep-space DTE architecture involves many interconnected elements. The communication chain includes radio-frequency generation and reception, amplification, filtering, switching, RF distribution and antenna paths. Redundancy and antenna diversity add complexity, but they improve robustness when nominal pointing or link conditions cannot be maintained.

For CubeSat-class spacecraft, the same functional needs remain, but the available mass, volume, power and accommodation margins are much more limited. The DTE architecture is therefore not only a link budget problem, but also an accommodation and implementation issue, since achievable telemetry performance depends on the RF power, antenna gain,

pointing capability and RF-chain complexity that can realistically be provided within the CubeSat platform.

For a DTE telemetry link, feasibility is assessed by comparing the received signal conditions predicted by the link budget with the requirements of the selected telemetry configuration, including the applicable design margin [17], [20]. In line with the TT&C signal structure discussed above, separate checks are performed for telemetry recovery and residual-carrier tracking. The detailed DTE telemetry-performance formulation used for these checks is defined in Chapter 3, following Consultative Committee for Space Data Systems (CCSDS) and European Cooperation for Space Standardization (ECSS) link budget conventions. At architecture level, this means that DTE feasibility depends on the complete communication chain, including spacecraft RF power, antenna gain, propagation loss, ground station receiver performance, modulation, coding and carrier-tracking conditions.

DTE feasibility therefore depends not only on the spacecraft transmitter, antenna and coding configuration, but also on the ground segment. The effective receive performance of the ground station is represented by its gain-to-noise-temperature ratio (G/T), while operational availability is represented through contact windows or an equivalent contact duty cycle [21], [22].

The DTE option provides a direct and independent path from each CubeSat to Earth, but its performance depends on the combined spacecraft, propagation and ground-segment conditions. These dependencies are formalised in Chapter 3 through the DTE link budget, feasibility and data-volume modelling framework.

2.3.2 Inter-Satellite-Link communication and relay architectures

In an Inter-Satellite-Link architecture, communication is established between spacecraft rather than directly between a spacecraft and the ground segment. The link can support telemetry, telecommand and, where required, radiometric tracking support between the spacecraft involved.

ISL architectures can be implemented in different operational forms. In a peer-to-peer ISL, the link supplements the communication capability of spacecraft that already have their own direct communication paths. Its purpose is to allow direct exchange between the spacecraft, for example when low-latency transfer or tasks coordination is required. In a mothercraft–daughter-spacecraft architecture, the link is asymmetric: a smaller element communicates locally with a more capable main spacecraft, which then provides the long-range Earth-return

path through its own communication system. This second case is the more relevant analogy for RAMSES, because the CubeSats are resource-constrained deployed elements operating near a main spacecraft with its own DTE capability.

Because an ISL is established over a local spacecraft-to-spacecraft distance, its propagation loss is much lower than that of a direct Earth link. This can reduce the need for high spacecraft RF power, high antenna gains and precise pointing capability. The received signal is also normally not as weak as in a long-range deep-space Earth link, therefore the ISL does not usually require the same residual-carrier signal structure and tracking techniques discussed in Section 2.3.1. However, the spacecraft-level communication needs remain, including reliable data exchange under degraded conditions and tracking support.

For the RAMSES case, this distinction is significant because the representative ISL range is of the order of 15 km, whereas the DTE range varies from approximately 32,000 km at closest approach to many millions of kilometres before and after the encounter.

Inter-spacecraft communication is not covered by generic communication standards in the same way as conventional DTE TT&C links. ISL implementations are often proprietary, adapted from terrestrial communication solutions [23], or tailored on application-specific standards such as CCSDS Proximity-1 for orbiter–lander proximity links [24] and CUCU for International space station (ISS)–Dragon communications [25]. Because ISL applications are diverse and lack a generally applicable standardisation framework, a mission-specific ISL analysis is usually based on the reference terminal and associated protocol definition.

For small-spacecraft applications, an ISL architecture also depends on the availability of compact terminal equipment compatible with the platform resources, accommodation constraints and mission schedule. This is particularly important for CubeSat-class spacecraft, where the communication concept must be based on equipment that can realistically be accommodated and implemented. The ISL performance assessment in this thesis is therefore based on the characteristics of an identified reference terminal and its associated protocol, rather than on the DTE telemetry-performance formulation used for ground-station links.

This type of mothercraft–daughter ISL is particularly relevant to small-body missions, where deployed elements are usually more constrained in mass, volume and power, while the mothercraft already requires a DTE subsystem for Earth return. For the RAMSES CubeSats, the ISL option therefore provides a relay-based alternative to DTE: it can reduce the CubeSat RF-power, antenna-gain and Earth-pointing need, but introduces dependence on the RAMSES spacecraft and on the local relay geometry. These dependencies are formalised in

Chapter 3 through the ISL received-power, terminal-threshold and data-volume modelling framework.

2.4Reference missions and technology heritage

Most CubeSat communication heritage has been gained in Low Earth Orbit, where link distances, ground station availability, mission duration and radiation environment differ substantially from deep-space operations [6], [7]. Communication solutions commonly used for low earth orbit (LEO) CubeSats cannot therefore be assumed to scale directly to deep-space missions such as RAMSES, where larger propagation losses, more demanding TT&C functions and a less protected space environment must be considered [8].

This section reviews missions that are more representative of the RAMSES CubeSat communication problem: CubeSat or small-spacecraft Direct-To-Earth communication beyond LEO, and mothercraft–daughter-spacecraft relay architectures. These examples frame the practical heritage behind the DTE and ISL options assessed in this thesis, without being treated as direct design templates.

Lunar CubeSats also provide useful beyond-LEO small-spacecraft communication heritage [26], but their relevance to RAMSES is partial. The Earth–Moon distance is much larger than LEO distances, but remains lower by orders of magnitude than the Earth–Apophis ranges relevant to the farthest parts of the RCS1 mission. For this reason, MarCO and LICIACube are used as the primary DTE references, while lunar CubeSats are treated only as supporting background.

MarCO and LICIACube: deep-space CubeSat DTE heritage

MarCO and LICIACube are the main flown references for CubeSat or small-spacecraft DTE communication beyond LEO. Both missions operated outside the near-Earth environment and returned data through their own communication systems. They are therefore relevant to the RAMSES CubeSat DTE option as practical examples of the resources, antenna accommodation and ground segment support associated with independent small-spacecraft Earth link.

The Mars Cube One mission, MarCO, consisted of two identical 6U CubeSats launched with the NASA InSight Mars lander in 2018. The description given here is based on the detailed MarCO telecommunications subsystem documentation [27]. The MarCO objective was to demonstrate that CubeSat-class spacecraft could operate in deep space and provide relay support during the NASA InSight Entry, Descent and Landing phase. Each spacecraft was

internally non-redundant; mission-level robustness was achieved by flying two identical spacecraft, either of which could support the demonstration. MarCO is therefore an important reference for RAMSES because it demonstrates DTE communication from a 6U-class spacecraft at interplanetary distances.

Figure 8 and Figure 9, reproduced from the MarCO telecommunications subsystem documentation, show the spacecraft accommodation and deployed reflectarray antenna.

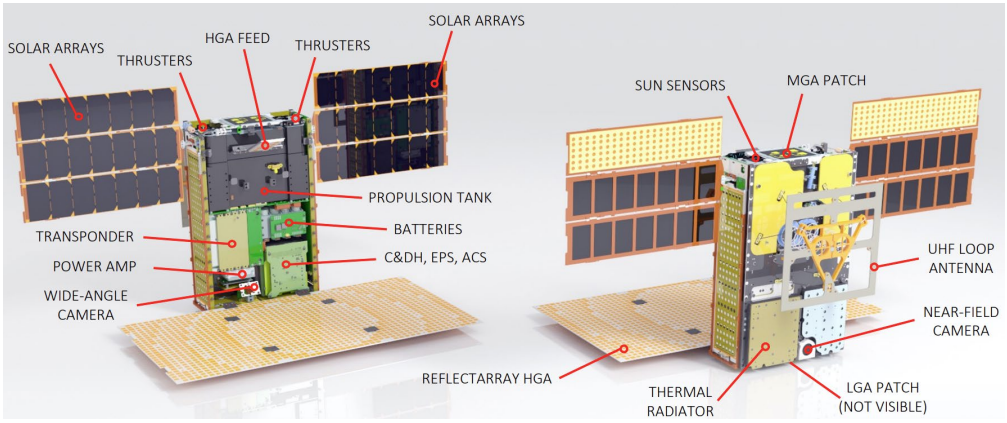


Figure 8: MarCO spacecraft accommodation. Source: [27]

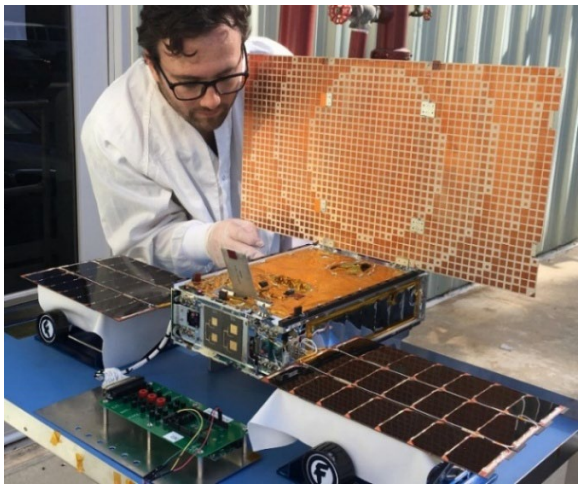


Figure 9: MarCO deployed reflectarray antenna. Source: [27]

Figure 8 and Figure 9 show that the MarCO DTE capability had a significant accommodation impact on the 6U spacecraft configuration. The internal volume was largely allocated to core spacecraft equipment and telecommunications hardware, with approximately 2U devoted to RF equipment, while the deployed configuration was dominated by the high-gain reflectarray and additional TT&C antennas. The MarCO telecommunications subsystem included a deep-space transponder, RF amplification, low-gain antennas and the deployable reflectarray, which achieved a tested X-band gain of approximately 29.2 dBi. At approximately 0.2 AU,

reported downlink data rates ranged from approximately 500 bps to 8 kbps, depending on antenna and offset-angle conditions.

For RAMSES, MarCO is therefore most relevant as a demonstrated 6U CubeSat DTE implementation, showing the high level of RF hardware and antenna accommodation needed for independent deep-space Earth return. It frames the level of resources and accommodation implications of an independent DTE solution, rather than providing a directly transferable RAMSES design.

LICIACube provides a later small-spacecraft DTE example in a planetary-defence mission context. It was carried by the DART spacecraft and released before impact with the Didymos system to observe the resulting ejecta from close proximity [28].

Figure 10 shows the main LICIACube spacecraft elements, including the communication antenna accommodation.

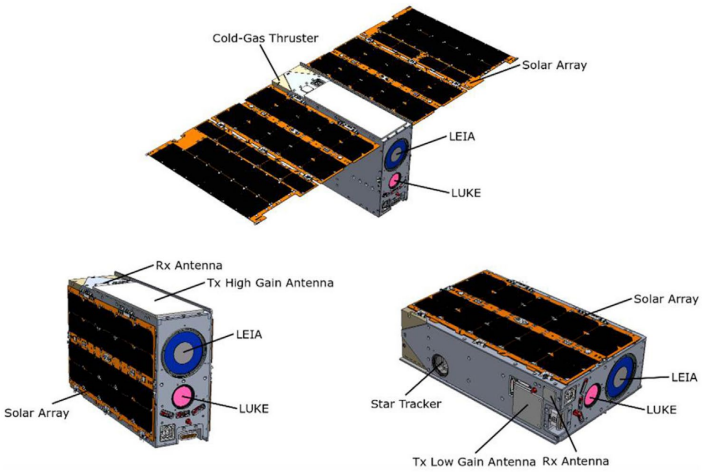


Figure 10: LICIACube spacecraft elements. Source: [28]

Figure 10 shows a more compact DTE accommodation approach than MarCO, with two low-gain antennas and a body-mounted 10 × 20 cm phased-array high-gain antenna rather than a large deployed reflectarray. Publicly available information on LICIACube communication performance is limited compared to MarCO, and its relevance is mainly at accommodation and implementation level rather than as a detailed performance benchmark. LICIACube mission indicates that a compact DTE antenna architecture can be implemented on a small spacecraft, while the RAMSES-specific performance must still be assessed through dedicated link modelling.

Both MarCO and LICIACube used Iris miniaturised deep-space transponders [27], [28], [29]. For RAMSES, however, European CubeSat-compatible deep-space TT&C equipment is

also relevant to the DTE baseline definition. A relevant example is the CubeSat Deep Space x-band Transponder (C-DST), developed to provide CCSDS- and ECSS-compatible deep-space TT&C functions within a CubeSat-compatible equipment envelope [30], [31]. Figure 11 shows the equipment configuration.

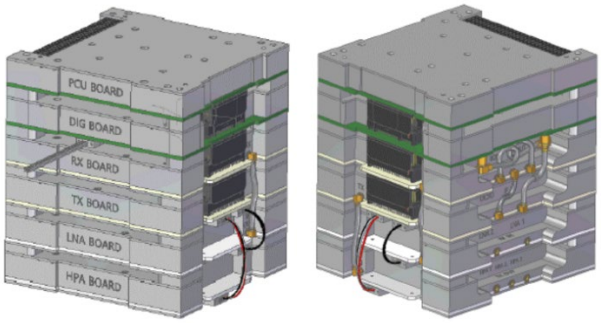


Figure 11: C-DST CubeSat Deep Space X-band TT&C transponder configuration. Source: [30].

Figure 11 shows a computer-aided design (CAD) representation of the C-DST, a modular CubeSat-compatible deep-space transponder developed by SITAEL. The C-DST configuration includes separate transponder, low-noise amplifier and high-power amplifier elements, with reported characteristics of approximately 1.5U total volume, mass slightly above 2 kg and RF output power up to 15 W [31]. For RAMSES, the C-DST provides a credible European technology reference for a CubeSat-class DTE implementation. The detailed DTE baseline assumptions used for the RAMSES CubeSat analysis are defined later in Section 2.5.

Rosetta/Philae, Hayabusa2/MASCOT and Hera: mothercraft–daughter ISL heritage

The most relevant ISL heritage for RAMSES is provided by the Rosetta/Philae, Hayabusa2/MASCOT and Hera missions. These missions all use a mothercraft–daughter-spacecraft architecture, in which a deployed element communicates locally with the main spacecraft rather than returning data directly to Earth. They therefore provide the main heritage for the relay-based ISL option considered for the RAMSES CubeSat cases.

Rosetta/Philae provides an important early example of this approach for a deployed surface element operating at a small Solar System body. The Philae lander, with a mass of approximately 100 kg, was deployed by Rosetta to the surface of comet 67P/Churyumov–Gerasimenko and communicated through the Rosetta orbiter rather than directly with Earth. The ISL was designed to support communication between the orbiter and lander over ranges

up to approximately 300 km [32]. However, the lander mass, ISL hardware scale and communication performance differ substantially from the 6U CubeSat-class accommodation considered for RAMSES, with the ROSETTA mission not providing a directly transferable design solution. Its relevance to RAMSES is therefore mainly architectural, as it demonstrates the use of a main spacecraft as the communication counterpart for a deployed surface element.

Figure 12 shows the Rosetta/Philae Inter-Satellite-Link hardware.

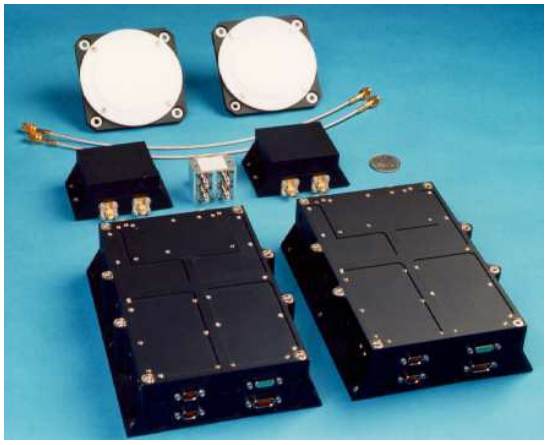


Figure 12: Rosetta/Philae Inter-Satellite-Link hardware. Source: [32]

Figure 12 illustrates the limitation of this heritage for the RAMSES CubeSat case. The hardware reflects a mission-specific relay implementation for a substantially larger lander and orbiter system, with the ISL terminal approaching the scale of a RAMSES CubeSat itself.

The Rosetta/Philae ISL system later evolved into a miniaturised lander implementation for the Japan Aerospace Exploration Agency (JAXA) Hayabusa2/DLR MASCOT mission [32]. In that case, however, the miniaturisation effort was driven by the extreme mass, volume and power constraints of a short-lived lander, with an operational lifetime of only a few hours. The resulting solution was therefore highly integrated and mission-specific. It is relevant to RAMSES as further evidence of mothercraft–lander relay communication in a small-body environment and of what may be achievable through dedicated hardware development. However, the time-critical nature of RAMSES limits the extent to which such a bespoke solutions can be considered. MASCOT therefore does not provide a directly suitable design reference for the RAMSES CubeSat platforms or for the longer RCS1 mission timeline.

Hera is the closest architectural analogue to RAMSES among the reference missions considered here. Like RAMSES, it is a deep-space small-body mission based on a main

spacecraft carrying two CubeSat-class deployed elements, Milani and Juventas, which communicate with the mothercraft through an Inter-Satellite-Link after deployment [33].

The Hera ISL is based on the GAMALINK terminal developed by TEKEVER, a European supplier, and provides a CubeSat-compatible implementation of the mothercraft–CubeSat relay concept [34], [35]. The mothercraft then provides the Earth-return path through its own DTE communication system.

After arrival at the Didymos binary asteroid system, whose secondary body Dimorphos was impacted by NASA DART, the CubeSats are released to operate in the asteroid environment. At this stage, the mission is at an Earth distance of the order of 200 million km, making independent CubeSat DTE communication significantly more demanding and supporting the use of a mothercraft–daughter relay architecture [33]. Hera is therefore particularly relevant to RAMSES because it combines a deep-space small-body mission context, two CubeSat-class deployed elements, and a CubeSat-compatible ISL implementation.

Figure 13 shows the internal accommodation of the Juventas CubeSat, one of the two CubeSats flying on Hera, including the Inter-Satellite-Link electronics.

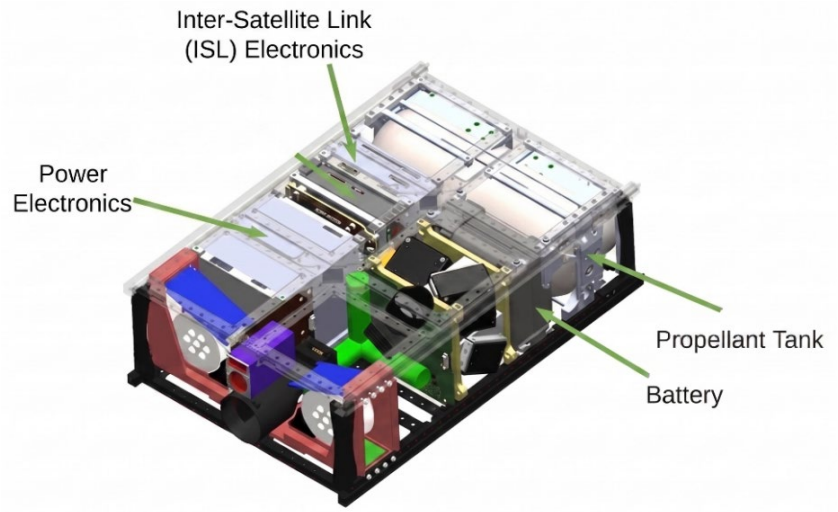


Figure 13: Juventas CubeSat internal accommodation. Source: [33].

Figure 13 shows the internal accommodation of the Juventas CubeSat, including the Inter-Satellite-Link electronics within the spacecraft configuration. The ISL hardware is accommodated alongside the spacecraft main subsystems, while approximately 3U of the 6U platform remains allocated to the payload section. This makes Hera a practical CubeSat-class example of the mothercraft–CubeSat relay architecture. However, the RAMSES target body, link geometries that include a lander, timeline and data-return requirements differ from Hera,

therefore RAMSES ISL performance must still be assessed through mission-specific modelling.

This heritage shows a progression from early small-body relay demonstrations to a CubeSat-compatible mothercraft–daughter implementation. Rosetta/Philae and Hayabusa2/MASCOT demonstrate the relay principle for deployed surface elements, while Hera provides the closest architectural and technological reference for the RAMSES ISL option. These examples support the inclusion of ISL as a credible alternative to DTE in the RAMSES CubeSat communication trade-off. Table 2.1 summarises the role of the communication heritage considered in this study, together with its applicability and main limitations with respect to the RAMSES CubeSat cases.

Table 2.1 – Summary of heritage applicable to the research

Heritage	Role in the study	Applicability to RAMSES	Main limitation
Lunar CubeSats	Supporting DTE heritage	Beyond-LEO CubeSat communication	Earth–Moon range below much of the RCS1 case
MarCO	Primary DTE reference	Independent deep-space CubeSat DTE and associated RF/antenna accommodation	Significant RF and antenna accommodation
LICIACube	Primary DTE reference	Compact DTE and body-mounted high-gain antenna accommodation	Limited public communication-performance information
Rosetta/Philae	Relay heritage	Small-body mothercraft–surface relay	Pre-CubeSat implementation
Hayabusa2/MASCOT	Relay heritage	Compact asteroid surface-element relay	Mission-specific implementation
Hera	Primary ISL reference	Closest CubeSat mothercraft–daughter architecture to RAMSES	Different mission geometry and operations

The reviewed heritage provides relevant references for the DTE and ISL options, but does not provide a directly transferable basis for selecting between the two architectures for the RAMSES CubeSats, particularly across the distinct orbiting RCS1 and landed RCS2 configurations. A mission-specific comparative assessment is therefore required to evaluate the two architectures under the conditions applicable to RAMSES.

2.5Scope, Assumptions and Analysis Requirements

This section defines the scope, assumptions and analysis requirements of the research, together with the main limitations that follow from them. The study is based on publicly

available information, relevant standards, literature sources and representative subsystem data. Where detailed mission, operational or subsystem information is not available, transparent engineering assumptions are adopted and, where relevant, their implications are examined through sensitivity or robustness analysis.

Within this boundary, the research focuses on the communication performance of the two RAMSES CubeSats under the candidate architectures. The scope is limited to telemetry return and communication feasibility. Telecommand link analysis, detailed operations planning and final flight-system design are outside the scope of the study. The baseline assumptions and input data used to support the analysis are defined in this section and then applied in the modelling framework introduced in Chapter 3.

2.5.1 Scope of the Analysis

The analysis focuses on the communication performance of the two RAMSES CubeSat elements during their main operational phases in the Apophis environment. RCS1 is the proximity-orbiting element and RCS2 the landed surface element; these configurations define the operational contexts for the DTE and ISL communication cases assessed.

Within this scope, DTE and ISL are assessed separately for RCS1 and RCS2, because the two CubeSats have major differences in their timelines and geometric constraints and roles. The analysis does not assume that both CubeSats must use the same communication architecture; the implications of common and differentiated solutions are considered in the final evaluation.

As the performance assessment focuses on telemetry return from the CubeSats, the main quantities considered are link feasibility, achievable information rates, communication coverage and cumulative data-volume. Nominal performance and selected sensitivity cases are used to assess communication robustness under the mission-specific conditions considered in the research.

Telecommand/uplink reception, detailed operations planning, final flight-system design, procurement definition and hardware verification are outside the scope of the thesis. The analysis evaluates communication capability rather than the scientific value, data prioritisation or detailed scheduling of the generated payload data.

2.5.2 Baseline Assumptions and Input Data

The analysis requires a defined set of mission, operational and communication-system assumptions and reference data. These inputs are drawn from RAMSES mission information,

ephemeris data, relevant standards, literature sources and representative subsystem data. Where CubeSat-specific design details are not publicly available, additional assumptions are introduced to define representative RCS1 and RCS2 configurations. Throughout the analysis, a distinction is maintained between published RAMSES information, heritage-derived values or configurations, and thesis-specific assumptions introduced where the available information is insufficient to define the analysis baseline. The source and type of individual baseline inputs are identified explicitly in Appendix I.

The detailed baseline assumptions and analysis requirements are provided in Appendix I containing the full numbered assumption set. They are organised into three groups: mission and geometric assumptions, communication-system assumptions and analysis requirements or implementation constraints.

Table 2.2 summarises the structure of the appendix and links each group to the corresponding appendix table.

Table 2.2 – Baseline assumption structure

Category	Appendix table	Prefix	Scope
Mission and geometry	**Table Ap.1**	**M**	Defines the RAMSES timeline, Earth–Apophis trajectory, CubeSat operational cases and landed-geometry representation.
Communication system	**Table Ap.2**	**C**	Defines the DTE and ISL link assumptions, ground-segment representation, link-performance criteria, data modes and antenna baselines.
Analysis requirements	**Table Ap.3**	**R**	Defines the resource, accommodation, technology-maturity and data-return assumptions used in the architecture assessment.

2.5.3 Simplifications and Limitations

The analysis is based on a modelling framework intended to support a preliminary performance assessment and architecture-level comparison rather than final mission design or operational planning. Representative subsystem assumptions are therefore used to consistently compare the DTE and ISL options, while maintaining a manageable level of complexity consistent with the public information available for the RAMSES CubeSat.

First, the mission geometry is represented through reference trajectories and operational configurations. RCS1 is analysed using a simplified circular-orbit representation around Apophis, rather than a full operational trajectory including dispersions, correction manoeuvres or recovery cases. RCS2 is analysed after landing; deployment, descent, touchdown dynamics and commissioning are not modelled. The results therefore describe

communication capability during the analysed operational phases, not the full CubeSat mission sequence.

Second, Apophis and its near-surface environment are represented at simplified geometric and dynamical fidelity. The asteroid is modelled using a triaxial ellipsoid and a prescribed kinematic rotation model. This captures the time-varying surface orientation needed for the communication analysis, without attempting to predict possible changes to the Apophis rotational state during or after the 2029 close encounter. Such changes remain uncertain and a broader sensitivity analysis would require additional poorly constrained parameters on top of the landed-geometry and communication variables already included. Detailed terrain, boulder-scale obstruction, local topography, surface scattering, multipath and diffraction are not modelled. Near-horizon and local-surface effects are instead represented through visibility constraints and minimum-elevation masking.

Third, the communication analysis does not model a complete operations schedule. Ground station support and contact availability are represented through the adopted profiles and duty-cycle abstractions defined in the baseline inputs. For DTE, RAMSES-compatible ground station support is assumed, but the detailed ground segment configuration, resource allocation and scheduling needed to realise simultaneous RAMSES/RCS support are not modelled. For ISL, the analysis is limited to the RCS-to-RAMSES link: RAMSES is assumed to be available as the relay endpoint according to the analysed geometry, and the RAMSES-to-Earth leg is not treated as the limiting segment for CubeSat data return.

Fourth, the quantitative communication analysis is focused on telemetry return, because this is the primary data-volume driver for the CubeSat communication architecture. Telecommand data-volumes and required data rates are expected to be much lower, and the uplink direction is not modelled as a separate link budget case, although command-link availability remains a necessary operational condition. For DTE, the CubeSat receiver is assumed to be compatible with the RAMSES deep space telecommand environment.

For ISL, the identified reference terminal provides equivalent communication capability and performance in both link directions, as specified in the supplier documentation [34], [35]. The ISL supports bidirectional communication, with transmission and reception separated through a time-division scheme rather than occurring simultaneously, using the same communication frequency and antenna for both directions. The link performance evaluated for the CubeSat-

to-RAMSES direction is therefore taken as representative of the opposite direction under equivalent link conditions.

These simplifications are consistent with the intended scope of the thesis: they allow the DTE and ISL options to be compared at performance and robustness level using the available RAMSES CubeSat information, without extending the work into final mission design, detailed operations planning or hardware verification.

2.6Chapter summary

This chapter defined the mission, communication and technology context used for the RAMSES CubeSat communication analysis. The RAMSES mission was introduced as a time-critical Apophis encounter mission involving a main spacecraft and two CubeSat elements with distinct operational roles deployed at the asteroid. RCS1 was defined as an orbiting CubeSat operating over an extended timeline, while RCS2 was defined as a landed CubeSat operating over a short interval around the close encounter.

The chapter then introduced the TT&C functions and communication architectures relevant to the analysis. DTE communication provides each CubeSat with an independent path to Earth, but requires the small spacecraft to close the full deep-space link using specialised standards, electronics and limited onboard resources. ISL communication reduces the need for the CubeSat to close the long-range Earth link directly by using a local link to RAMSES, but introduces dependence on relay geometry, availability and the mothercraft DTE system.
Relevant mission and technology heritage was reviewed to frame the communication design space. MarCO and LICIACube were identified as the main flown references for small-spacecraft DTE beyond LEO, while C-DST provided a European CubeSat-compatible deep-space TT&C reference. Rosetta/Philae and Hayabusa2/MASCOT provided mothercraft–daughter relay heritage, while Hera and its CubeSat and ISL implementation were identified as the closest architectural and technological reference for the RAMSES ISL option.

Finally, the chapter defined the scope, baseline assumptions and limitations used to frame the RAMSES CubeSat communication problem. These provide the reference context for the modelling framework developed in Chapter 3 and support Objective 1.

Chapter 3 Communication modelling framework

3.1Introduction

This chapter defines the analytical and computational framework adopted to evaluate the telecommunications performance of the RAMSES CubeSat elements (RCS1 and RCS2) during the 2029 Apophis close-encounter mission phase. The objective is to develop a validated computational framework capable of quantifying link feasibility and cumulative data return for both Direct-To-Earth (DTE) and Inter-Satellite-Link (ISL) architectures under the specific CubeSat role (orbiter or lander).

The framework is organised as a sequence of modelling layers. The chapter begins with the definition of mission geometry and reference frames, including the Earth–Apophis relative trajectory and the kinematic representation of asteroid rotation. It then defines the propagation and visibility assumptions used to distinguish geometric line-of-sight conditions from RF link performance. The DTE and ISL link budget models are introduced next, using their distinct performance metrics. The remaining sections describe the computational implementation and verification of the models, before defining the CubeSat platform, antenna and communication-architecture design space used in the later analyses.

In relation to the research objectives defined in Chapter 1, the chapter addresses Objective 1 and provides the modelling basis for the nominal-performance, statistical-characterisation and sensitivity analyses presented in Chapters 4 and 5.

3.2Mission Geometry and Reference Frames

This section defines the geometric framework used to model the relative motion of Earth, Apophis, the RAMSES spacecraft and the two CubeSats. The objective is to establish a consistent reference basis for visibility evaluation, antenna orientation and subsequent link budget computation.

The Earth–Apophis relative trajectory is generated using the Jet Propulsion Laboratory (JPL) Horizons ephemerides service through the *astroquery.jplhorizons* Python interface [36]. Detailed information on this service can be found on NASA/JPL Horizons System dedicated webpage available at URL “ssd.jpl.nasa.gov/horizons/”. State vectors for Apophis (object ID: 99942, small-body designation) and Earth (object ID: 399) are retrieved with respect to the Solar System barycentre (location code 500@0).

The closest-approach epoch is fixed at 13 April 2029 21:46:00 universal time coordinated (UTC), consistent with the adopted ephemeris and the RAMSES mission-analysis baseline [10], [36]. Cartesian state vectors are queried over a time window wider than the RAMSES CubeSat mission durations, using a 1-minute step size. Horizons outputs are returned in astronomical units and converted to kilometres using:

$$1 \text{ AU} = 149\,597\,870.700 \text{ km}$$

The numerical model is evaluated on discrete simulation epochs t, defined from the adopted RAMSES mission timelines and operational windows. The simulation time step Δt is defined as the interval between consecutive epochs.

The time discretisation is case-specific. The RCS1 analyses use mission-phase-dependent time steps, with 0.5-day sampling away from close encounter and 300 s sampling during the rapidly varying close-encounter interval, consistent with the CReMA mission-analysis baseline [10]. The RCS2 surface-operation analyses use a fixed 100 s time step over the analysed surface operations window.

Equation (3.1) defines the Earth position relative to Apophis as

$$\mathbf{r}_{EA}(t) = \mathbf{r}_{Earth}(t) - \mathbf{r}_{Apophis}(t) \tag{3.1}$$

where $\mathrm{r}_{Earth}(t)$ and $\mathrm{r}_{Apophis}(t)$ are the position vectors of Earth and Apophis, respectively, with respect to the Solar System barycentre at time (t). This gives the Earth trajectory in an Apophis-centred inertial frame.

For visualization purposes, a best-fit plane is computed for the Earth–Apophis relative trajectory using singular value decomposition (SVD). The coordinate system is then rotated so that the trajectory lies in the XY-plane. This rotation is used only for visualisation and does not affect the range or relative-direction calculations used in the subsequent analyses.

The Earth–Apophis range at a given time, $R_{EA}(t)$, is then computed from Equation (3.2) as

$$R_{EA}(t) = \| \mathbf{r}_{EA}(t) \| \tag{3.2}$$

This quantity provides the time-dependent geometric driver for the DTE free-space path loss introduced later in this chapter.

The Earth–Apophis range evolution is used as the main scalar geometric input for the DTE analysis. The time evolution of the Earth–Apophis range over the months around Closest Approach is shown in Figure 14**,** where also key markers for RCS1 and RCS2 operations are reported. A zoom-in view of the range around Closest Approach (CA) is shown in Figure 15

and used to better visualise the short-timescale geometric particularly relevant to the RCS2. The Earth–Apophis range evolution shown in Figure 14 and Figure 15 is derived from the JPL Horizons ephemerides used in the geometric model [36].

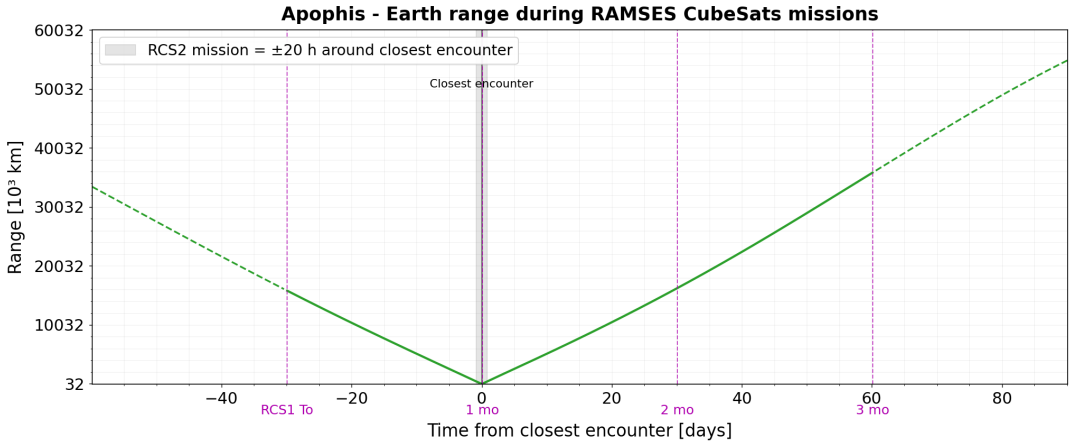


Figure 14: Apophis–Earth range during RAMSES CubeSat missions

Figure 14 shows the Earth–Apophis range evolution around the closest-approach epoch. The range profile is approximately symmetric around closest approach, but the plotted window extends further into the post-encounter phase because of the longer RCS1 analysis duration. Over this period, the range varies by several orders of magnitude, producing the strongly time-varying free-space path-loss regime that governs the DTE link.

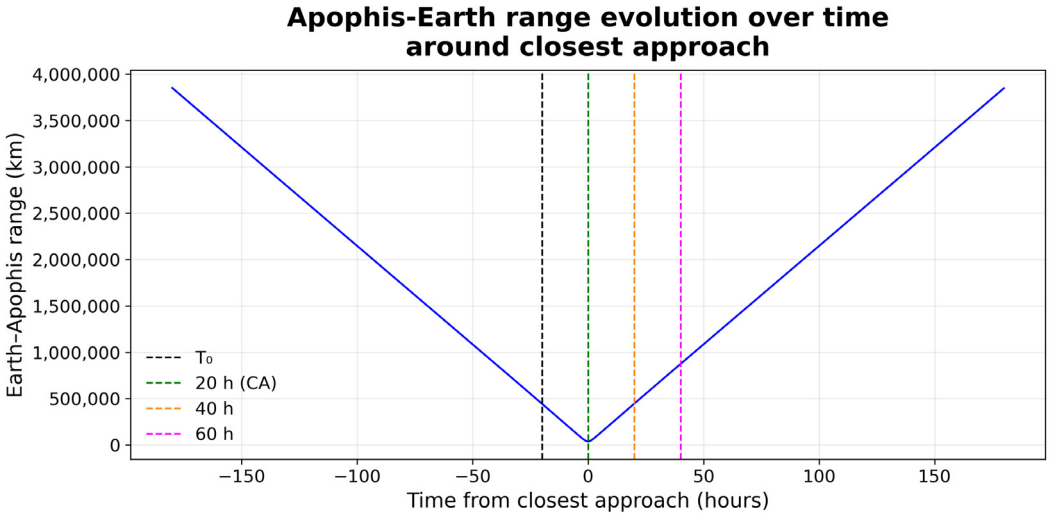


Figure 15: Earth–Apophis range over time around CA

Figure 15 shows the Earth–Apophis range evolution over the RCS2 surface-operation interval, with closest approach as the zero-time reference. At the reference landed time of 20 h before closest approach, the range is of the order of 500,000 km. The range then decreases

to its minimum at closest approach and increases again during the post-encounter surface windows at a nearly symmetric rate and approaching 900,000 km by 40 h after closest approach.

Figure 16 shows the close-encounter geometry in the rotated Apophis-centred reference frame with a 3D view, with markers for closest approach and the representative RCS2 surface-operation times.

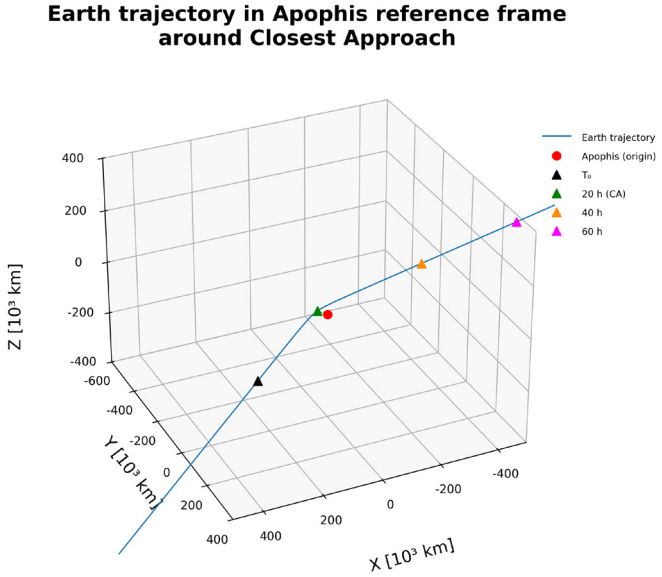


Figure 16: 3D Earth trajectory around CA (±100 h) and RCS2 Surface operations markers

The operational geometry is defined separately for RCS1 and RCS2. RCS1 is modelled in close proximity to Apophis on a circular Apophis-centred orbit with a radius of 1.2 km, consistent with the RCS1 operational baseline in Table Ap.1, M6. The RAMSES spacecraft follows a prescribed proximity trajectory around the asteroid. For the ISL analysis, the instantaneous link geometry is defined by the line-of-sight vector between the RAMSES and RCS1 positions, as illustrated in Figure 17.

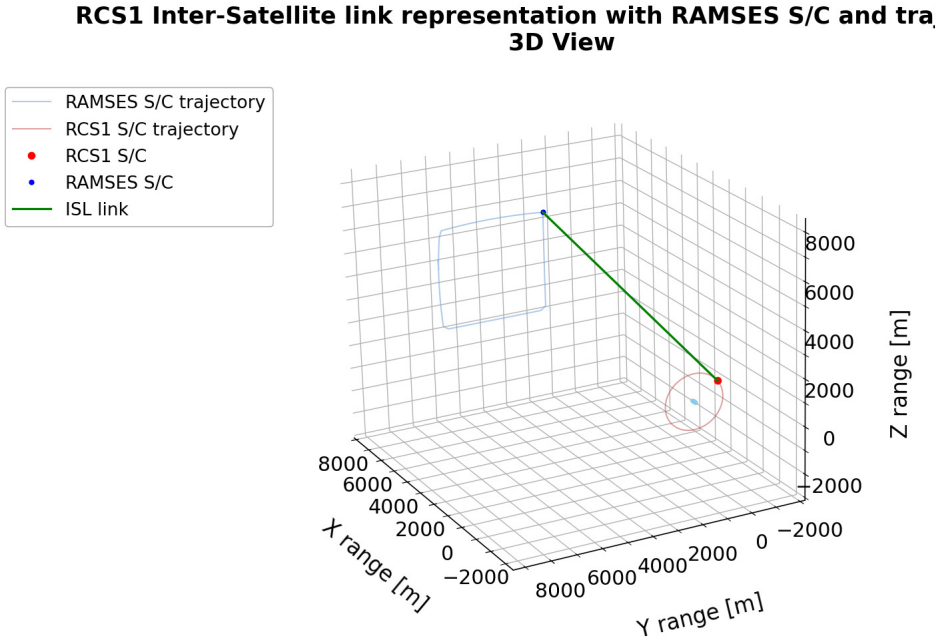


Figure 17: RCS1 ISL representation with RAMSES S/C and trajectories (3D view)

As stated in section 2.2.1, for the DTE analysis, the RCS1–Earth range is approximated by the Earth–Apophis range because the RCS1 orbital radius is small compared with the Earth–Apophis distance. Even at closest approach, the 1.2 km orbit radius is negligible relative to the minimum Earth–Apophis range of approximately 32,000 km, and the resulting range error has no meaningful effect on the DTE link-distance calculation. The RCS1 resulting range profile is therefore the one shown already in Figure 14.

RCS2 is modelled as a surface asset fixed to the asteroid surface, without burial or penetration. Its location is therefore constant in the asteroid body-fixed reference frame, while its inertial attitude evolves with the asteroid rotation. The RCS2 DTE geometry is defined by the line-of-sight vector from the landed surface point to the Earth direction in the inertial frame and varies with both Earth relative motion and asteroid rotation. For the RCS2 DTE analysis, the link distance is computed between the landing location and the ground station (G/S) on the Earth surface. Figure 18 shows, with a 3D view, the instantaneous DTE link between RCS2 and a G/S during the surface-operation window.

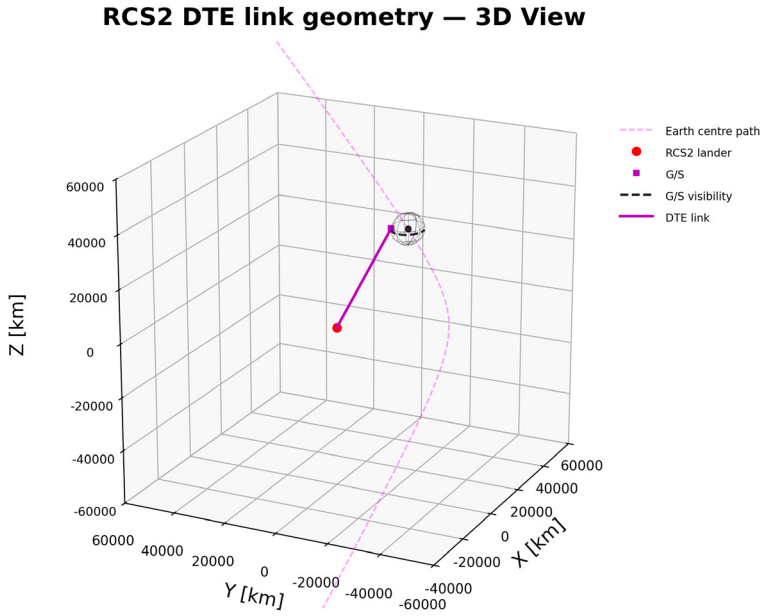


Figure 18: RCS2 DTE geometry under TS-axis rotation (3D view)

For the RCS2 ISL case, the geometry is defined by the line-of-sight vector between the landed surface point and the RAMSES spacecraft along its proximity pattern. The link geometry varies with both asteroid rotation and RAMSES spacecraft motion. Figure 19 shows the RCS2 on the asteroid surface, the RAMSES S/C proximity pattern and the instantaneous ISL link.

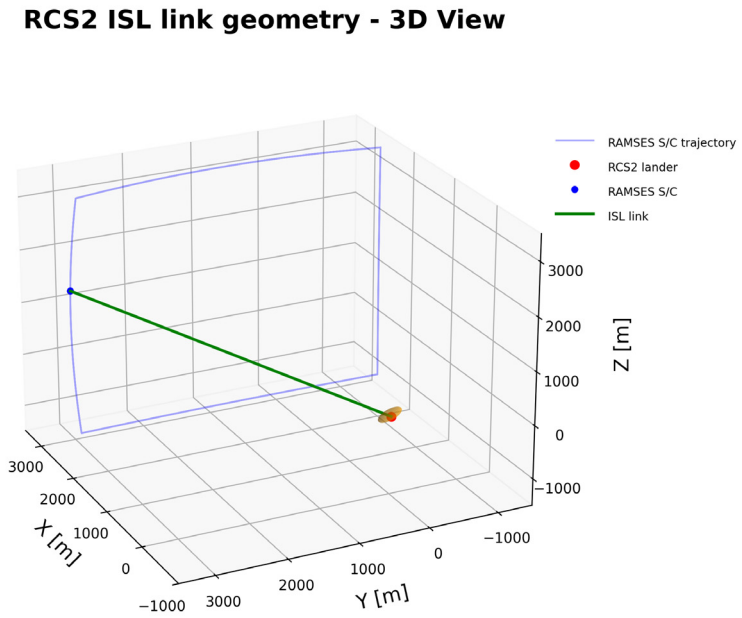


Figure 19: RCS2 ISL geometry under TS-axis rotation (3D view)

Apophis is represented in this work as a simplified triaxial ellipsoid with principal-axis lengths of approximately 375 m × 170 m × 170 m, consistent with the shape and rotation baseline defined in Table Ap.1, and with published Apophis shape/radar constraints [12], [13]. In the adopted geometric model, this corresponds to semi-axes:

$$a = 187.5 \text{ m}, \qquad b = 85 \text{ m}, \qquad c = 85 \text{ m}$$

defined in the body-fixed reference frame aligned with the principal axes of the ellipsoid as shown in Figure 20.

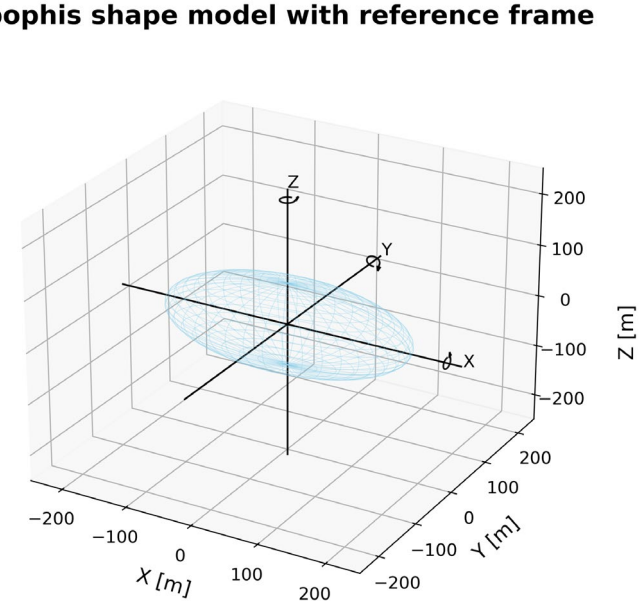


Figure 20: Apophis ellipsoidal model with body-fixed reference frame

This ellipsoidal approximation does not aim to reproduce fine-scale topography but provides a deterministic and analytically tractable surface model suitable for visibility and antenna-orientation analysis.

For RCS2, performance is evaluated statistically across the asteroid surface, as defined in Table Ap.1, M14. A Monte Carlo population of 200 candidate landing locations is generated on the ellipsoidal surface and reused consistently across all simulations. This avoids artificial bias from structured grids aligned with the asteroid principal axes or rotational phase, and reduces the sensitivity of the results to arbitrary choices of grid resolution or orientation. The body-frame coordinates of the selected landing locations are listed in Appendix II for traceability and reproducibility. The selected population size is supported by the Monte Carlo convergence assessment reported in Appendix III, where the main data volume indicators are

shown to stabilise with respect to the defined criteria. Figure 21 shows the generated candidate landing locations on the ellipsoidal asteroid surface.

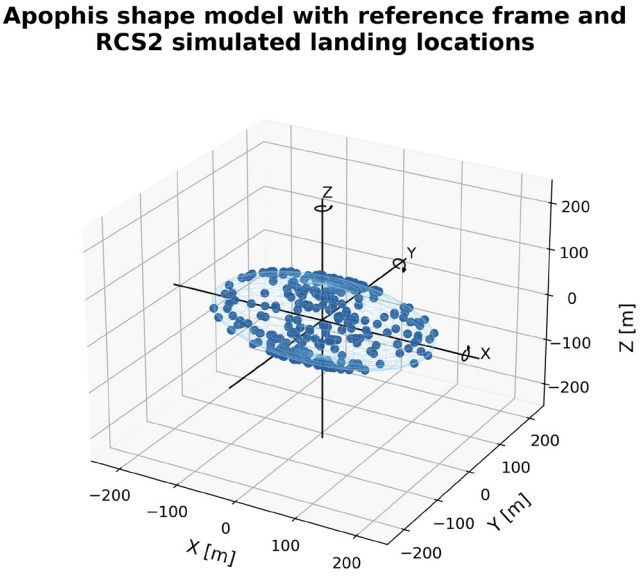


Figure 21: Apophis model with RCS2 candidate landing sites

The post-encounter rotational state of Apophis is uncertain because its spin state may evolve during the close encounter, as described in [1]. The current best knowledge of the Apophis rotational state is described in [11] and later updated in [12] where it is defined as a complex non-principal-axis (NPA) rotation, commonly described as tumbling, with uncertainties remaining especially in the phase angle.

Consistent with the scope defined in Section 2.5.3, a parametric kinematic Tilt-Spin with Precession (TS-axis) attitude model is adopted to generate repeatable surface-orientation histories for visibility, antenna-pattern and polarisation-loss calculations. The model is based on published Apophis shape and non-principal-axis rotational-state solutions [11]-[13], but is simplified for communication-geometry analysis. The asteroid rotational state is defined by three parameters:

- a principal spin motion with period

$$T_{\text{spin}} = 30.56 \text{ h}$$

  corresponding to angular rate $\omega = 2\pi/T_{\text{spin}}$

- a representative fixed tilt angle

$$\alpha = 30^{\circ}$$

defining the inclination between the instantaneous spin axis and the reference direction,

- a long-period precession term with period

$$T_{\text{prec}} = 264.0 \text{ h}$$

corresponding to angular rate $\Omega = 2\pi/T_{\text{prec}}$

At time t, the transformation from the asteroid body-fixed frame to the inertial frame is represented by the rotation matrix

$$\mathbf{R}_{TS}(t) = \mathbf{R}_z(\Omega t)\mathbf{R}_y(\alpha)\mathbf{R}_x(\omega t) \tag{3.3}$$

where $\mathbf{R}_x$, $\mathbf{R}_y$ and $\mathbf{R}_z$ denote right-handed elementary rotations about the corresponding body-frame axes. Figure 22 shows a 3D representation of the implemented kinematic rotational model.

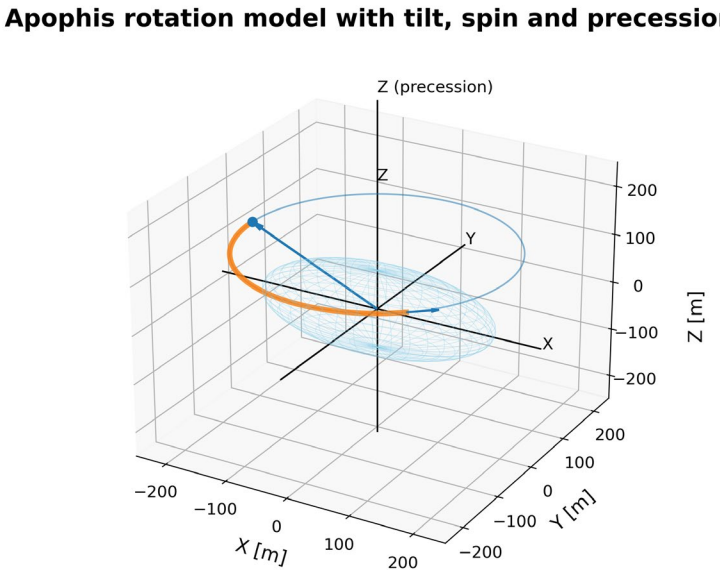


Figure 22: Apophis rotation model with tilt, spin and precession

The resulting axis-orientation evolution is propagated over the simulated timeline and used to compute the inertial orientation of surface normals and local zenith directions.

The adopted TS-axis model is purely kinematic and is not intended to reproduce the full rotational dynamics of Apophis. It provides a deterministic and controllable representation of the time-varying surface orientation needed for visibility, antenna-pattern and polarisation-loss analyses.

Being the dominant spin period of Apophis (≈ 30.56 h) of the same order as the RCS2 surface-operation duration (≈ 40 h), link visibility and antenna geometry are sensitive to the initial rotational state at the start of surface operations. An initial phase parameter ϕ_0 is therefore introduced as shown in Figure 23.

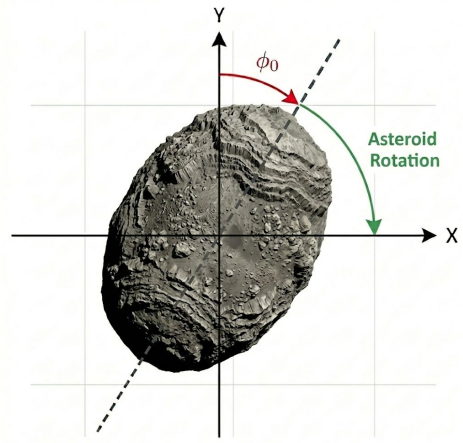


Figure 23: Initial Apophis rotational phase offset

The phase ϕ_0 defined as the initial rotation angle of the body-fixed (X)-axis projected onto the inertial (XY)-plane at the beginning of the simulation window. In the rotation sequence, this initial phase is represented by a rotation about the body-fixed (z)-axis, giving the orientation operator in Equation (3.4)

$$\mathbf{R}(t, \phi_0) = \mathbf{R}_{TS}(t)\, \mathbf{R}_{\text{init}}(\phi_0) \tag{3.4}$$

with $\phi_0 \in [0°, 360°]$ and sampled in $5°$ increments in this research. This allows the analysis to represent uncertainty or sensitivity to the initial rotational state of Apophis at the start of the RCS2 surface operations.

The inertial position of a surface point and the corresponding surface normal are obtained from Equation (3.5) by applying the asteroid rotation operator to the body-fixed vectors

$$\mathbf{x}(t) = \mathbf{R}(t, \phi_0)\mathbf{x}_{body} \tag{3.5}$$

where $\mathbf{x}_{body}$ represents either the surface point or its outward normal expressed in the body-fixed frame.

3.3Propagation assumptions and geometrical visibility

This section defines how geometric visibility is determined for DTE and ISL links before RF link budget modelling is applied. Visibility is treated as a purely geometric condition,

based on line-of-sight and elevation constraints, and is kept separate from the RF feasibility conditions introduced later in Section 3.5.

For RCS2 surface operations, visibility is determined from the zenith angle θ_z. This angle is defined between the local surface normal, or local zenith direction, and the direction to the communication counterpart: Earth for DTE and RAMSES for ISL. Figure 24 illustrates this geometry for both communication cases.



Figure 24: Zenith angle θ_z definition

A zenith angle of 0° corresponds to the overhead direction, while zenith angles $\geq 90^\circ$ correspond to direction below the local horizon. Although the small dimensions of Apophis can result in a geometric horizon extending beyond a zenith angle of 90° for an observer above the surface, the RCS2 antenna altitude relative to the surface is sufficiently small for this effect to be considered negligible. The local horizon is therefore approximated by the tangent plane at the landing location, corresponding to a zenith angle of 90°.

The theoretical geometric visibility condition for either a DTE or ISL link is that the zenith angle satisfy the expression:

$$\theta_z < 90^\circ$$

The far-field criterion is based on the Fraunhofer distance, defined in [37], Chapter 3 Section 3, and is computed according to Equation (3.6) as:

$$R_{\mathrm{FF}} = \frac{2D^2}{\lambda} \tag{3.6}$$

where D is the maximum antenna dimension and λ the wavelength.

For the RAMSES cases, the resulting far-field distances are several orders of magnitude smaller than the operational link ranges, which are of the order of kilometres for ISL and at

least 10^4 km for DTE. Antenna gain-based modelling with plane-wave propagation is therefore applicable.

As further discussed in [37], Chapter 8 Section 2.2, to be in free-space propagation conditions, the first Fresnel zone must be unobstructed. The first Fresnel radius at an intermediate point between transmitter and receiver is computed from Equation (3.7) as

$$r_F = \sqrt{\frac{\lambda d_1 d_2}{d_1 + d_2}} \tag{3.7}$$

where λ is the wavelength, d_1 and d_2 denote the distances from the considered point to the transmitter and receiver, respectively.

Figure 25 shows the RCS2 to RAMSES S/C ISL geometry case, including the line-of-sight height above the surface Z and the elevation angle ∈.

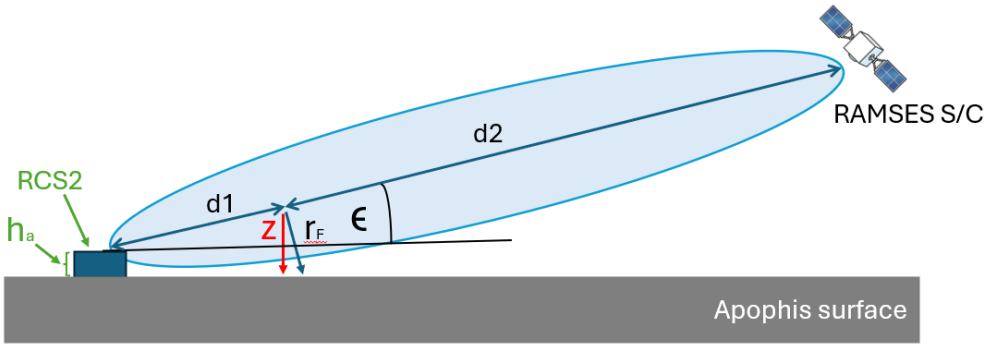


Figure 25: RCS2 first Fresnel ellipse

For the RCS2 surface cases, the antenna phase centre is assumed to be located at a height $h_a = 20cm$ above the local surface. For a line of sight with elevation angle ϵ, the height of the ray above the local tangent plane at a d_1 from RCS2 is given by Equation (3.8) as

$$z(d_1) = h_a + d_1 \sin \epsilon \tag{3.8}$$

where d_1 is measured along the horizontal projection of the link path.

To ensure that the lower boundary of the first Fresnel zone remains above the surface along the full path, Equation (3.9) imposes the condition:

$$z(d_1) - r_F(d_1, D - d_1) \geq 0 \forall\, d_1 \in [0, D] \tag{3.9}$$

where D is the total link distance.

Substituting the Fresnel radius from Equation (3.7) into Equation (3.9) gives the Equation (3.10):

$$h_a + d_1 \sin\epsilon \geq \sqrt{\frac{\lambda\, d_1 (D - d_1)}{D}} \forall\, d_1 \in [0, D] \tag{3.10}$$

Requiring this condition to be satisfied for all intermediate points gives the conservative minimum-elevation condition defined in Equation (3.11) as

$$\sin\epsilon_{\min} \geq \frac{\lambda}{4h_a} - \frac{h_a}{D} \tag{3.11}$$

and therefore, the minimum elevation angle is obtained from Equation (3.12) as

$$\epsilon_{\min} \geq \arcsin\left(\frac{\lambda}{4h_a} - \frac{h_a}{D}\right) \tag{3.12}$$

With the RCS2 ISL case assumptions

$$D = 5\,\text{km}, \qquad f = 2.4\,\text{GHz},\ h_a = 0.20\,\text{m}$$

$$\lambda = \frac{c}{f} \approx 0.125\,\text{m}$$

Then

$$\epsilon_{\min} \approx \arcsin\left(\frac{0.125}{4 \cdot 0.20} - \frac{0.20}{5000}\right) \approx \arcsin(0.1562) \approx 9.0^\circ \tag{3.13}$$

A minimum elevation of approximately 9° is therefore sufficient to maintain first-Fresnel-zone clearance for the assumed RCS2 ISL local geometry.

For the DTE case, the much longer link distance and higher carrier frequency reduce the corresponding Fresnel-clearance elevation requirement. The local-clearance constraint is therefore governed by the ISL case.

The visibility condition is therefore expressed using a minimum-elevation mask, defined in Equation (3.14) as

$$\theta_z < 90^\circ - \varepsilon_{\min} \tag{3.14}$$

where $\varepsilon_{\min}$ is the minimum elevation angle, as shown in Figure 26.

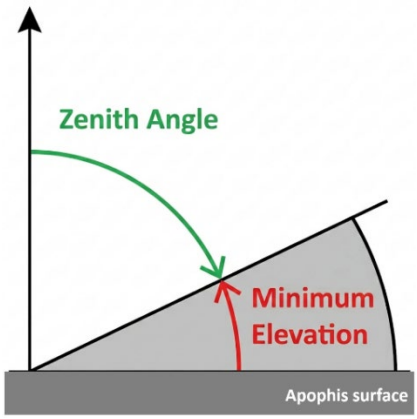


Figure 26: Conceptual representation minimum elevation

A minimum-elevation constraint represents local horizon masking and near-horizon degradation in the absence of a detailed terrain and electromagnetic surface model [21]. At low elevation angles, the link becomes more sensitive to local obstruction, Fresnel-zone blockage, multipath and surface-scattering effects. An elevation mask therefore provides a first-order way to represent these effects without modelling the local terrain explicitly, as also discussed for planetary surface communication cases in [37], [38].

For Apophis, the relatively small and elongated shape of the asteroid implies that the local horizon is close and strongly site-dependent. In the absence of a high-fidelity local terrain model and a dedicated electromagnetic scattering model, a baseline minimum elevation of 10° is adopted as a conservative modelling assumption, consistent with the first-Fresnel-clearance result derived above.

In addition to the minimum-elevation constraint, local surface inclination modifies the effective zenith direction and therefore the antenna boresight orientation. At each candidate landing site, a local surface reference frame is defined from the ellipsoidal asteroid geometry. The outward-pointing surface normal defines the local zenith direction, denoted $\hat{z}$.

Two additional unit vectors are then constructed in the tangent plane to the local surface. These vectors, denoted $\hat{e}_1$ and $\hat{e}_2$, are orthogonal to $\hat{z}$ and mutually perpendicular. The triad, $(\hat{e}_1, \hat{e}_2, \hat{z})$, therefore forms a right-handed orthonormal coordinate system centred at the landing site.

The landed spacecraft attitude is defined relative to this local surface frame. In addition to landing location and initial rotational phase, the spacecraft may have an azimuthal heading about the local zenith axis. A heading parameter ϕ is therefore introduced to rotate the spacecraft body frame within the local tangent plane. This rotation does not change the local

zenith direction, but changes the azimuthal orientation of the antenna pattern, relative to the surface frame. In the nominal analysis of Chapter 5, the heading is treated as an uncontrolled realised condition. In the sensitivity analysis, it is used to represent possible controlled landing capability over the final landed azimuthal orientation.

In the nominal landed configuration, the antenna boresight is aligned with the local zenith. The tilt/slope sensitivity case represents local surface-orientation uncertainty by rotating the boresight away from the zenith by a prescribed tilt magnitude δ within the local tangent plane. For a given tilt direction, the tilted boresight vector is defined by Equation (3.15) as

$$\mathbf{b}_{\text{tilted}} = \cos\delta\,\hat{\mathbf{z}} + \sin\delta\,\hat{\mathbf{e}}_t \qquad \hat{\mathbf{e}}_t \in \{\pm\hat{\mathbf{e}}_1, \pm\hat{\mathbf{e}}_2\} \tag{3.15}$$

Tilt magnitudes of $\delta = 15^\circ$ and $\delta = 30^\circ$ are adopted as mild and severe perturbations. For each landing site and tilt magnitude, four orthogonal tilt directions are evaluated and treated as equally probable, $\pm\hat{\mathbf{e}}_1$ and $\pm\hat{\mathbf{e}}_2$. Visibility and link budget quantities are then computed with respect to the tilted boresight direction, as illustrated in Figure 27.

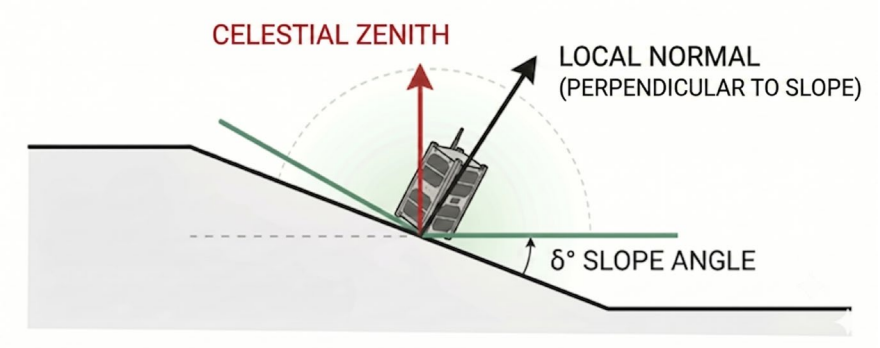


Figure 27: Effect of local surface slope on zenith-angle definition

The tilt/slope model is not intended to reproduce or replace missing detailed topography models but to quantify sensitivity to plausible local slope-induced antenna misalignment.

For RCS1 orbital operations, the DTE link may be interrupted when Apophis blocks the line of sight between the spacecraft and Earth. Because the Earth–Apophis distance is much larger than both the asteroid dimensions and the RCS1 orbital radius, the Earth direction can be treated as quasi-fixed over a single RCS1 orbit. This geometry is illustrated in Figure 28.

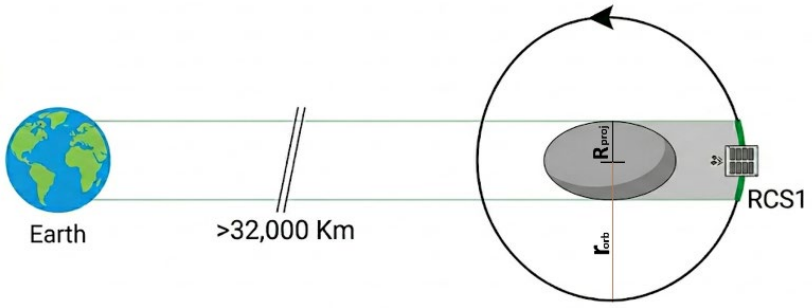


Figure 28: Geometric occultation of Earth by Apophis for RCS1 orbiter

The occurrence and duration of occultation would depend on the orientation of the RCS1 orbital plane relative to the Earth direction. As Apophis moves along its heliocentric orbit, this relative geometry would vary with time, producing a time-dependent occultation pattern. In this research, a deliberately conservative assumption is adopted: the RCS1 orbital plane is assumed to remain aligned with the Earth direction, so that an occulted arc is present on every orbit whenever the spacecraft passes behind the asteroid limb.

Consistent with the RCS1 operational and DTE-geometry baselines in Table Ap.1, M6 and M8, the RCS1 orbit is modelled as circular with uniform angular velocity. Under this approximation, the occulted part of each orbit is represented by the angular sector subtended by the asteroid projection onto the orbital path. For an effective projected radius R_{proj}, the occultation fraction is computed from Equation (3.16) as

$$f_{\text{occ}} = \frac{1}{\pi}\arcsin\left(\frac{R_{\text{proj}}}{r_{\text{orb}}}\right) \tag{3.16}$$

Bounding R_{proj} between the smallest and largest semi-axes of the ellipsoid (85 m and 185 m) gives

$$f_{\text{occ}} \approx 2.3\% \text{ to } 4.9\%$$

The median value is adopted in this research as a representative conservative approximation

$$f_{\text{occ}} = 3.5\%$$

The corresponding occultation penalty factor is obtained from Equation (3.17) and is applied multiplicatively to the nominal DTE visibility time.

$$\eta_{\text{occ}} = 1 - f_{\text{occ}} = 0.965 \tag{3.17}$$

3.4Link Budget, Data-volume Models for Direct-To-Earth and Inter-Satellite-Links

This section defines the link-performance and data-return models used for the RAMSES CubeSat DTE and ISL cases. Both architectures are assessed using the same general logic: link feasibility is evaluated over the mission timeline, the applicable telemetry or data mode is selected from a discrete set, and the resulting data return is accumulated over the available communication periods.

The two architectures differ in how link performance is evaluated. For DTE, the assessment uses a conventional space-to-ground telemetry link budget, consistent with the standards and link budget references introduced in Section 2.3.1. For ISL, the assessment is based on the selected reference terminal and its supplier-provided performance data, in line with the relay-link context introduced in Section 2.3.2.

Direct-To-Earth link modelling

The DTE model computes the achievable telemetry performance and returned data-volume as functions of mission time, accounting for spacecraft antenna configuration, propagation effects, ground segment performance, and modulation and coding assumptions. Following the formulation described in [18], Chapter 5, the telemetry performance is expressed by Equation (3.18) as

$$\frac{E_b}{N_0} = \frac{C}{N_0} - 10\log_{10}(R_b) \tag{3.18}$$

where C/N_0 is the carrier-to-noise density ratio expressed in dBHz and R_b is the useful information bit rate in bits per second. In the link budget formulation used here, C/N_0 is used as the standard received-signal reference quantity from which the telemetry E_b/N_0 is derived. The resulting E_b/N_0 is the performance quantity used to assess the feasibility of each candidate telemetry rate.

For each simulation epoch, C/N_0 is computed from Equation (3.19), following the standard downlink link budget formulation described in [39], Chapter 6.

$$\frac{C}{N_0} = \mathrm{EIRP_{SC}} - L_{\mathrm{FSPL}} - L_{\mathrm{atm}} - L_{\mathrm{pol}} - L_{\mathrm{mod,\,TM}} + \left(\frac{G}{T}\right)_{\mathrm{GS}} - L_{\mathrm{impl}} - k \tag{3.19}$$

with all terms expressed in decibels. The spacecraft equivalent isotropic radiated power, $\mathrm{EIRP_{SC}}$, includes the RF high-power-amplifier (HPA) output power after onboard insertion and mismatch losses, antenna gain and any applicable pointing losses.

Atmospheric losses, L_{atm}, are included as a constant loss term covering gaseous, rain and cloud attenuation. The value is derived according to [40] under adverse conditions, as a simplifying assumption; the difference between favourable and adverse atmospheric cases is marginal for the RAMSES link budgets. A fixed implementation loss is also included to represent the demodulation-performance degradation of the actual ground station modem relative to an ideal theoretical receiver, following the approach described in [22]. The final k term is the Boltzmann constant, with value -228.60 dB.

Consistent with the propagation assumptions defined earlier and following the formulation defined in [18], Chapter 4, the free-space path loss (FSPL) is computed using Equation (3.20)

$$L_{\text{FSPL}} = 20log_{10}(R) + 20log_{10}(f) + 92.45\ dB \tag{3.20}$$

where R is the Earth–spacecraft range in kilometres, f is the downlink frequency in gigahertz and the 92.45 dB term is the unit-conversion constant for this form of equation.

The polarisation mismatch loss, L_{pol}, is computed according to [39], Chapter 6, using Equation (3.21)

$$L_{\text{pol}} = 10\log_{10}\left(\frac{(1+\rho_{\text{SC}}^2)(1+\rho_{\text{GS}}^2)}{(\rho_{\text{SC}}+\rho_{\text{GS}})^2}\right) \tag{3.21}$$

where ρ_{SC} and ρ_{GS} are the spacecraft and ground station antenna axial ratios, respectively. The ground station ρ_{GS} is represented using the fixed axial-ratio value from [14], consistent with the constantly pointed large aperture antenna toward the RAMSES CubeSat. The spacecraft axial ration ρ_{SC} depends on the specific case and is defined from the RCS1 and RCS2 antenna assumptions described later in this section.

For telemetry recovery, the modulation loss represents the apportionment of the total RF power between the telemetry, ranging and residual-carrier components. It is computed using the definition given in [39], Chapter 6, by Equation (3.22)

$$L_{\text{mod, TM}} = -[10\log_{10}(\sin^2 m_{\text{TM}}) + 20\log_{10} | J_0(m_{\text{RNG}}) |] \tag{3.22}$$

where m_{TM} and m_{RNG} are respectively the telemetry and ranging modulation indices. $J_0(\cdot)$ is the zeroth-order Bessel function.

The link budget result is first expressed as a rate-independent normalised quantity, corresponding to the E_b/N_0 that would be available at a reference information bit rate of 1 bps. This quantity is defined by Equation (3.23).

$$\left(\frac{E_b}{N_0}\right)_{\text{norm}} = \frac{C}{N_0} \tag{3.23}$$

The actual E_b/N_0 for a telemetry mode with information bit rate R_b is then obtained from Equation (3.24), consistent with the standard link budget formulation described in [39], Section 6.6.5.

$$\frac{E_b}{N_0} = \left(\frac{E_b}{N_0}\right)_{\text{norm}} - 10\log_{10}(R_b) \tag{3.24}$$

The rates R_b applied in the modelling are taken from the DTE TM-rate set defined in the communication-system baseline in Table Ap.2, C3, and reported in full in Appendix IV.

Each telemetry rate is associated with a required value $(E_b/N_o)_{req}$

$$M = \frac{E_b}{N_0} - \left(\frac{E_b}{N_0}\right)_{\text{req}} \tag{3.25}$$

Imposing the design-margin requirement through Equation (3.26)

$$M \geq M_{\text{target}} \tag{3.26}$$

gives the feasibility condition given by Equation (3.27)

$$\frac{E_b}{N_0} - \left(\frac{E_b}{N_0}\right)_{\text{req}} - M_{\text{target}} \geq 0 \tag{3.27}$$

where $M_{\text{target}} = 3$ dB is adopted as the nominal deep-space telemetry design margin, as recommended in [17], Chapter 8.

Among the feasible telemetry modes, the selected telemetry rate R_{sel} at epoch t, is determined using Equation (3.28)

$$R_{\text{sel}}(t) = \max_{i \in \mathcal{F}(t)} R_i \tag{3.28}$$

where:

- R_i is the information bit rate associated with telemetry mode i
- $\mathcal{F}(t)$ is the set of feasible telemetry modes at epoch t, defined as the modes satisfying the margin condition in Equation (3.27).

In addition to telemetry feasibility assessment, a residual carrier-only C/N_0 is computed separately. In this case, C denotes the residual carrier power component of the composite phase-modulated downlink. This component is obtained by separating the carrier contribution from the telemetry and ranging modulation components using the applicable modulation

indices. The corresponding residual-carrier modulation loss is defined in [39], Chapter 6, and computed by Equation (3.29)

$$L_{mod,C} = -[10\log_{10}(\cos^2 m_{TM}) + 20\log_{10} \mid \mathrm{J}_0(m_{RNG}) \mid] \quad (3.29)$$

The residual-carrier C/N_0 is then computed using Equation (3.30). The same link budget structure used for the telemetry case in Equation (3.19) is applied, with $\mathrm{L}_{\mathrm{mod,\,TM}}$ replaced by $\mathrm{L}_{\mathrm{mod,\,C}}$

$$\left(\frac{C}{N_0}\right)_{carrier\text{-}only} = EIRP_{SC} - L_{FSPL} - L_{atm} - L_{pol} - L_{mod,C} + \left(\frac{G}{T}\right)_{GS} - L_{impl} - k \quad (3.30)$$

This quantity is computed over the simulated timeline to evaluate the feasibility of carrier recovery, which is required for telemetry recovery.

In addition to the residual carrier C/N_0, an indicative measure of carrier-tracking robustness is derived by estimating the maximum Phase-Locked Loop (PLL) tracking bandwidth. A larger $B_{\mathrm{PLL,\,max}}$ indicates greater residual-carrier margin and a less restrictive tracking condition supported by the available carrier power. This quantity is used for interpretation only and does not affect link feasibility in this thesis.

The maximum PLL noise bandwidth is estimated from the excess residual-carrier C/N_0 above the carrier-tracking threshold, using the PLL-bandwidth formulation in Equation (3.31), following [39], Section 5.6

$$B_{\mathrm{PLL,max}} = 10^{\frac{C/N_0 - C/N_{0,\mathrm{thr}}}{10}} \quad (3.31)$$

where $C/N_{0,\mathrm{thr}} = 17$ dBHz is the carrier-tracking threshold as defined in [14], Chapter 4. The resulting $B_{\mathrm{PLL,\,max}}$ provides a first-order indication of the carrier-tracking capability of the receiver.

The equations above define the generic DTE telemetry and residual-carrier feasibility checks. In the RAMSES CubeSat cases, the same formulation is applied using the case-specific geometry, antenna, pointing and ground station assumptions defined in Sections 2.2.1 and 3.6.2. The resulting feasible telemetry rate is then used in the data-volume calculation introduced below.

DTE data-volume is computed by integrating the selected useful information bit rate over the applicable time interval. Since the selected telemetry rates are expressed as useful information bit rates, R_{b}, the effect of the coding rate is already included in the data-rate definition. Differences between the Turbo 1/2 and Turbo 1/4 options are therefore reflected

directly in the computed cumulative data-volume, consistent with standard telemetry-system engineering practice [16]. Ground station contact time is represented through an effective contact duty cycle $D(t) \in [0,1]$, rather than through an explicit contact/gap time series, to avoid embedding a specific DTE operations concept.

The incremental and cumulative DTE data-volumes are computed using Equations (3.32) and (3.33), respectively.

$$\Delta V = R_b \cdot \Delta t \cdot D(t) \tag{3.32}$$

$$V_{\text{cum}} = \sum \Delta V \tag{3.33}$$

where Δt is the simulation time step introduced in Section 3.2 and $D(t)$ is the effective fraction of that time step during which communication is available.

For RCS1, $D(t)$ is derived from the effective contact-duty-cycle profile defined in the DTE ground station support baseline in Table Ap.2, C4. The RCS1 occultation penalty defined in Equation (3.17) is then applied as a fractional reduction of effective communication time.
For RCS2, ground station support is assumed continuously available, in line with Table Ap.2, C4. The RCS2 communication opportunity is then governed by the landed CubeSat-to-Earth geometries and resulting visibility.

Inter-Satellite-Link modelling

The ISL model evaluates a short-range link between the CubeSat and the RAMSES spacecraft. Performance is assessed directly in terms of received power relative to receiver sensitivity thresholds, as defined by the supplier of the reference terminal. At each simulation epoch, the received power is computed following standard satellite link budget practice as described in [18], Chapter 4, and according to Equation (3.34)

$$P_{\text{RX}} = P_{\text{TX}} + G_{\text{TX}}(t) + G_{\text{RX}}(t) - L_{\text{FSPL}}(t) - L_{\text{onboard}} - L_{\text{pol}}(t) \tag{3.34}$$

where P_{TX} is the CubeSat transmit power, G_{TX} and G_{RX} are the transmit and receive antenna gains, L_{onboard} represents constant onboard insertion losses (e.g. RF cable, connectors, combiner/splitter), and L_{pol} is the polarisation mismatch loss, computed with the same formulation defined in the DTE case by Equation (3.21). Because the ISL is a space-to-space link, no atmospheric losses are included.

Under the same free-space propagation model assumption, the ISL free-space path loss is computed using Equation (3.35), also in this case formulated in [18], Chapter 4.

$$L_{\mathrm{FSPL}} = 20\log_{10}(R_{\mathrm{km}}) + 20\log_{10}(f_{\mathrm{MHz}}) + 32.44 \tag{3.35}$$

where R_{km} is the inter-satellite distance in kilometres and f_{MHz} is the ISL carrier frequency in megahertz.

ISL link feasibility is evaluated against a discrete set of data modes, each characterised by a specific data rate and a receiver sensitivity threshold. ISL mode feasibility is evaluated through Equation (3.36)

$$P_{\mathrm{RX}} \geq P_{\mathrm{sens},m} + M_{\mathrm{ISL}} \tag{3.36}$$

where $P_{\mathrm{sens},m}$ is the sensitivity threshold of data mode m defined in the communication-system baseline in Table Ap.2, C12, and reported in full in Appendix IV.

The reference ISL transceiver is able to adapt the data mode in use, autonomously detecting the faster rate satisfying the link conditions [34], [35] and the implementation margin, M_{ISL}, is set to 1 dB.

As the ISL waveform does not involve apportionment of RF power to different signal types (residual-carrier, Telemetry, Ranging), no modulation losses or residual-carrier recovery margin are computed.

The instantaneous and cumulative ISL data-volumes are computed using Equations (3.37) and (3.38), without the introduction of a duty-cycle abstraction as is the case for the Direct-To-Earth:

$$\Delta V = R_b \cdot \Delta t \tag{3.37}$$

$$V_{\mathrm{cum}} = \sum \Delta V \tag{3.38}$$

Periods of geometric outage naturally result in zero contribution to data-volume.

The equations above define the common ISL received-power, mode-feasibility and data-volume formulation, using Equations (3.34), (3.36), (3.37) and (3.38). Their application differs between RCS1 and RCS2 because the two cases use different geometry inputs and visibility conditions.

For RCS1 ISL, the inter-spacecraft range and off-boresight profiles are taken from the external mission-analysis inputs described in Section 3.2. These profiles define the link distance and angular separation from the antenna boresight, but not the azimuthal orientation

of the antenna pattern. Multiple antenna-azimuth cuts are therefore sampled to avoid making the result dependent on a single arbitrary pattern cut.
For RCS2 ISL, the model uses the landed CubeSat-to-RAMSES geometry defined in Section 3.2 and the visibility conditions defined in Section 3.3. The link is available only when the RAMSES spacecraft is above the local minimum-elevation mask. During those visible intervals, the same ISL formulation is applied using the instantaneous range, antenna gain and polarisation mismatch terms.

This provides a consistent basis for comparing the DTE and ISL cases while keeping the performance metrics appropriate to each link type. The differences between RCS1 and RCS2 are represented through their case-specific geometry, pointing and uncertainty assumptions, without imposing a detailed operational schedule.

3.5 Implementation Framework and Computational Architecture

The modelling described in Sections 3.2 to 3.4 is implemented through a modular, time-dependent computational framework developed in Python using scientific numerical libraries (NumPy) with extensive use of array-based operations, allowing quantities to be evaluated over complete time arrays rather than only through step-by-step iteration.
The implementation is structured to keep geometry and visibility modelling, link budget evaluation and data-volume calculation separated, while connecting them through explicit data dependencies.
The overall structure of the framework is illustrated in Figure 29 for RCS1 and Figure 30 for RCS2. The two figures show parallel but distinct computational chains, reflecting the different operational configurations of the two RAMSES CubeSats.

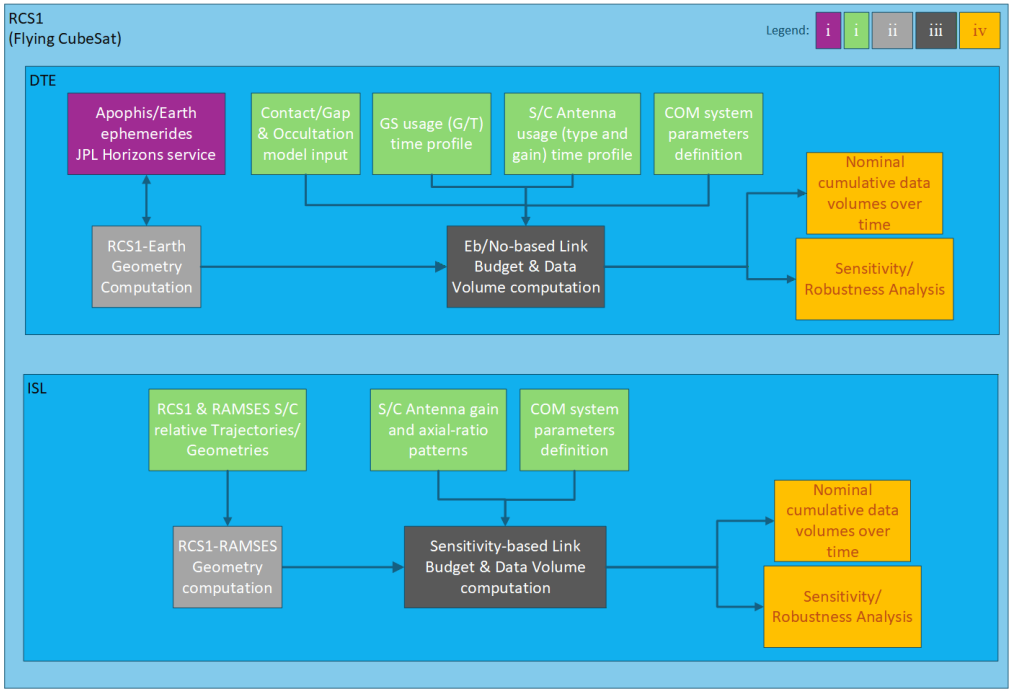


Figure 29: RCS1 Implementation Framework (DTE and ISL branches)

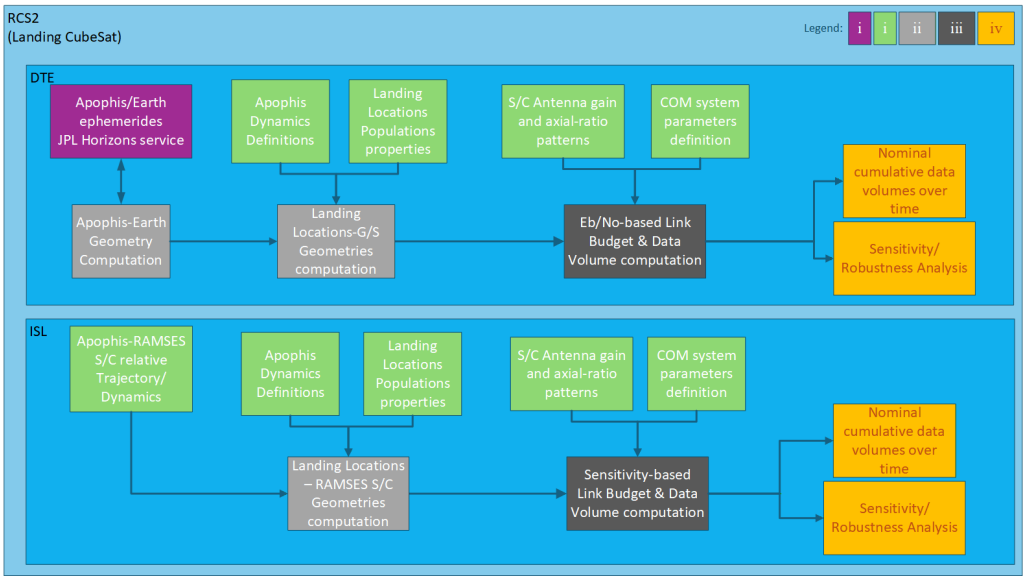


Figure 30: RCS2 Implementation Framework (DTE and ISL branches)

In both Figure 29 and Figure 30, the architecture is organised into four functional stages: (i) external definitions and operational inputs, (ii) geometry computation, and (iii) link feasibility and data-volume integration, followed by performance outputting for post-processing plotting and evaluation (iv).

In Figure 29, the RCS1 DTE branch begins with external inputs that include the Apophis–Earth ephemerides obtained from the online JPL HORIZONS service, the contact/gap abstraction model, the ground station G/T time profile, the spacecraft pointing strategy and

antenna usage schedule, and the communication system configuration parameters (data rates, modulation indexes, Eb/No thresholds, RF power outputs, onboard losses to mention the main ones). These blocks represent time-dependent or configuration-dependent definitions that are not computed internally by the link budget engine.

The ephemerides feed the "RCS1–Earth Geometry Computation" block, which generates time series of range and relative viewing angles. The Contact/Gap & Occultation model, ground station usage profile, and antenna usage schedule are applied at the link evaluation stage, where they modulate the availability and performance of the link.

The geometry outputs, together with antenna gain and axial-ratio characteristics and the communication subsystem parameters, are then passed to the "Eb/No-based Link Budget & Data-volume computation" block. As described in Section 3.4, this block computes the E_b/N_0, evaluates the link margin for each candidate data mode, and selects the highest feasible rate at each time step. The instantaneous rate is integrated over the effective communication time to produce the nominal cumulative data volume quantities. Sensitivity and robustness analyses are obtained by varying selected operational inputs and repeating the computations.

Still in Figure 29, for the RCS1 ISL branch, the external geometry input consists of relative trajectories between RCS1 and the RAMSES spacecraft. These are combined with antenna gain and axial-ratio pattern files and communication system definitions. The geometry computation block computes the range and angular quantities between the two spacecraft as a function of time. The link feasibility is evaluated using a sensitivity-based approach, in which the received power is compared against discrete modem sensitivity thresholds, as detailed in Section 3.4. The selected mode determines the instantaneous data rate, and cumulative data-volume is obtained through temporal integration.

Figure 30 presents the corresponding architecture for RCS2, which differs primarily in the geometry layer. In the DTE branch, the external inputs include Apophis–Earth ephemerides, Apophis rotational dynamics definitions (including the TS–Precession model and the initial phase resolution steps), landing-location population properties, antenna gain and axial-ratio patterns, and communication system parameters.

The "Apophis Dynamics Definitions" block contains the kinematic rotational model described in Section 3.2 and provides the time-dependent orientation of the asteroid. The landing-location population properties define the set of candidate surface sites over which

geometry is evaluated. These inputs feed the “Landing Locations–GS Geometries computation” block, which generates, for each site and time step, range and zenith-angle histories toward the Earth.

The resulting geometry is then passed to the “Eb/No-based Link Budget & Data-volume computation” block. As in the RCS1 DTE case, link feasibility is evaluated through E_b/N_0 margin comparison. The final outputs consist of nominal cases cumulative data-volume curves and distributions. Sensitivity and robustness analyses are obtained by varying selected operational inputs and repeating the computations.

Still on Figure 30, in the RCS2 ISL branch, the geometry is driven by the relative motion between Apophis and the RAMSES spacecraft, combined with the asteroid rotational state and landing-site definitions. The “Landing Locations – RAMSES spacecraft (S/C) Geometries computation” block generates time series of range and angular quantities for each candidate site. This step is done consistently with the DTE case to ensure comparability. Link feasibility is then assessed through a sensitivity-based approach, comparing received power against modem thresholds. As for the DTE case, cumulative data-volumes and statistical envelopes are produced, followed by sensitivity and robustness analysis on selected parameters.

The implementation is therefore modular and hierarchical. External definitions (ephemerides, rotational parameters, antenna patterns, operational assumptions) are treated as inputs to the computational chain. The geometry blocks compute the corresponding positions, ranges, and angles as functions of time. Link budget blocks apply RF modelling consistently with the theoretical formulation presented earlier. Output blocks aggregate the results into performance metrics used in the following assessment.

The framework is deterministic: identical input files and configuration parameters produce identical outputs, ensuring reproducibility of all reported results.

Some statistical visualisations require a representative case in addition to the best and worst cases. In this thesis, the representative median case is selected as the landing location whose cumulative data distribution is closest to the combined distribution across all candidate locations, using the minimum Kolmogorov–Smirnov distance [41]. This method is used only for data visualisation purposes and does not affect the underlying geometry or link budget computations.

3.5.1 Verification of the Implementation

The computational framework was subjected to systematic verification to ensure that the Python implementation correctly reproduces the geometry and link budget formulations introduced in Sections 3.2 to 3.4. This subsection summarises the approach used to confirm that the Python code developed for this thesis performs the intended calculations.
The verification combined independent replication of results, incremental activation of model components and controlled boundary-case analysis.

As a first step, a Direct-To-Earth link analysis was performed for the RAMSES spacecraft case using an implementation of the core framework developed for RCS1 DTE, with a dedicated output format matching an independent industry-standard analysis produced in the frame of the RAMSES project and reported in [42]. The two implementations were compared under identical geometric and RF assumptions. Agreement within numerical precision confirmed the correctness of the core geometry and link budget models implementation.

Second, the model was verified incrementally by activating one dynamic element at a time. Starting from static geometries with fixed range and fixed antenna gain, the framework was tested against analytically predictable outcomes. Subsequently, individual time-varying effects were introduced separately, including:

- Range variation only (with fixed antenna gain and no rotations),
- Antenna gain variation only (with fixed range),
- Single-axis asteroid rotation (all axis),
- Combined rotation and range variation.

At each stage, intermediate key quantities such as free-space loss, received power, C/N_0, and E_b/N_0 were plotted independently to confirm that variations followed the expected analytical results.

Third, for RCS2, controlled geometric edge cases were analysed by positioning landing locations at the extremities of the triaxial ellipsoid to verify correct zenith-angle computation and visibility masking. Simultaneously, the asteroid rotation model was decomposed into single-axis rotations to trace expected markers of specific geometric and dynamic cases, such as complementary behaviour at opposite landing locations, symmetry effects, or singular cases such as landing locations aligned with the rotation axis.

Finally, the time integration of data-volume was verified using simplified constant-rate cases, for which the expected result can be computed directly.

The verification strategy therefore combined analytical cross-checks, incremental dynamics tests, boundary-case evaluation and inspection of intermediate computed quantities. This approach verifies that the framework correctly implements the geometry and link budget models presented earlier, and that the reported performance metrics are produced by the intended numerical logic rather than by implementation artefacts. Validation of the complete mission-specific model against an independent representative case was not possible because no suitable reference was available; the modelling therefore applies established communication theory to the RAMSES mission assumptions defined in this thesis.

3.6Design Space Definition and Baseline Configuration

To apply the modelling framework defined in the previous sections to the RCS1 and RCS2 performance analyses, representative spacecraft and communication-system baselines must be defined.

For the RAMSES CubeSats, the communication design space is constrained by link performance but also by the technology maturity, accommodation, power availability and suitability for deep-space implementation. Many CubeSat communication solutions developed for LEO cannot be assumed to scale directly to the RAMSES case as discussed in Chapter 2. The baseline design space is therefore limited to solutions that remain compatible with CubeSat resources while being suitable for deep space and with the RAMSES mission development schedule.

This section defines the platform resource envelope, antenna architecture and representative subsystem characteristics adopted for RCS1 and RCS2. These definitions provide the design-space boundary used in the performance assessments presented in Chapters 4 and 5.

3.6.1 Platform Resource constraints

The platform resource constraints define the mechanical and power boundaries within which the communication subsystem baselines are selected. The CubeSats' mechanical envelope is assumed to be compatible with the ISISPACE Deep Space Deployer developed for ESA's Hera mission, a qualified deployer designed specifically for low-velocity staged deployment in deep-space proximity operations [43]. The maximum external dimensions and allowable mass constrain the volume and mass available for the communications subsystem, including RF electronics, antenna and RF harnessing. These constraints directly influence the

feasible RF power class, subsystem architecture, and antenna configuration, consistent with standard spacecraft-systems engineering allocation practice [6].

In particular, the limited internal volume and mass allocation preclude the accommodation of large high-power amplification solutions such as travelling-wave tube amplifiers (TWTAs). Redundant transmit chains and complex RF switching matrices are not included in the baseline configuration. The communication architecture is therefore composed of a simplified RF chain, sized to remain compatible with the CubeSat resource constraints while still supporting the DTE and ISL performance comparisons developed in this thesis.

RCS1 Configuration

RCS1 is modelled as a 6U Extended-length CubeSat form factor (XL) platform (approximate dimensions 30 x 20 x 10 cm [43]) with two deployable solar array wings, each of approximate dimensions 30 × 60 cm, mounted on opposite long faces of the bus. The Hera CubeSat Milani, which is also a 6U XL design with 2 deployable solar array wings, is used as a reference and shown in Figure 31.

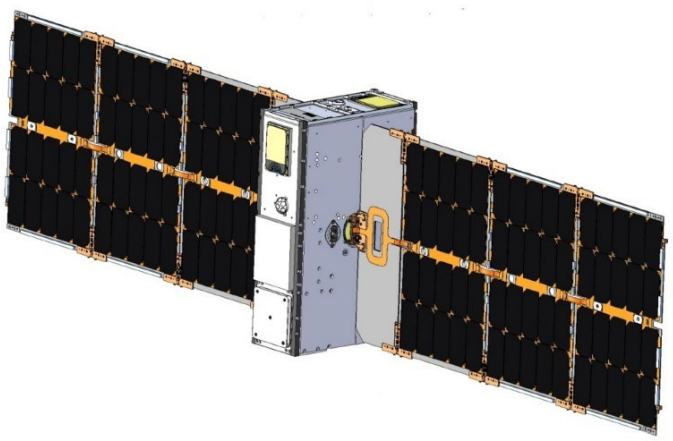


Figure 31: RCS1 reference CubeSat design: ESA Hera Milani. ESA. Source [33]

A simplified CAD model based on this reference is used for antenna radiation pattern simulation, and includes:

- The 6U XL bus,
- Two deployable solar array wings,
- Two opposite low-gain antennas (LGA),
- One body-mounted medium-gain antenna (MGA) with aperture 10x20cm.

The CAD model is simplified for electromagnetic simulation purposes, and does not include fine details, harness routing or internal elements. The model is shown in Figure 32, where also the applied locations of the antenna emission are visible.

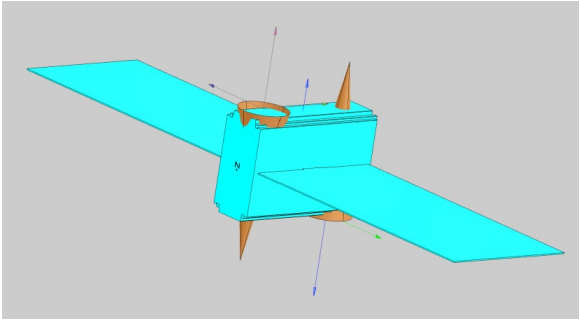


Figure 32: Simplified RCS1 CAD model for antenna simulations

Electrical Power Envelope — RCS1

The available DC power for the RCS1 communication subsystem is estimated from the assumed solar-array configuration and a representative allocation of the generated electrical power to communications. This provides the simplified power-budget reference used to assess compatibility with requirement R2 in Table Ap.3.

The maximum solar power generation capability is estimated from the illuminated cell area.

Assuming constant Sun-pointing of the cells, the peak generated power at 1.1 AU is first estimated using Equation (3.39), following the standard solar-array power relation formulated in [7], Chapter 21.

$$P_{\text{gen,max}} = A_{\text{SA}} \cdot \frac{S_0}{(1.1)^2} \cdot \eta_{\text{cell}} \cdot \eta_{\text{pack}} \cdot \cos\beta \tag{3.39}$$

where:

- A_{SA}= total illuminated cell area (before cells packing factor)
- S_0= solar constant at 1 AU ≈ 1361 W/m²
- $1/(1.1)^2$ accounts for 1.1 AU, the worst-case Sun–RCS1 distance during operations
- η_{cell} and η_{pack} are cell efficiency and cell packing factors
- β= Sun incidence angle (≈ 0 when pointing at the Sun)

Assuming $\eta_{\text{cell}} = 0.30$ and $\eta_{\text{pack}} = 0.85$ as representative values from [44], the solar flux at 1.1 AU is computed from Equation (3.40) as

$$S(1.1) = \frac{S_0}{(1.1)^2} = \frac{1361}{1.21} \approx 1125\ W/\text{m}^2 \tag{3.40}$$

The total geometric solar array wings area is computed from Equation (3.41) as

$$A_{\mathrm{SA}} = 2 \times (0.30 \times 0.60) = 0.36\ \mathrm{m}^2 \tag{3.41}$$

Assuming Sun-pointing ($\beta \approx 0 \Rightarrow \cos\beta \approx 1$), the incident optical power on the wings is computed from Equation (3.42) as

$$P_{\mathrm{in}} = A_{\mathrm{SA}} \cdot S(1.1) \cdot \cos\beta \approx 0.36 \times 1125 \times 1 \approx 405\ \mathrm{W} \tag{3.42}$$

The corresponding peak generated electrical power is obtained from Equation (3.43) as

$$P_{\mathrm{gen,max}} = P_{\mathrm{in}} \cdot \eta_{\mathrm{cell}} \cdot \eta_{\mathrm{pack}} \approx 405 \times 0.30 \times 0.85 \approx 103\ \mathrm{W} \tag{3.43}$$

Only a fraction of the generated electrical power can be allocated to communications after accounting for conversion losses and other platform subsystems, consistent with standard spacecraft power budgeting practice: in preliminary spacecraft power budgets for planetary missions, TT&C subsystems typically represent 15–25% of mission-average electrical power [6]. With the stated assumption of constant Sun-pointing, the peak and average are equal, and therefore the power available is

$$P_{\mathrm{com,avg}} \approx 103\ \mathrm{W} \times 0.15 - 0.25\ \approx\ 15 - 25\ \mathrm{W} \tag{3.44}$$

The value is used as the reference for the available DC power for the communication system in relation to requirement R2 in Table Ap.3.

The communication subsystem design is therefore constrained to RF output power levels compatible with this envelope.

From Table Ap.2, C4, RCS1 is assumed to transmit continuously whenever geometrically feasible during the close-encounter downlink window. The identified 15–25 W range is therefore used as a reference average communication-power allocation rather than as a peak subsystem-power limit. Short-duration demand above this range may require available spacecraft power margin or battery support.

RCS2 Configuration

RCS2 is modelled as a compact 6U XL platform without deployable appendages, with two LGA on the opposite largest panels. The simplified CAD representation used for antenna simulations is shown in Figure 33.

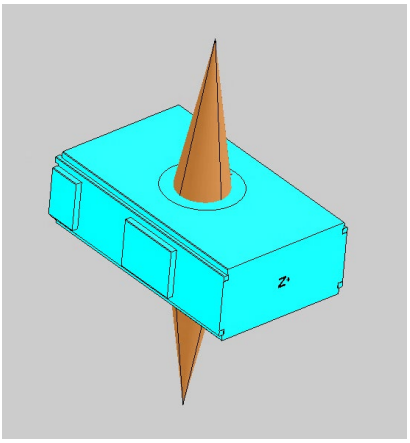


Figure 33: Simplified RCS2 CAD model for antenna simulations

The configuration includes:

- The 6U bus,
- Two opposite body-mounted low-profile antennas located at the centre of the nominal zenith and nadir faces.

No deployable panels or other directive antenna is assumed for RCS2.

Electrical Power Envelope — RCS2

RCS2 is assumed to operate on a primary battery during the surface phase, potentially supported by limited body-mounted solar panels. The surface mission duration (tens of hours) implies that both instantaneous power and total available energy are extremely constrained. For the communication-energy assessment, the communication subsystem is conservatively assumed to remain active throughout the considered RCS2 operational lifetime, rather than being switched according to predicted communication opportunities. Unlike RCS1, the generic spacecraft power-allocation rule is not applied to RCS2, as its short-duration, primarily battery-powered operation is not considered representative of the spacecraft-level assumptions underlying that approach.

The corresponding communication-energy demand is estimated as $E_{\mathrm{com}} = P_{\mathrm{com}} \cdot T$, where P_{com} is the communication-subsystem DC power and T is the assumed operating duration. Based on the reference-equipment peak DC powers adopted in Section 3.6.2, 12 W for DTE and 5.5 W for ISL, the corresponding communication-energy demands over 20 h, 40 h and 60 h are 240 Wh, 480 Wh and 720 Wh for DTE, and 110 Wh, 220 Wh and 330 Wh for ISL, respectively. The ISL communication-energy demand for an equal operating duration is approximately 46% of that of DTE, corresponding to a reduction of approximately 54%. The absolute demand scales linearly with the assumed RCS2 operational lifetime. A numerical

spacecraft-level battery-energy closure is not reported because the detailed RCS2 battery capacity and spacecraft energy allocation are not available within the public mission information used as the basis of this thesis, and a representative battery-energy capacity for the small lander cannot be established from the available reference information. The assessment is therefore limited to the communication-system energy demand. Since the two architectures are evaluated over the same operating duration, this limitation does not affect their relative energy-demand comparison for the assumed reference equipment.

Due to the combined mass, volume and energy constraints, high-power amplification solutions are not considered feasible in the baseline design, and the lowest available RF output-power option for the reference CubeSat-class deep-space transponder is therefore adopted.

3.6.2 Communication Subsystem Sizing and Antenna Architecture

This section defines the representative onboard communication configurations adopted for RCS1 and RCS2, based on the communication-system baseline and resource constraints introduced in Section 2.5, and reported in Table Ap.2 and Table Ap.3, and on the operational roles described in Chapter 2. The section first introduces the reference DTE and ISL equipment classes, then defines the RCS1 and RCS2 communication architectures, and finally describes the antenna radiation-pattern assumed for this thesis, simulated to account for the spacecraft accommodation effects.

Communication baseline

Consistent with the communication architectures and reference implementations introduced in Sections 2.3 and 2.4, the communication baseline combines two representative subsystem options: a CubeSat-compatible X-band deep-space TT&C transponder for DTE, and a compact S-band ISL terminal for relay communication. The DTE baseline uses the telemetry-rate set, coding assumptions, ground-segment support and E_b/N_0-based feasibility criterion defined in Section 2.5 and modelled in Section 3.4. The ISL baseline uses the Hera/GAMALINK-type CubeSat relay terminal introduced in Section 2.4, with the data modes, dynamic bandwidth configuration, receiver-sensitivity thresholds and dual CubeSat support defined in Section 2.5 and modelled in Section 3.4.

The same representative communication equipment options are used for both CubeSats, but with different antenna configurations. Both CubeSats include an LGA-based configuration using two opposite low-profile circularly polarised patch antennas, providing wide angular

coverage for reduced-pointing or geometry-driven cases. This LGA configuration forms the antenna baseline for both ISL architectures and for the RCS2 DTE architecture.

For RCS1 DTE, the nominal far-range configuration uses a body-mounted medium gain antenna (MGA) to support higher-gain communication when Earth-pointing capability is available. This is consistent with the RCS1 DTE antenna baseline in Table Ap.2, C8, and the operational profile described in Section 2.2 and shown in Figure 5. During close-encounter conditions or reduced-pointing, where Earth pointing is not assumed, the LGA configuration is used instead as the fallback communication case.

RCS1 communication configuration

The RCS1 communication configuration is organised around two architecture-dependent RF chains: the X-band DTE chain and the S-band ISL chain.

For DTE, the baseline includes the transponder, RF high-power amplification stage, MGA and LGA antennas, and distribution network providing the required RF paths. The X-band transponder is represented using the European C-DST reference discussed in Section 2.4, with an assumed RF output power of 5.8 dBW, corresponding to 3.8 W, at the amplifier output [30]. The associated 31 W DC power reported in Table 3.1 represents the peak RX + TX subsystem demand and is not assumed as a continuous or mission-average power consumption.

The MGA is body-mounted on the RCS1 outer structure, meaning that antenna pointing requires reorientation of the full spacecraft. The small external panel area limits the antenna aperture, making a medium-gain solution more compatible with the CubeSat accommodation envelope than a larger, more complex deployable high gain antenna (HGA). A heritage-inspired topology consistent with the LICIACube X-band communication system is used as the representative reference [28]. Based on that reference, an antenna envelope of approximately 20 x 10 cm is assumed in this study, allowing a representative 4 x 12 radiating-element arrangement and a representative boresight gain of 20 dBi with circular polarisation. This gain figure already includes pointing losses assuming a pointing error of 0.5°. RCS1 is assumed to provide pointing capability compatible with this requirement, which is well within the 10 arcsec capability reported for LICIACube [28].

The resulting CubeSat-compatible X-band deep-space TT&C system architecture is shown in Figure 34.

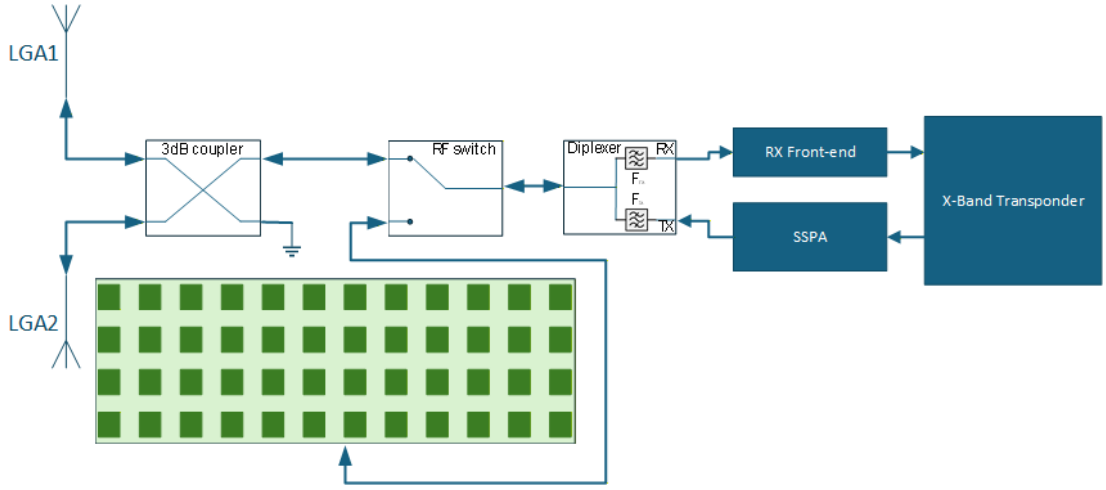


Figure 34: RCS1 DTE communication architecture

Table 3.1 summarises the properties of the resulting RCS1 DTE communication system and its main elements.

Table 3.1 - RCS1 DTE communication subsystem baseline

Parameter	Baseline definition
RF subsystem	
Transponder class	CubeSat-compatible X-band deep-space TT&C transponder
Downlink frequency	8436 MHz, X-band
RF output power	5.8 dBW (3.8 W) at TX RF port
DC power (peak)	31 W, RX + TX
Accommodation	
Mass	< 3.0 kg, complete DTE subsystem
Volume	1.5U, transponder only; antenna and harness excluded
RF losses	
Onboard losses	2 dB
VSWR	1.15
Modulation	
TM modulation index	1.2 rad
RNG modulation index	0.2 rad
Antenna configuration	

Antennas	Combined two-LGA path and body-mounted MGA
Coding and telemetry rates	
FEC options	Turbo 1/2, (Eb/N0) req. = 1.35 dB; Turbo 1/4, (Eb/N0) req. = 0.25 dB
Rate set	25 TM modes, TM01–TM25; see Appendix IV, Table Ap.5

For the ISL architecture, RCS1 uses a compact S-band ISL terminal together with the LGA-based antenna configuration. The terminal is represented using the Hera/GAMALINK-type reference introduced in Section 2.4, with the main modem characteristics taken from the supplier documentation [34], [35], [45]. Although the terminal supports dynamic bandwidth allocation among concurrent links, this capability is not exploited in the baseline configuration. A fixed symmetric allocation is instead adopted for RCS1, with the corresponding data rates defined in Appendix IV. This is a conservative choice, as the higher-priority allocation could be used for most of the RCS1 mission with the exception of the short period of RCS2 operations.

To connect the two LGAs, a passive splitter/combiner is assumed. This increases RF insertion loss compared with a three-port RF-switch design but provides a simple and robust configuration not requiring active selection between antennas. The resulting ISL architecture is shown in Figure 35.

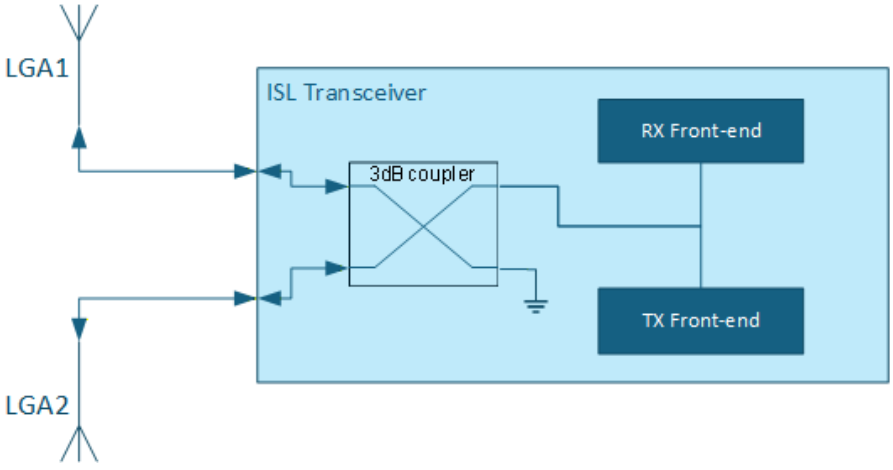


Figure 35: RCS1 ISL communication architecture

The corresponding subsystem parameters and modem operating modes adopted for the RCS1 ISL analysis are summarised in Table 3.2.

Table 3.2 - RCS1 ISL communication subsystem baseline

Parameter	Baseline definition
RF subsystem	
Terminal class	Compact S-band inter-satellite-link terminal
Carrier frequency	2430 MHz, S-band
RF output power	27 dBm (0.5 W), at ISL terminal RF port
DC power (peak)	5.5 W maximum
Accommodation	
Mass	approximately 0.8 kg, complete ISL subsystem
Volume	0.5U, antenna excluded
Antenna configuration	
Antennas	Combined two-LGA path
RF distribution	Passive splitter/combiner architecture
RF losses	
Onboard losses	7.5 dB total: 1.5 dB cables/connectors + 3 dB TX splitter + 3 dB RX splitter
Mode-selection basis	
Modem representation	Adaptive data modes with dynamic bandwidth allocation capability
Dynamic bandwidth configuration	Fixed symmetric allocation for RCS1 baseline
Mode selection	Adaptive: highest feasible mode satisfying the receiver-sensitivity threshold
Data modes	Six data modes defined by data rate, bandwidth-allocation configuration and receiver-sensitivity threshold; see Appendix IV, Table Ap.6

For the ISL mode-selection approach defined in Table 3.2, the highest feasible mode is selected at each simulation step based directly on the corresponding receiver-sensitivity threshold. This provides a simplified representation of achievable link performance based on the deterministic link conditions calculated from the defined geometry and link parameters. The model therefore does not represent the short-term link-quality fluctuations that could

cause repeated transitions between adjacent modes around a threshold in the real world. In the actual ISL reference terminal, data-mode transitions are managed by link-state mechanisms that monitor link quality and adapt between higher-throughput and more robust modes as conditions vary [35].

RCS2 communication configuration

For RCS2, the same two communication options are considered as for RCS1: the X-band DTE chain and the S-band ISL chain. The landed role of RCS2 leads to a simpler antenna configuration, based on the two-LGA architecture for both DTE and ISL, without an additional MGA for the DTE. This reflects the fixed landed attitude after touchdown, which does not allow a directive antenna to be pointed towards Earth.

For DTE, the architecture is based on the same CubeSat-compatible X-band deep-space transponder reference adopted for RCS1. However, the lowest available RF output power option is selected to reduce DC power demand and limit the RCS2 communication-energy requirement discussed in Section 3.6.1. The RF chain therefore includes the transponder, high-power amplification, RF distribution network and combined LGA antenna path. Figure 36 shows the resulting RCS2 DTE communication architecture.

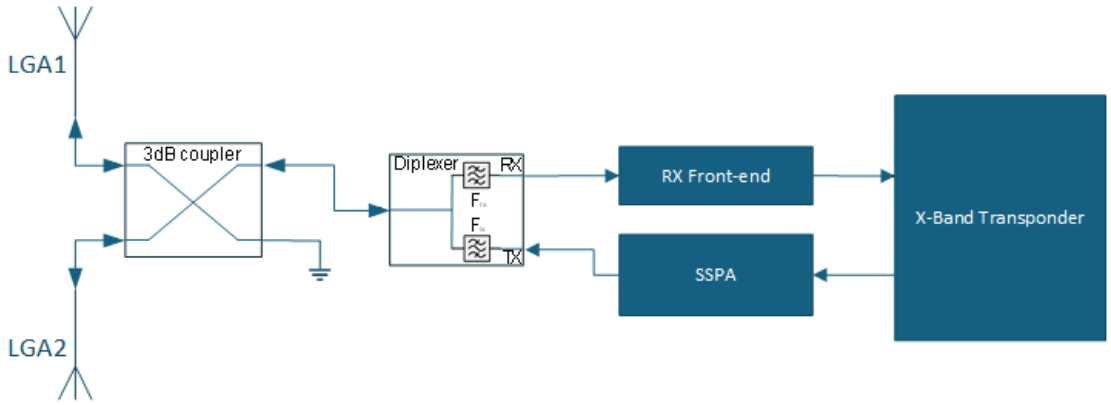


Figure 36: RCS2 DTE communication architecture

The subsystem baseline adopted for the RCS2 DTE analysis is summarised in Table 3.3.

Table 3.3 - RCS2 DTE communication subsystem baseline

Parameter	Baseline definition
RF subsystem	
Transponder class	CubeSat-compatible X-band deep-space TT&C transponder
Downlink frequency	8436 MHz, X-band

RF output power	0 dBW (1 W) at TX RF port
DC power (peak)	12 W, RX + TX
Accommodation	
Mass	< 3.0 kg, complete DTE subsystem
Volume	1.5U, transponder only; antenna and harness excluded
RF losses	
Onboard losses	2 dB
VSWR	1.15
Modulation	
TM modulation index	1.2 rad
RNG modulation index	0.2 rad
Antenna configuration	
Antennas	Combined two-LGA path
Coding and telemetry rates	
FEC options	Turbo 1/2, (Eb/N0) req. = 1.35 dB; Turbo 1/4, (Eb/N0) req. = 0.25 dB
Rate set	25 TM modes, TM01–TM25; see Appendix IV, Table Ap.5

For the ISL architecture, RCS2 uses the same compact S-band terminal and two-LGA antenna configuration adopted for RCS1 ISL. Unlike the RCS1 baseline, which uses a symmetric bandwidth allocation, the RCS2 baseline follows the high-priority allocation defined in Appendix IV, reflecting the short surface-operation window and the time-critical nature of the landed data return.

The RCS2 ISL subsystem baseline is summarised in Table 3.4.

Table 3.4 - RCS2 ISL communication subsystem baseline

Parameter	Baseline definition
RF subsystem	
Terminal class	Compact S-band inter-satellite-link terminal
Carrier frequency	2430 MHz, S-band

RF output power	27 dBm (0.5 W), at ISL terminal RF port
DC power (peak)	5.5 W maximum
Accommodation	
Mass	approximately 0.8 kg, complete ISL subsystem
Volume	0.5U, antenna excluded
Antenna configuration	
Antennas	Combined two-LGA path
RF distribution	Passive splitter/combiner architecture
RF losses	
Onboard losses	7.5 dB total: 1.5 dB cables/connectors + 3 dB TX splitter + 3 dB RX splitter
Mode-selection basis	
Modem representation	Adaptive data modes with dynamic bandwidth allocation capability
Dynamic bandwidth configuration	High-priority allocation for the RCS2 baseline
Mode selection	Adaptive: highest feasible mode satisfying the receiver-sensitivity threshold
Data modes	Six data modes, defined by: data rate, dynamic bandwidth configuration and receiver-sensitivity threshold; see Appendix IV, Table Ap.6

The mode-selection approach follows the same simplified representation defined for RCS1 ISL in Table 3.2 and discussed above.

Antenna radiation-pattern modelling

The antenna radiation patterns used in this thesis are derived from simulations performed using TICRA ESTEAM, a software tool used to model the effects of the spacecraft structure on antenna gain and axial ratio. This type of approach has been applied in comparable CubeSat platform-scattering and structure-antenna coupling performance studies [46], [47].

The simplified RCS1 and RCS2 spacecraft CAD models defined earlier in this section are imported into the simulation environment. The antenna elements are represented as radiation-pattern, consistent with the ESA Hera mission ISL antenna [48], avoiding the need to reconstruct their detailed electrical design.

Metallic spacecraft surfaces are modelled as perfect electric conductors. This is consistent with the aluminium spacecraft structures and microwave frequencies of ISL and DTE, for which conductive surfaces are expected to dominate the antenna–structure interaction. The approximation simplifies the electromagnetic model while preserving the main effects relevant to this level of communication-performance analysis [49].

The simulations account for antenna–spacecraft coupling effects and structural shadowing. The resulting far-field gain and axial-ratio maps are exported in spherical coordinates (θ, ϕ) and used as inputs to the link budget model described in Section 3.4.

Four CubeSat simulation cases are defined: ISL and DTE low-gain-antenna configurations for both RCS1 and RCS2. Figure 37 to Figure 40 show the simulation results as polar cuts.

Antenna angular coverage is reported, defined as the fraction of the full sphere where the gain exceeds selected thresholds, with sin θ weighting applied.

Figure 37 shows the simulated RCS1 ISL antenna gain pattern used for the ISL link budget analysis.

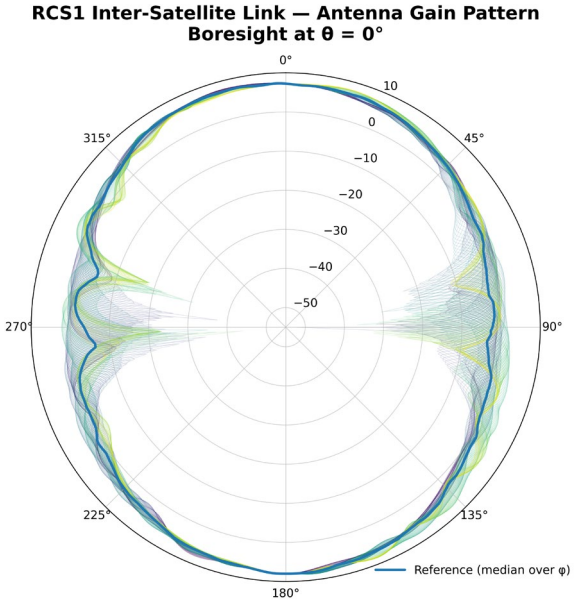


Figure 37: RCS1 ISL antenna gain pattern

The pattern remains broadly quasi-omnidirectional but includes structure-induced notches and variations between azimuth cuts. Around the $\theta = 90°$ plane, the pattern also shows constructive and destructive coupling between the two LGAs. RCS1 ISL provides gain above −10 dBi over approximately 96.6% of the spherical solid angle and above −12 dBi over approximately 97.9%. These variations motivate the statistical antenna-cut treatment adopted in the RCS1 ISL analysis.

Figure 38 shows the RCS2 ISL antenna gain pattern for the compact landed-CubeSat configuration.

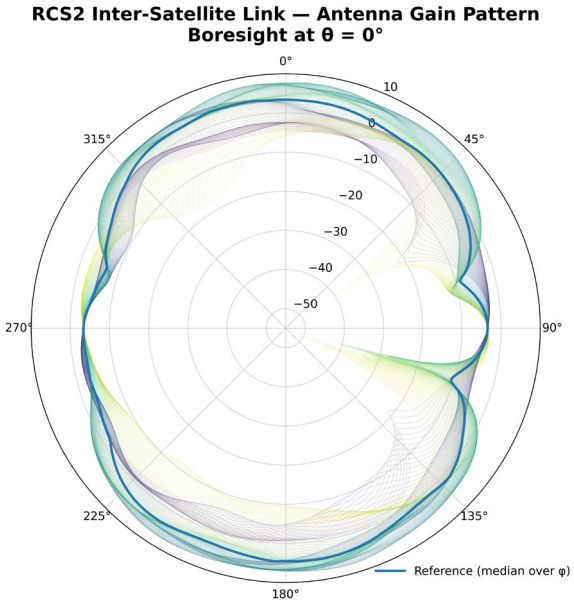


Figure 38: RCS2 ISL antenna gain pattern

Compared with RCS1, the compact 6U configuration without deployable solar-array wings produces a different response to the spacecraft structure, with more pronounced asymmetry and local shadowing. RCS2 ISL provides gain above −10 dBi over approximately 95.0% of the spherical solid angle and above −12 dBi over approximately 97.0%. This pattern is used as the antenna input for the landed ISL analysis.

Figure 39 shows the RCS2 DTE combined-LGA antenna gain pattern used for the landed DTE analysis.

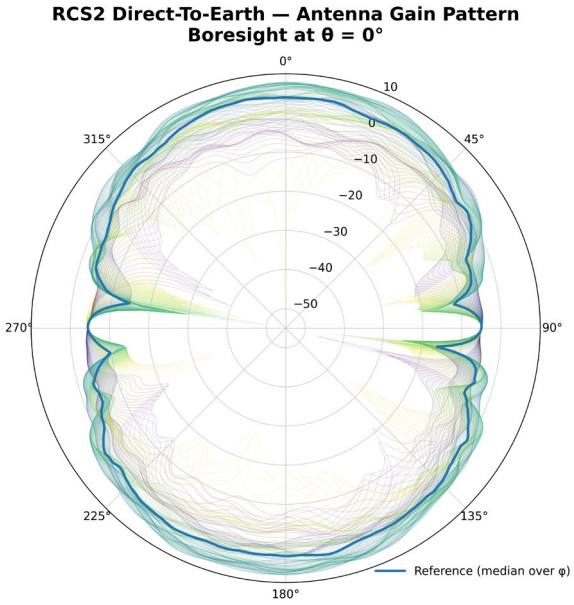


Figure 39: RCS2 DTE combined-LGA gain pattern

Also in the case, the pattern reflects the combined response of the two LGA elements, including spacecraft shadowing and coupling effects. Compared with the S-band ISL patterns, the X-band DTE pattern shows a more structured angular response because the two-antenna combination produces more rapid phase variation with direction at the higher frequency. RCS2 DTE provides gain above −10 dBi over approximately 90.0% of the spherical solid angle and above −12 dBi over approximately 93.5%. This pattern is used as the antenna input for the landed DTE analysis.

Figure 40 shows the simulated RCS1 DTE combined-LGA antenna gain pattern.

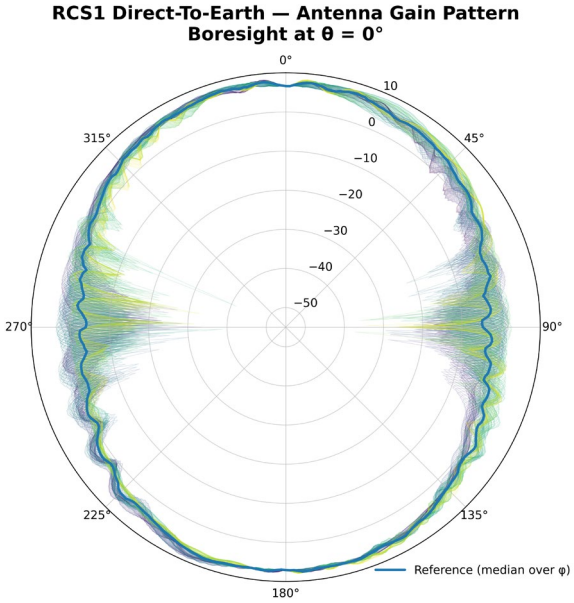


Figure 40: RCS1 DTE combined-LGA gain pattern

Compared with the RCS2 configuration, the deployed solar-array wings of RCS1 provide a better isolation among the LGAs resulting in less pronounced and different distribution of local notches across the angular domain. RCS1 DTE provides gain above −10 dBi over approximately 97.0% of the spherical solid angle and above −12 dBi over approximately 98.0%.

For the RCS1 MGA, the DTE option uses the representative boresight gain defined in the RCS1 DTE baseline in Table 3.1, consistent with the antenna-use profile introduced in Table Ap.2, C8. Since the MGA is used only in the nominal Earth-pointed configuration, the boresight gain value is assumed applicable. A separate spacecraft-antenna pattern simulation is therefore not introduced for the MGA, as structure effects are not expected to be present in the boresight region of this aperture-type antenna.

For the ISL architecture, the RAMSES spacecraft antenna also contributes to the local link budget. Figure 41 shows the isolated antenna pattern used to represent the RAMSES-side ISL antenna.

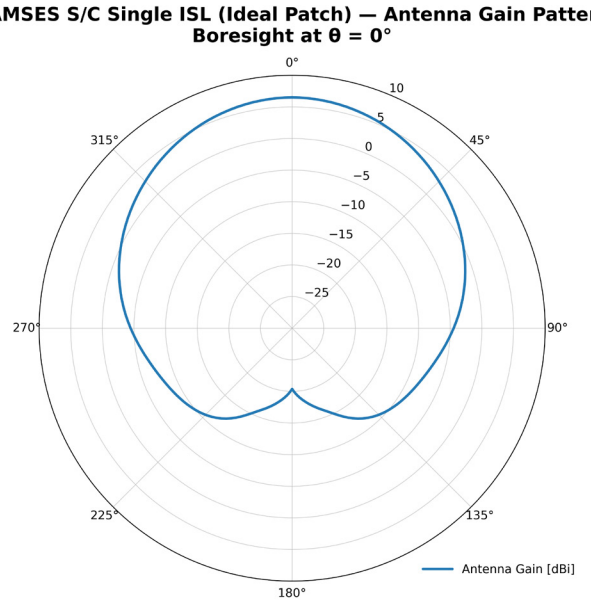


Figure 41: RAMSES-side ISL isolated antenna gain pattern

The RAMSES spacecraft ISL antenna is treated in the same simplified way as the RCS1 MGA. A CAD model of the RAMSES spacecraft was not available for this research work. The antenna pattern was therefore represented without spacecraft-mounted effects. RAMSES is assumed to maintain an asteroid-pointing attitude during proximity operations, with the ISL antenna located on the asteroid-facing side of the spacecraft. Under this condition, the RCS1 and RCS2 link directions are assumed to remain close to the RAMSES ISL antenna boresight over the analysed geometries, therefore the isolated boresight gain is applied. The validity of this assumption is checked for the RCS1 Section 4.2.2 covering also RCS2.

3.7Model Applicability and Limitations

This section consolidates the main applicability limits of the modelling framework defined in the preceding sections. It summarises how the adopted geometric, operational and communication-system assumptions affect the interpretation of the result.

The propagation model uses the standard free-space path-loss formulation, with the analysed links verified to operate in the antenna far-field region as discussed in Section 3.3. This is appropriate for the first-order link budget analysis, but does not include local propagation effects such as multipath reflections, terrain-edge diffraction, small-scale fading or detailed electromagnetic interaction with the asteroid surface. Surface-related properties and uncertainty are instead represented through simplified geometric sensitivity cases, including minimum-elevation masking and local tilt/slope perturbations.

Apophis rotation is represented through the parametric kinematic Tilt/Spin–Precession model introduced earlier in this chapter. This captures the time-varying surface-orientation effects relevant to link geometry but does not solve the full rigid-body dynamics or propagate possible rotational-state changes caused by the 2029 close encounter.

Operational behaviour is represented at architecture-assessment level. Contact availability, visibility gaps and telemetry-rate changes are modelled through contact duty cycles, visibility masks and time-dependent rate profiles rather than through a detailed operations schedule. This keeps the analysis focused on the main communication-architecture drivers without introducing dependence on a specific operational timeline.

The antenna modelling uses simplified spacecraft geometries for the two RAMSES CubeSats. Fine details, tolerances, thermal deformation and statistical antenna-performance variations are not included. The RAMSES spacecraft ISL antenna is also represented through its assumed boresight behaviour rather than through a spacecraft-mounted antenna-pattern simulation.

For RCS2 surface operations, the results depend on the adopted sample of landing locations and initial rotational phases. They should therefore be interpreted as representative of the defined landed-case population, rather than as exhaustive coverage of all possible touchdown conditions.

With these limitations, the framework provides a consistent basis for comparing the DTE and ISL options and identifying the main robustness drivers for each CubeSat case. The results should be interpreted as architecture-assessment outputs under the stated assumptions, rather than as final mission-operations predictions.

3.8Chapter summary

A common geometric basis was established using JPL HORIZONS ephemerides in an Apophis-centred inertial frame. The time-varying orientation of Apophis was represented using a triaxial ellipsoid and a parametric Tilt/Spin–Precession kinematic model, allowing surface visibility and antenna orientation to be evaluated without introducing a full rigid-body rotational model.

Propagation and visibility were treated through the verified far-field free-space link regime, line-of-sight conditions and minimum-elevation masks ($\varepsilon_{\min} = 10°$). Additional geometric effects were represented through simplified but explicit modelling assumptions, including

tilt/slope sensitivity cases for RCS2 and a 3.5% RCS1 occultation penalty due to blockage by Apophis.

The link budget formulation was defined separately for the two communication regimes. DTE feasibility was evaluated using E_b/N_0- based telemetry criteria, while ISL feasibility was assessed using receiver-sensitivity thresholds for discrete data modes. In both cases, the highest feasible rate or mode was selected over time and integrated to obtain cumulative data-volume.

The framework was implemented in Python as a modular time-dependent analysis chain, with separate geometry, visibility, link budget and data-volume calculations. Verification checks confirmed consistency between the implementation, analytical formulations and independent reference cases.

The communication design space was then defined for RCS1 and RCS2, including platform resource constraints, representative DTE and ISL subsystem baselines, and spacecraft-mounted antenna-pattern inputs.

This chapter therefore establishes the quantitative basis for the RCS1 and RCS2 performance analyses in Chapters 4 and 5. In relation to the research objectives defined in Chapter 1, it addresses Objective 1 by developing the communication-architecture modelling framework, and provides the modelling basis required for the nominal-performance, statistical-characterisation and sensitivity analyses that follow.

Chapter 4 Orbiting CubeSat (RCS1) Communication analysis

4.1Introduction

This chapter presents the communication analysis results for the orbiting RAMSES CubeSat, RCS1, using the modelling framework developed in Chapter 3. The analysis covers the two candidate architectures defined earlier in the thesis: Direct-To-Earth (DTE) communication from RCS1 to Earth ground stations, and Inter-Satellite-Link (ISL) communication via the RAMSES spacecraft.

Both architectures are first evaluated under nominal mission conditions, followed by sensitivity cases representing operational uncertainties or degraded link conditions. The analysis focuses on communication feasibility, data-rate performance and cumulative data-volume over the RCS1 mission timeline. For DTE, data-rate performance is expressed through the achievable telemetry rate, with residual-carrier tracking checked separately. For ISL, it is expressed through the achievable data mode.

Section 4.2 presents the nominal DTE and ISL performance, while Section 4.3 investigates sensitivity to degraded link conditions, operational constraints and uncertainties. The key observations are summarised in Section 4.4 and used in Chapter 6 for the architecture-level evaluation.

4.2RCS1 Nominal Performance over the proximity-operations

4.2.1 Direct-To-Earth Communication Performance

The nominal DTE case uses the RCS1 communication configuration defined in Section 3.6.2 together with the operational-profile assumptions introduced in Section 2.2.1.

Figure 42 shows the time-dependent input profiles used for the DTE link budget calculation: spacecraft transmit-antenna gain, ground station receive performance and contact duty cycle over the RCS1 mission timeline.

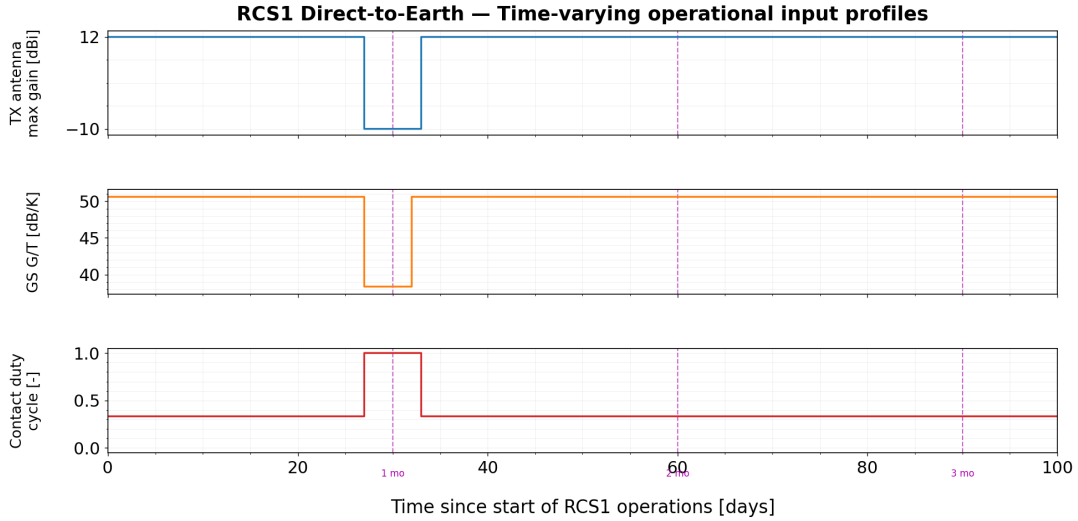


Figure 42: RCS1 DTE time-varying operational input profiles

Figure 42 shows the time-dependent DTE inputs used in the nominal RCS1 calculation. The transmit-antenna gain is generally set to the MGA value, with a short interval around closest approach based on a conservative LGA gain value. The same interval also uses a reduced ground station receive-performance G/T, while the equivalent contact duty cycle increases to 100% to represent enhanced ground station coverage during the close-approach period. These profiles define the operational inputs used for the nominal DTE results that follow.

Figure 43 shows the Earth–Apophis range evolution and the corresponding free-space path loss (FSPL) over the entire RCS1 operational timeline. As defined in Chapter 3, the Earth–Apophis range is computed from Equation (3.2), while the DTE free-space path loss is computed from Equation (3.20).

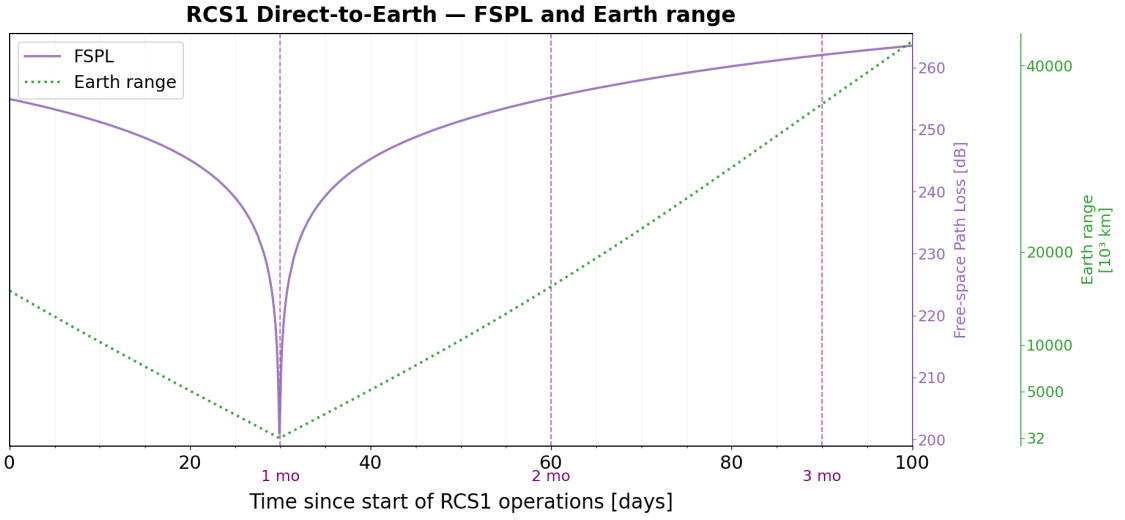


Figure 43: RCS1 DTE Free Space Path Loss evolution

Figure 43 shows the rapid reduction in Earth–Apophis range during the first month of RCS1 operations, from approximately 15 million km to approximately 32,000 km at closest approach. This produces a reduction of approximately 55 dB in free-space path loss from the initial 255 dB. After closest approach, the range increases again, reaching more than 40 million km by the end of the analysed period, and the FSPL rises correspondingly by approximately 60 dB to values close to 263 dB.

This large FSPL variation sets the main trend in the available DTE link budget, affecting the resulting telemetry feasibility and cumulative data return.

Figure 44 presents the residual carrier C/N_0 over the RCS1 timeline, together with the 17 dBHz carrier-tracking thresholds defined in Section 3.4 and the corresponding maximum PLL tracking bandwidth. The residual-carrier C/N_0 is computed from the modulation-loss and carrier-only link budget formulations of Equations (3.29) and (3.30), while the PLL bandwidth is estimated using Equation (3.31).

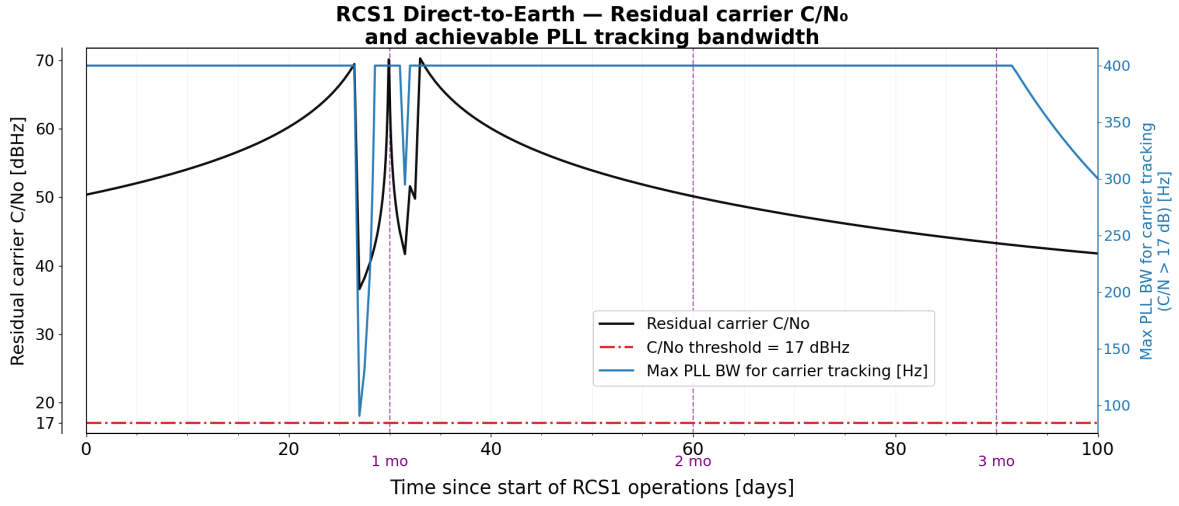


Figure 44: RCS1 DTE Residual carrier C/No and PLL tracking bandwidth

Figure 44 shows that the residual carrier C/N_0 remains above the required tracking threshold throughout the analysed timeline. Carrier acquisition and tracking are therefore not a limiting condition for the nominal RCS1 DTE case, allowing the subsequent assessment to focus on telemetry feasibility and data return.

The residual-carrier trend follows the FSPL evolution shown in Figure 43. As the Earth–Apophis range decreases before closest approach, the reduction in propagation loss increases the received carrier power and therefore the residual carrier C/N_0. After closest approach, the

increasing range produces the opposite effect, with a gradual reduction in residual carrier C/N_0.

The step changes visible in Figure 44 around the first month of RCS1 operation, associated approximately to the Apophis closest approach, follow from the time-dependent antenna-gain and ground station G/T profiles shown in Figure 42. During this interval, the conservative LGA gain value and reduced ground station receive performance lower the residual-carrier margin, even though the Earth–Apophis range is at its most favourable. Carrier tracking remains feasible, but the maximum PLL tracking bandwidth is reduced during this constrained communication interval.

For most of the RCS1 mission timeline, the maximum PLL bandwidth reaches the 400 Hz upper limit considered in the analysis, indicating that carrier tracking is not a limiting condition. During the short degraded intervals around closest approach, the bandwidth decreases to values of the order of tens of hertz. This remains sufficient for carrier tracking in the nominal case, but shows that Doppler-rate variations and other frequency-instability effects would become more relevant during these constrained periods.

Carrier recovery is therefore not expected to limit nominal RCS1 DTE performance. The following figures examine the corresponding telemetry-link margin, achievable data rate and cumulative data return.

Figure 45 presents the telemetry-link feasibility margins for the RCS1 DTE case over the operational timeline. The margins are computed using the DTE feasibility formulation of Equations (3.23)–(3.27), while the selected telemetry mode is determined using Equation (3.28). Figure 46 provides a zoomed view around closest approach, where the fastest variations in link performance occur.

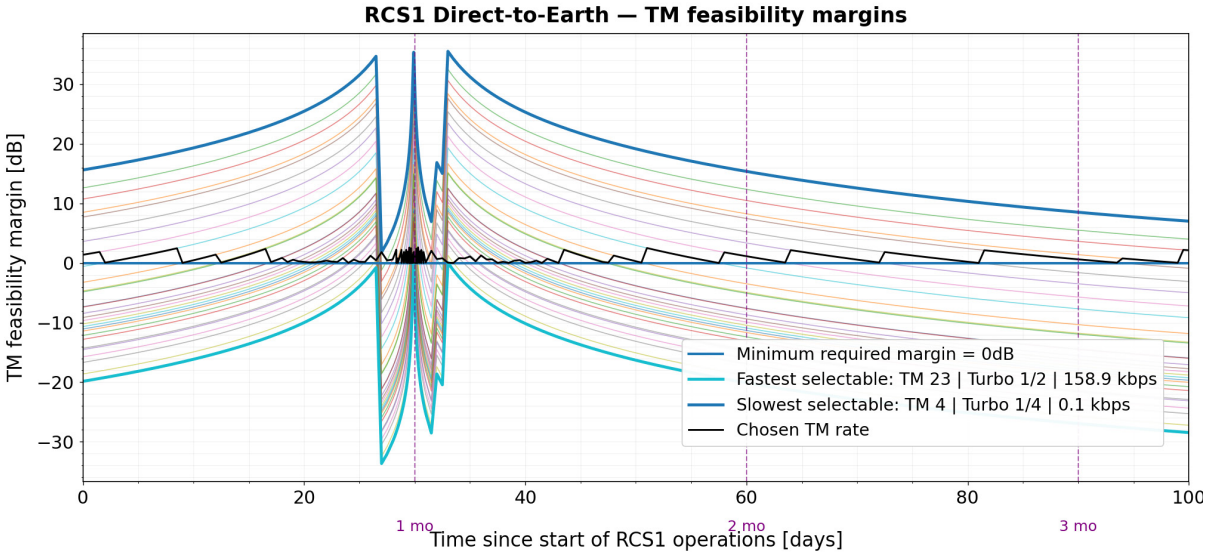


Figure 45: RCS1 DTE TM feasibility margins over mission time

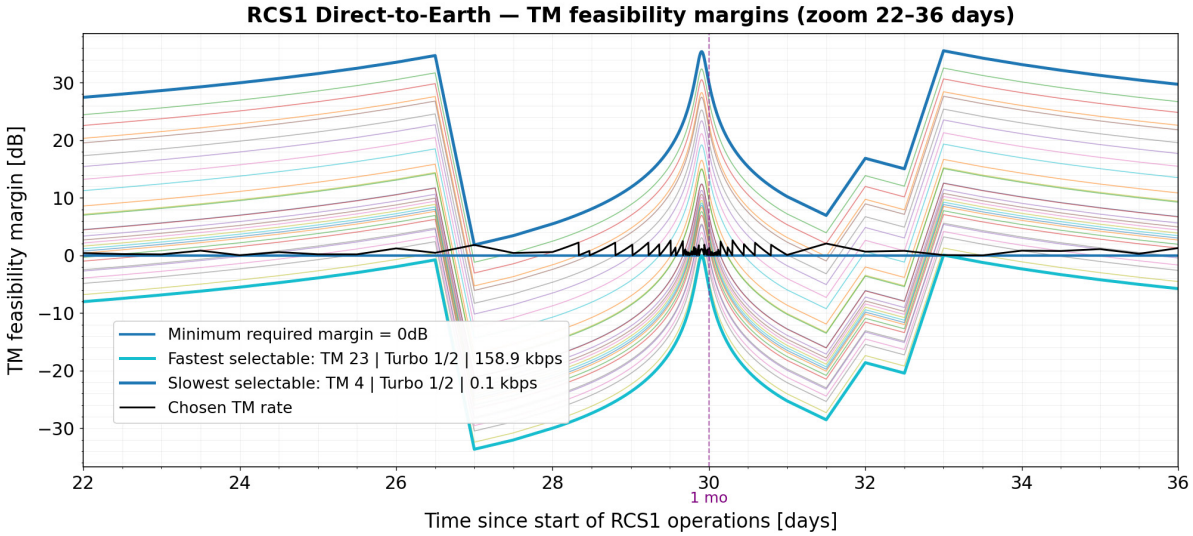


Figure 46: RCS1 DTE TM feasibility margins around closest approach

Figure 45 and Figure 46 show that, similarly to what was observed for the residual carrier C/N_0, the feasibility margins also exhibit a global evolution that closely follows the variation of FSPL together with the changes in onboard antenna gain and ground station G/T. As the RCS1–Earth distance decreases, the reduction in FSPL increases the available link margin, allowing higher telemetry rates to become feasible.

Around the closest approach the rapid reduction in propagation loss produces a sharp increase in the achievable link margin. During this short interval of a few days, the fastest telemetry configuration available becomes feasible, reaching positive margins above the minimum required level for all the available rates, without the possibility to further increase the

communication speed. However, such favourable link conditions are concentrated within a relatively limited time window around closest approach.

A sharp reduction in achievable telemetry rate also occurs at the edges of the close-encounter interval, both before and after minimum distance. During these periods, the reduction of onboard antenna gain and ground station receive performance coincide with still relatively high FSPL. The resulting link budget reduction temporarily limits the DTE link to very low telemetry rates.

After the close-encounter phase, the recovery of higher onboard antenna gain and improved ground station performance compensates for the increasing propagation loss. However, the continuing growth in Earth–spacecraft distance reduces the available link margin and gradually lowers the selectable telemetry rate.

Figure 45 and Figure 46 also show frequent changes in the highest feasible telemetry configuration over time. When RCS1 is far from Earth, the achievable telemetry rate remains stable over several days because the link budget varies relatively slowly. Near closest approach, the rapid change in propagation geometry, combined with the time-dependent onboard antenna gain and ground station G/T, causes the fastest selectable telemetry rate to change over much shorter intervals. This indicates that telemetry-rate planning becomes more time-critical during the close-approach phase.

Overall, the feasibility margins remain positive throughout the RCS1 mission timeline. As the link budget evolves, the selected telemetry configuration follows the highest rate compatible with the required margin, with the fastest available mode becoming feasible during the short period of favourable geometry around closest approach.

Detailed DTE telemetry link budget plots for the representative RCS1 cases are provided in Appendix V.

The computed feasibility margins are used to select the instantaneous maximum feasible telemetry rate, which is then integrated using Equation (3.32) and Equation (3.33) to obtain cumulative data return. Figure 47 shows the cumulative data-volume obtained from the selected RCS1 DTE telemetry rates over the operational timeline. Figure 48 provides a zoomed view around the close-encounter phase, where the largest short-term changes in data accumulation occur.

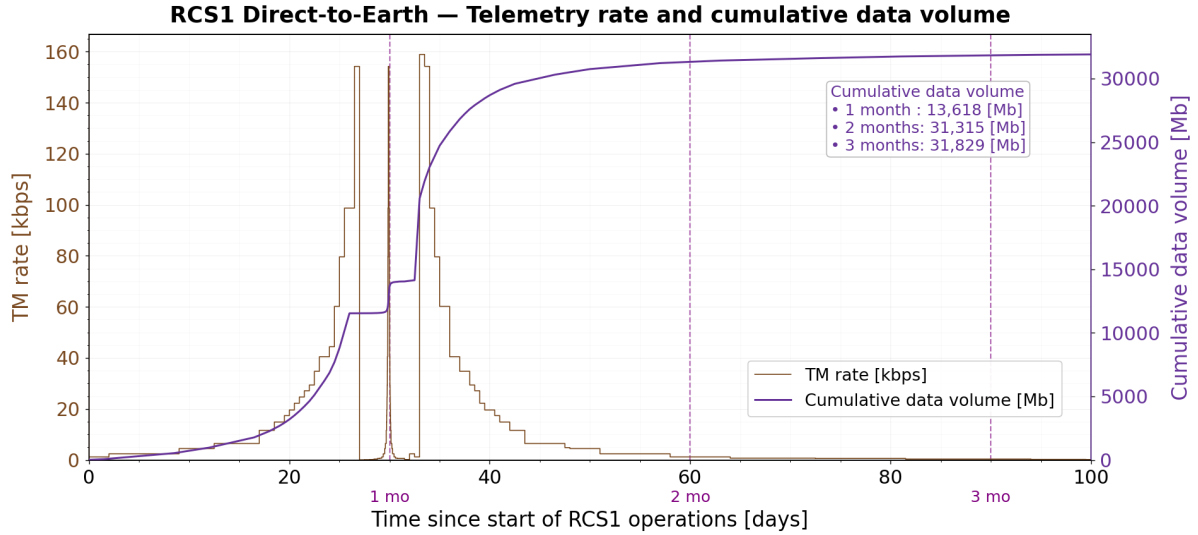


Figure 47: RCS1 DTE TM rate and cumulative data-volume

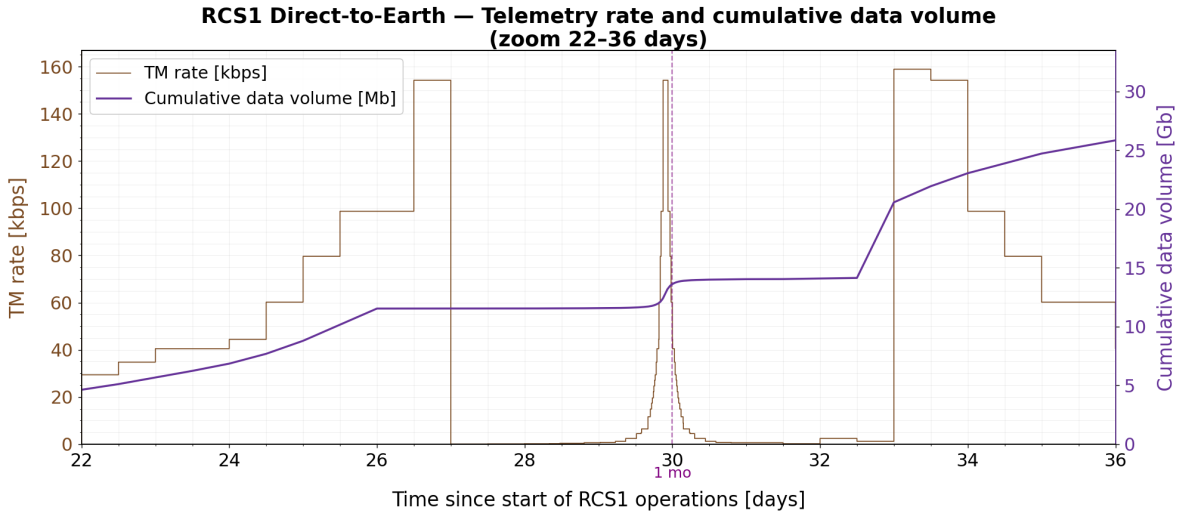


Figure 48: RCS1 DTE TM rate and cumulative data-volume – zoomed around closest approach

Figure 47 and Figure 48 show that RCS1 DTE data return is concentrated within a relatively short interval around closest approach. As the Earth–Apophis range decreases, higher telemetry rates become feasible, and the cumulative data-volume rises rapidly. Most of the total volume of data is returned within approximately ten days before and ten days after closest approach.

The instant of closest approach makes only a limited contribution to the cumulative data return, even though it corresponds to the minimum Earth–Apophis range. The highest-rate condition is very short-lived, while the main DTE data return comes from the surrounding pre- and post-encounter periods, where high telemetry rates remain feasible for longer durations.

DTE nominal communication summary

Overall, the nominal DTE analysis shows that RCS1 communication performance is driven mainly by the Earth–Apophis range evolution, with antenna gain, ground station performance and contact duty-cycle profiles superimposed on this trend. Carrier tracking remains feasible throughout the analysed timeline, and telemetry feasibility margins remain positive over the full mission duration. The resulting data return is strongly concentrated around the close-approach period, where the most favourable range conditions allow higher telemetry rate, with the main contribution coming from the surrounding pre- and post-encounter intervals, where high rates remain feasible for longer.

4.2.2 Inter-Satellite-Link Communication Performance

Unlike the DTE case, RCS1 ISL performance is governed mainly by the local geometry between RCS1 and the RAMSES spacecraft. The key variables are the inter-satellite distance (ISD) and the antenna viewing angles at both ends of the link. These quantities determine the received power and therefore the feasible ISL data mode, using the received-power and mode-feasibility relations defined in (3.34) to (3.36). The selected data mode is then integrated over time to obtain the resulting cumulative data return.

Based on the RCS1 ISL geometry defined in Section 2.2.1 and the model described in Section 3.2, Figure 49 shows the evolution of the inter-satellite distance together with the antenna angles at the RAMSES spacecraft and at RCS1 over the full RCS1 operational timeline.

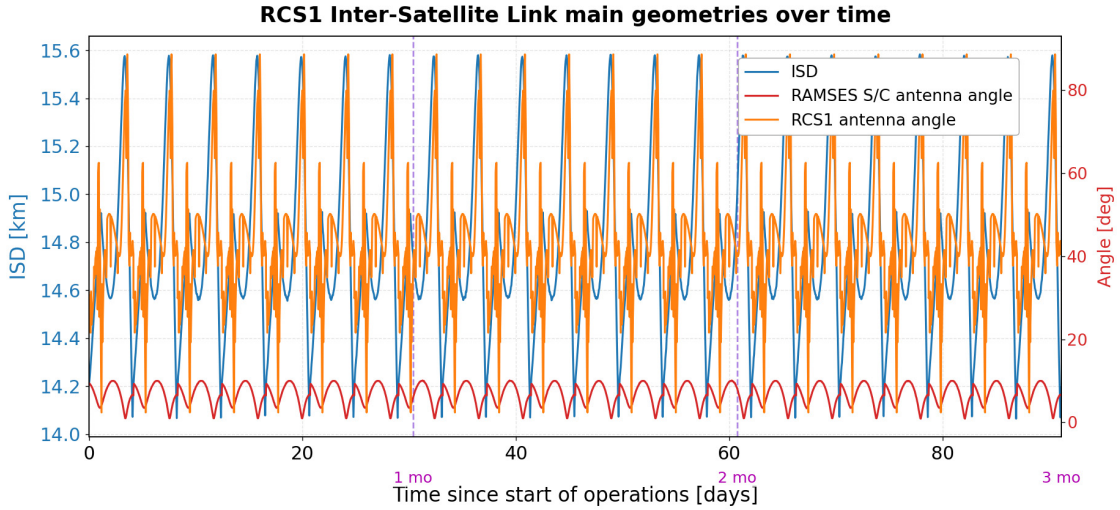


Figure 49: RCS1 ISL Link evolution over mission time

Figure 49 highlights a strongly repetitive geometric pattern over time. ISD remains within a relatively narrow interval around 15 km, varying only between approximately 14.1 km and 15.6 km.

The antenna angle at the RAMSES spacecraft remains close to boresight, typically within only a few degrees of the ISL antenna boresight. By contrast, the corresponding angle at RCS1 spans a much wider range, from near-boresight conditions to values approaching 90°. This asymmetry reflects the relative geometry adopted for the nominal ISL scenario: from the RAMSES spacecraft, RCS1 remains within a limited angular sector, whereas from the RCS1 perspective, the RAMSES spacecraft sweeps across a much broader part of the antenna pattern. The main geometric variation in the nominal ISL case is therefore associated with antenna viewing angle rather than inter-satellite distance.

Figure 50 provides a zoomed view of the repetitive geometric pattern for the first three cycles, allowing these variations to be inspected in more detail.

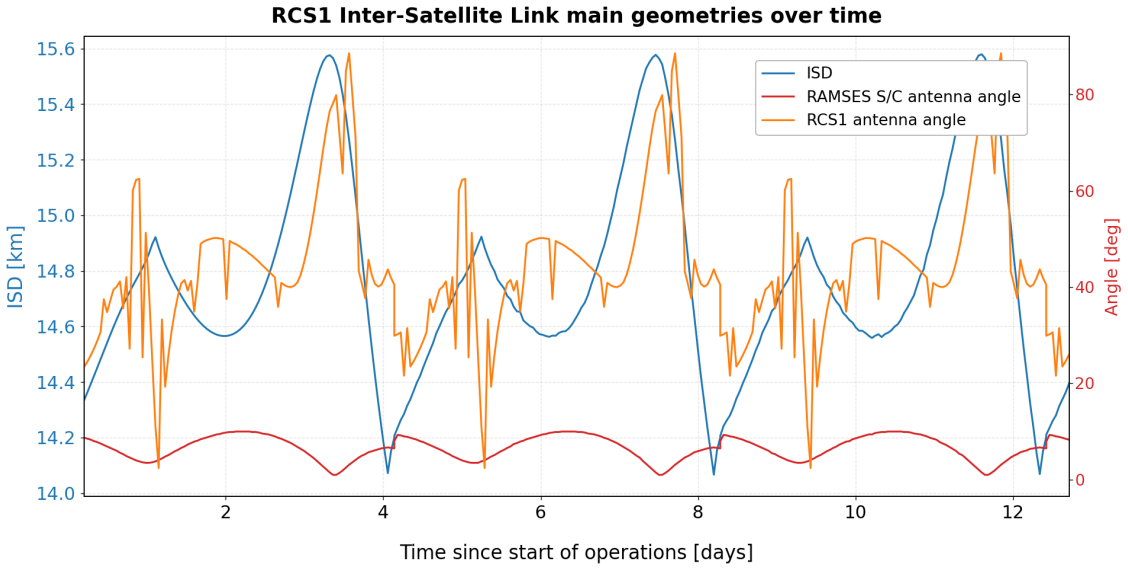


Figure 50: RCS1 ISL Link evolution zoom-in

Figure 50 confirms, over the first three cycles of the mission timeline, the same repeating geometry observed in Figure 49. The cycles have a characteristic duration of approximately four days. Over each cycle, the inter-satellite distance varies only modestly around the nominal 15 km separation, while the antenna viewing angle at RCS1 changes over a much wider range. The RAMSES spacecraft antenna angle remains comparatively close to boresight. The zoomed view confirms that the short-timescale geometric variation is mainly associated with the RCS1 viewing geometry rather than with changes in spacecraft separation.

To identify the main drivers to the ISL link budget variation, Figure 51 shows the time evolution of the varying terms contributing to the received-power computation of Equation (3.34) over the full RCS1 operational timeline. The plotted case represents one RCS1 antenna-orientation realisation from the antenna-cut sampling introduced in Section 3.6.2. Figure 52 provides a zoomed view of the same contributor terms over the first 14 days of the mission timeline.

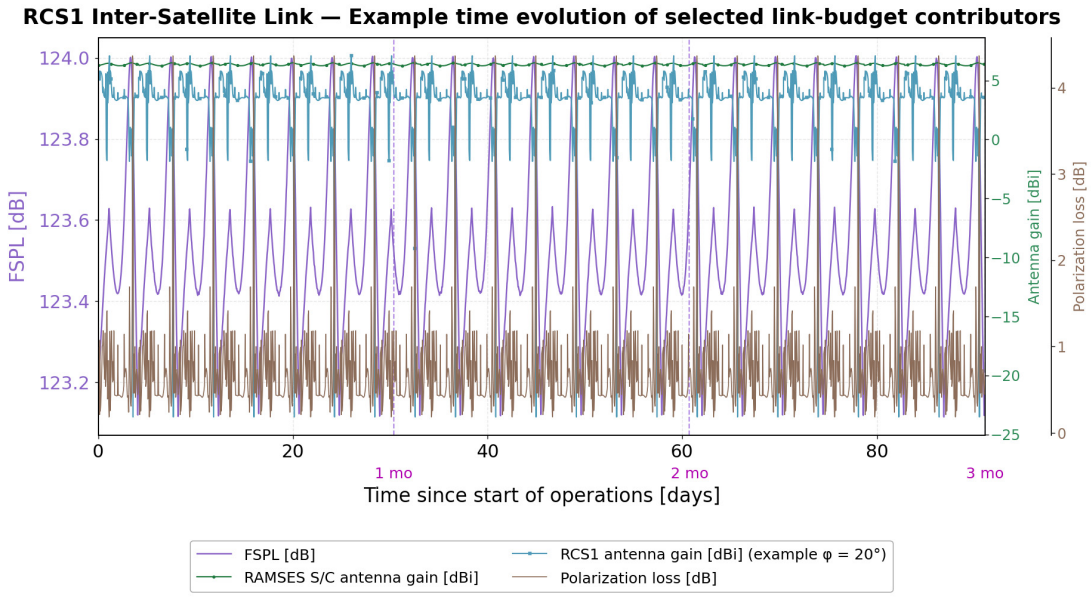


Figure 51: RCS1 ISL link budget contributors over mission time

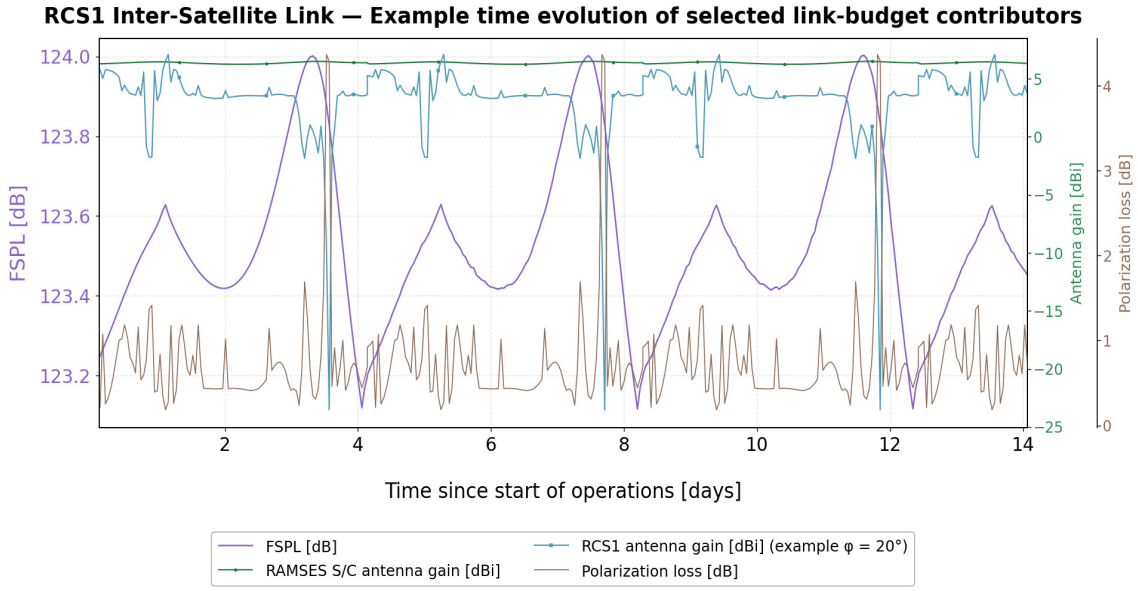


Figure 52: RCS1 ISL link budget contributors over mission time (zoom-in)

Figure 51 and Figure 52 show that the Free-Space Path Loss (FSPL), computed according to Equation (3.35), remains nearly constant over the full mission timeline, with peak-to-peak variation of less than 1 dB. This follows from the limited inter-satellite distance variation observed in Figure 49 and Figure 50. The RAMSES spacecraft antenna gain, derived from the pattern shown in Figure 41, also remains close to its boresight value, with only minor periodic fluctuations.

The main link budget variations therefore come from the RCS1 antenna gain and the polarisation loss. The RCS1 materialized antenna gain exhibits repeated excursions linked to the broad sweep of antenna angles described previously. The polarisation loss, computed according to Equation (3.21), follows the same periodic structure and increases during the intervals where the RCS1 antenna angle approaches the region in which the two combined low-gain antennas interfere more strongly. The resulting CubeSat antenna-gain variation is of the order of 30 dB, while the polarisation loss reaches peaks of almost 4.5 dB.

The resulting instantaneous ISL received power and achievable data mode are shown in Figure 53 and Figure 54 over the full mission timeline for two selected cases. Figure 53 reports a favourable received-power and data-mode behaviour, while Figure 54 reports an adverse case. The two cases correspond to two different values of antenna-azimuth, φ, introduced in Section 3.6.2 as the unresolved RCS1 antenna heading. In each figure, the received-power profile is computed using Equation (3.34), and the overlaid achievable ISL data mode is selected using the mode-feasibility criterion of Equation (3.36).

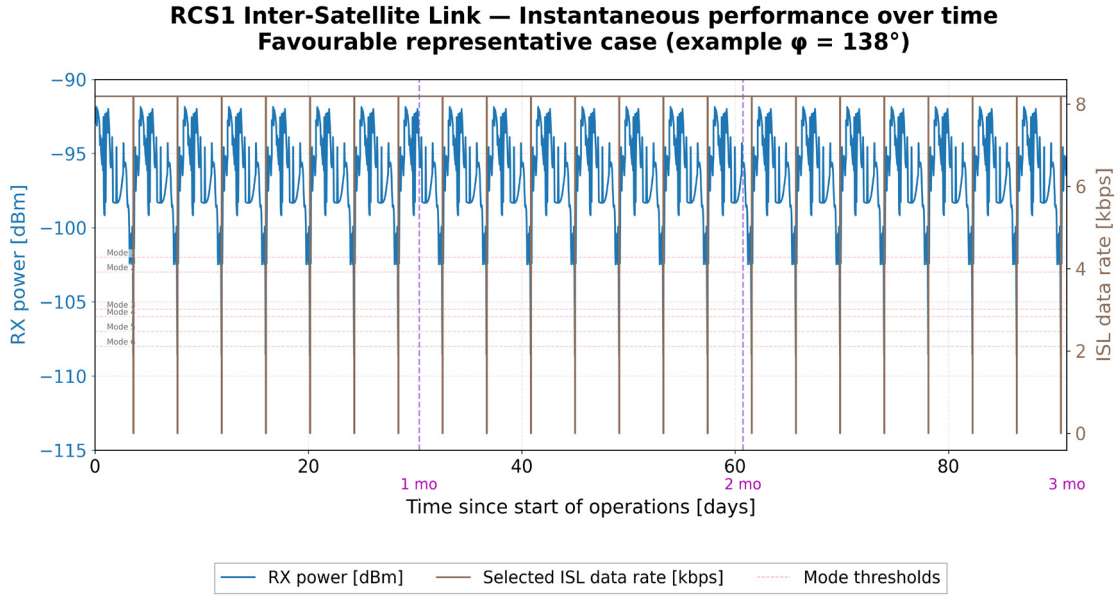


Figure 53: RCS1 ISL received power and instantaneous rate – favourable representative case

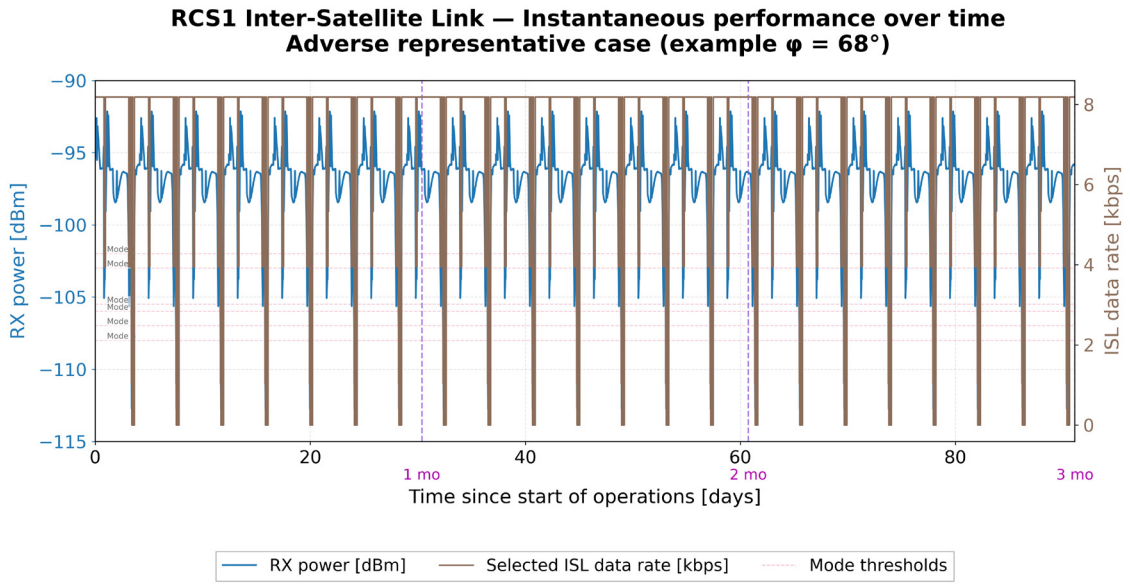


Figure 54: RCS1 ISL received power and instantaneous rate – adverse representative case

Figure 53 and Figure 54 show that the received-power behaviour remains strongly repetitive over the full RCS1 mission timeline, consistent with the periodic geometry and link budget contributors identified above. The favourable case remains at higher data modes for a larger fraction of the timeline, while the adverse case shows more frequent reductions in received power and drops in data mode.

A clearer comparison of these repeated variations is provided by the zoomed views in Figure 55 and Figure 56 for the same favourable and adverse representative cases over the first 14 days of RCS1 operations.

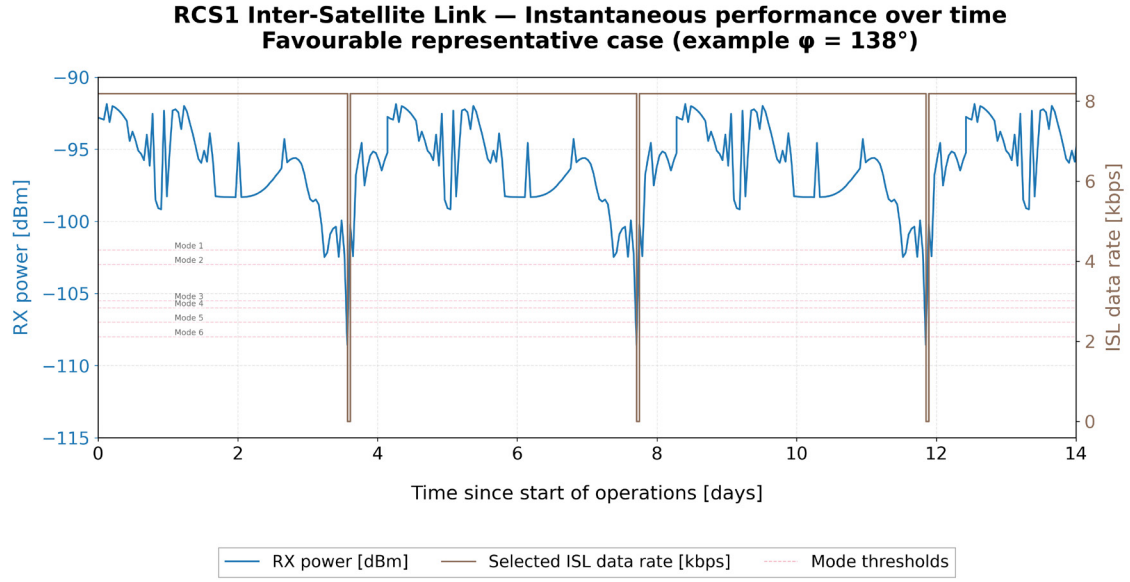


Figure 55: RCS1 ISL received power and instantaneous rate – favourable representative case (zoom-in)

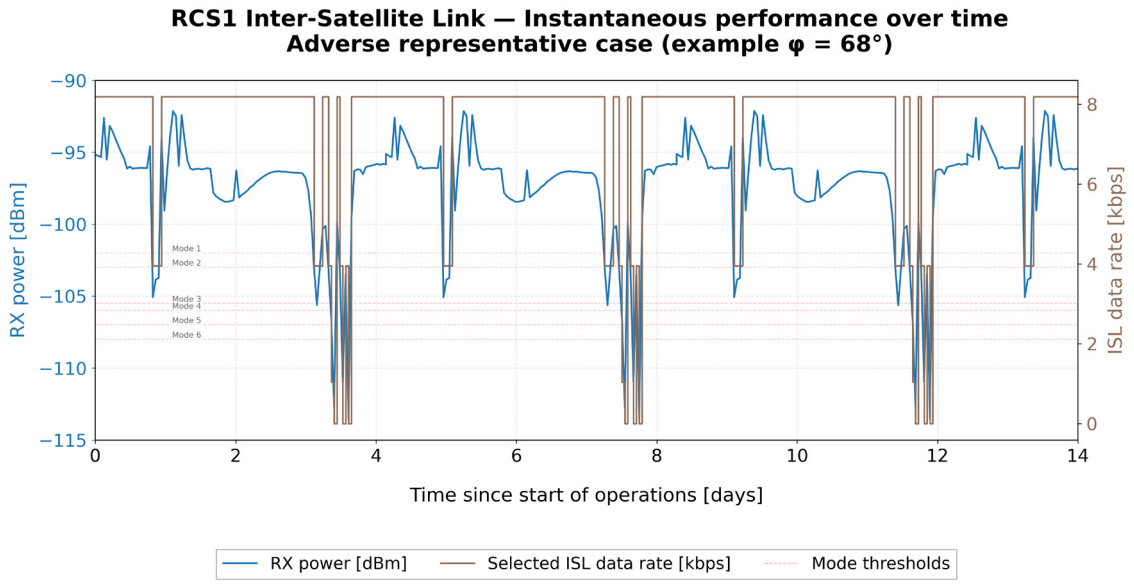


Figure 56: RCS1 ISL received power and instantaneous rate – adverse representative case (zoom-in)

Figure 55 shows that, for the favourable case, the received power remains above the highest-mode threshold for most of each cycle. Only short drops to lower data modes occur during the deepest periodic minima, and outage intervals remain limited to approximately 1% of the analysed time. In the adverse case shown in Figure 56, the same underlying cycle is visible, however, the received-power curve lies closer to the mode thresholds. This produces more frequent mode changes and repeated outage intervals, with individual outages lasting up to approximately 7 h.

The comparison indicates that the difference between the favourable and adverse cases is mainly caused by how the selected RCS1 antenna-azimuth case maps the repeated viewing-angle cycle onto antenna gain and polarisation loss. The ISL link therefore retains a stable periodic structure, but the specific antenna-azimuth case determines the fraction of time spent in the highest data mode, in lower modes or in outage.

Figure 57 shows the cumulative data-volume evolution obtained from the instantaneous ISL data rates across all antenna φ cases, using Equations (3.37) and (3.38). Best, median and worst realisations are highlighted.

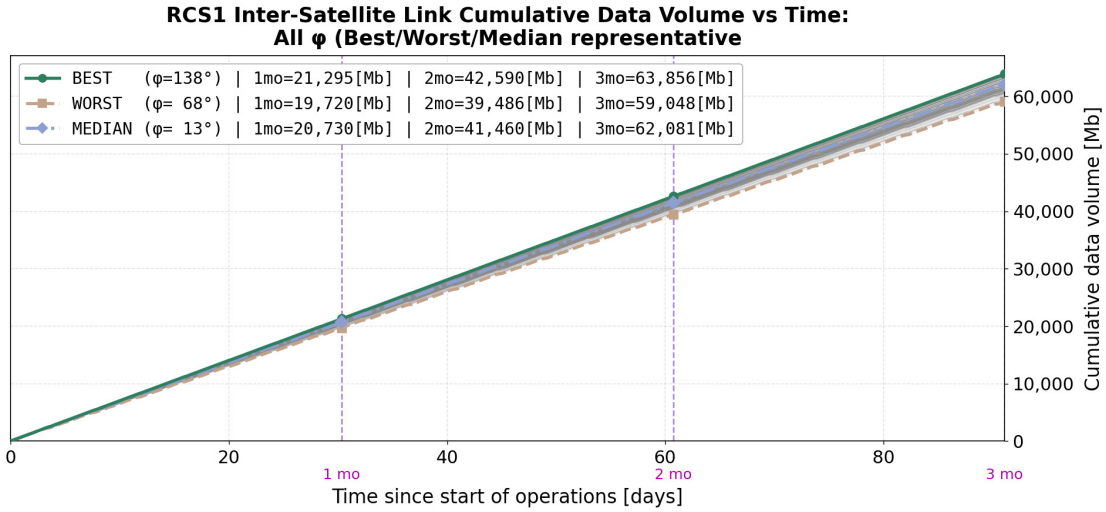


Figure 57: RCS1 ISL cumulative data-volume with antenna φ dispersion

Figure 57 shows that the cumulative data-volume curves grow almost linearly over the full RCS1 mission duration, with no visible plateaus or repeated-cycle signatures from the local geometry. The separation between the different antenna φ realisations therefore increases gradually as data-volume is accumulated.

Table 4.1 summarises the cumulative data-volume values reported for the three highlighted cases, including the absolute best–worst difference and the same difference expressed relative to the worst case.

Table 4.1 - RCS1 ISL cumulative data-volume at monthly mission markers

Time marker	Worst [Mb]	Median [Mb]	Best [Mb]	Best–Worst [Mb]	Best–Worst difference relative to worst [%]
1 month	19720	20730	21295	1575	8.0
2 months	39486	41460	42590	3104	7.9
3 months	59048	62081	63856	4808	8.1

Table 4.1 shows that the spread between best, median and worst φ realisations remains visible but limited throughout the RCS1 mission. The difference between the best and worst cases grows with time as small per-cycle differences in data rate and outage duration accumulate. However, the spread remains of the order of 8% relative to the worst case at the three monthly markers, indicating that the antenna-azimuth case introduces a secondary, rather than dominant, variation in nominal ISL performance.

Figure 58 shows the cumulative distribution function (CDF) of the cumulative ISL data-volume across the RCS1 antenna-azimuth population, at the three principal RCS1 time markers.

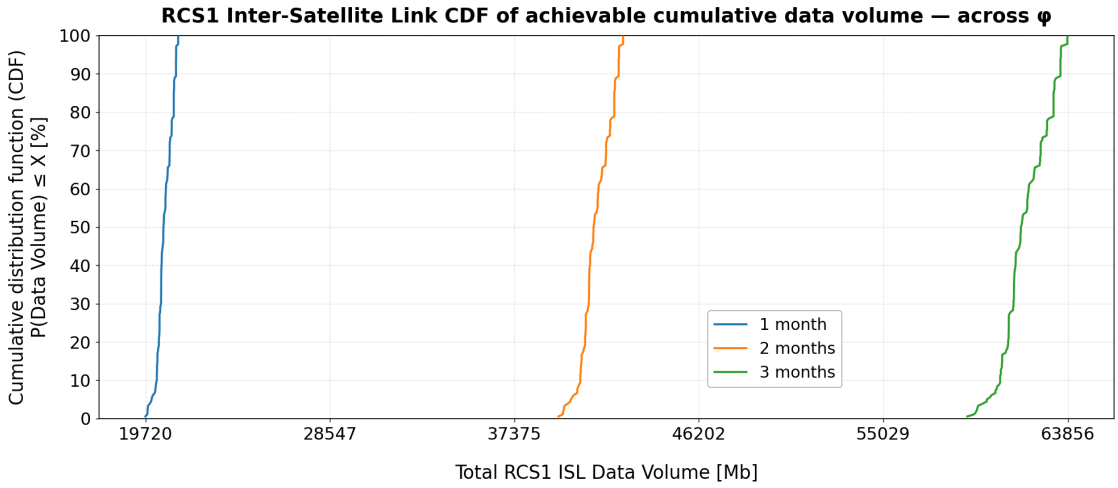


Figure 58: RCS1 ISL CDF of achievable cumulative data-volume across φ

Figure 58 shows that for each mission time marker, the spread remains noticeable but limited, confirming that the sampled φ case affects the final data return but does not dominate the nominal ISL performance.

The distributions also show the sustained accumulation of data-volume over the RCS1 mission duration, with the full population shifting to higher cumulative values at each successive time marker.

ISL nominal communication summary

Overall, the nominal ISL analysis shows that RCS1 link behaviour is driven mainly by local antenna geometry rather than by range evolution. The inter-satellite distance remains almost constant while the received power, selected data mode and cumulative data-volume are mainly affected by the variations in RCS1 antenna viewing. The sampled antenna-azimuth case introduces a moderate but non-negligible dispersion in nominal performance.

4.3RCS1 sensitivity analysis

The nominal results provide a reference for RCS1 communication performance, but operational constraints or possible degraded conditions can change the achievable return. For DTE, the sensitivity analysis considers RCS1 pointing capability that would affect the effective antenna gain and limits on the frequency at which the telemetry rate can be updated in the transmitter and receivers. For ISL, it considers how increased inter-satellite distance and less favourable relative antenna orientation affect relay-link performance.

Direct-To-Earth cases

The Direct-To-Earth communication link is driven by the Earth–spacecraft geometry together with the effective antenna gain and ground-segment performance. The nominal analysis showed how the available link margin changes over the RCS1 timeline; the following cases examine how this behaviour is affected when selected operational conditions are degraded or constrained.

In this context, two major aspects are investigated for the DTE case. First, operational constraints related to telemetry-rate selection are considered, since the telemetry configuration cannot be continuously adapted in practice and must instead be updated at discrete time intervals [20]. Second, the impact of degraded RCS1 antenna pointing conditions is analysed to assess the robustness of the link behaviour under partial or complete loss of attitude control.

Case 1: TM planning constraints

Telemetry-rate selection is subject to operational planning constraints as the telemetry configuration cannot be changed continuously with the instantaneous link condition; it must be updated at discrete times and kept consistent between the spacecraft and ground segment.

No mission-defined telemetry-rate update interval is available for the RAMSES CubeSats, nor could an applicable value be derived from the publicly available reference information. Therefore, 8 h, 12 h and 24 h intervals are introduced as sensitivity cases rather than as defined operational values. The 24 h case represents a daily operational timescale, while the 8 h and 12 h cases represent more frequent rate-update opportunities. Their purpose is to quantify the sensitivity of DTE performance to the update cadence rather than to reproduce a defined RAMSES operational strategy.

Figure 59 shows the unconstrained rate-selection case together with cases where telemetry-mode updates are allowed only every K hours. Using the DTE rate-selection and data-volume formulation of Equations (3.28) to (3.33), the figure shows the resulting cumulative data-volume and relative performance loss. Figure 60 provides a zoomed view of the closest-approach interval, where the link conditions vary most rapidly.

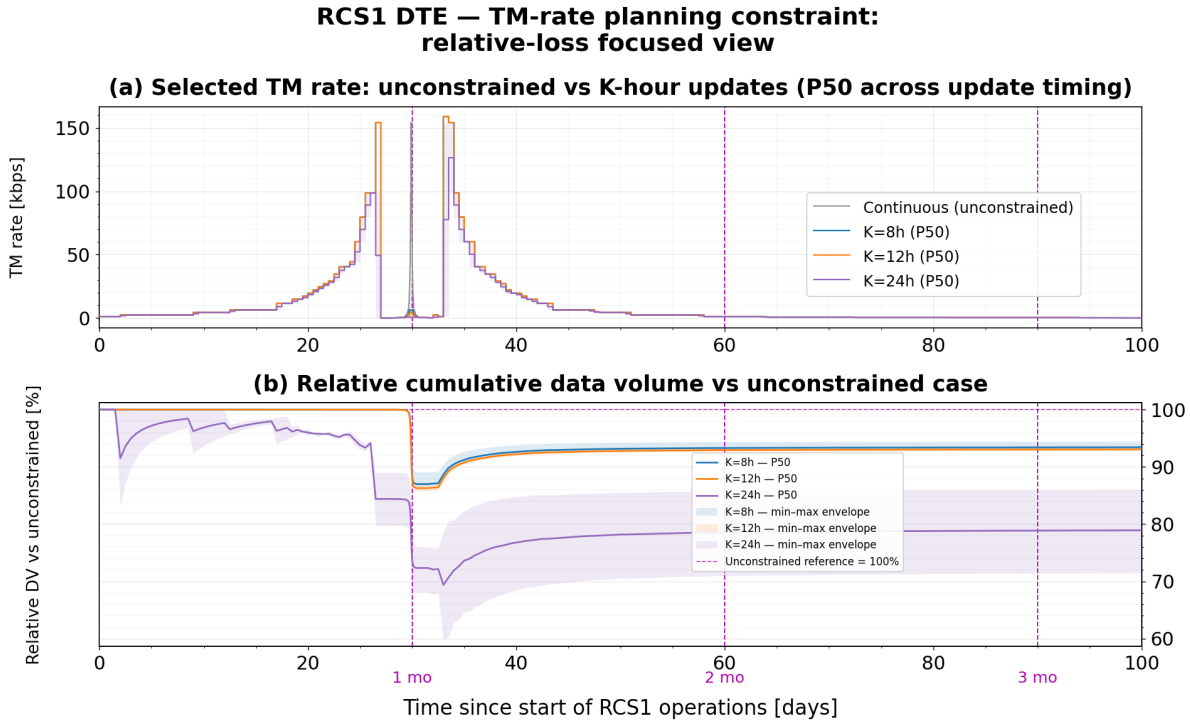


Figure 59: RCS1 DTE TM rate planning relative losses

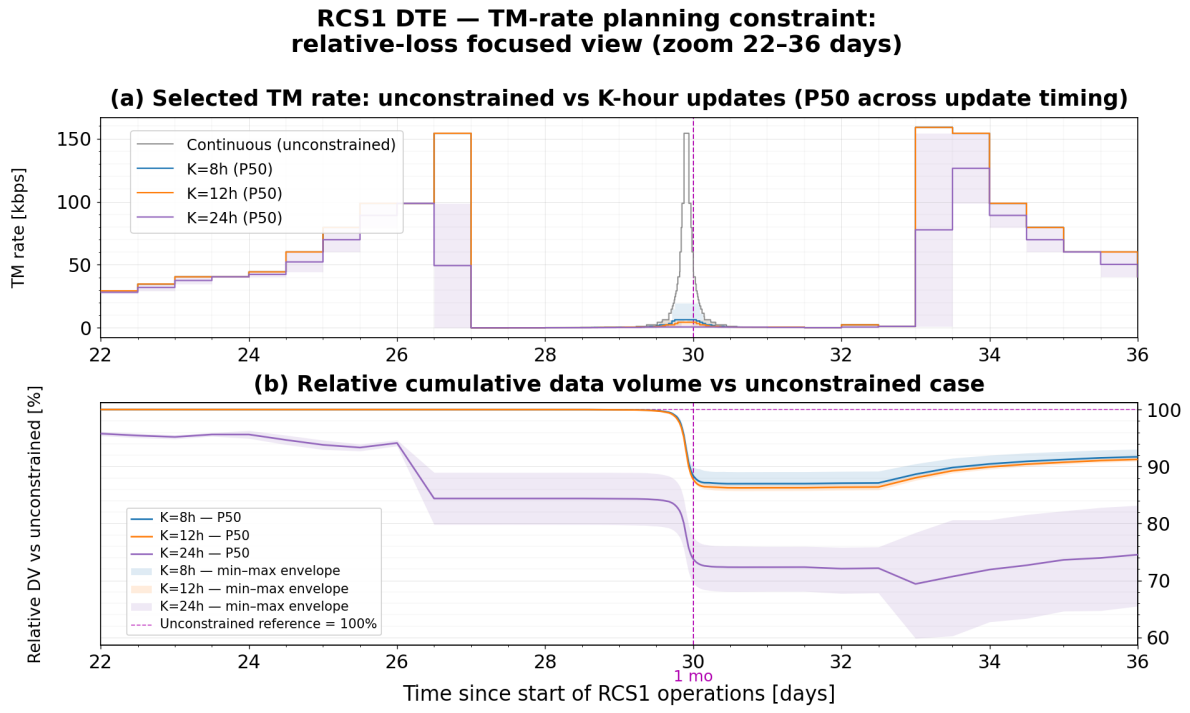


Figure 60: RCS1 DTE TM rate planning relative losses – zoomed timespan

Figure 59 and Figure 60 show that limiting the telemetry-rate update frequency mainly affects the short interval around closest approach, where the DTE link budget changes

rapidly. In the unconstrained case, the selected telemetry rate follows the available margin continuously. With *K* hour planning, the selected rate can only be changed at discrete update times. When the link is improving, a higher rate cannot be applied until the next allowed update instant. When the link is degrading, the rate must be reduced early enough to remain feasible until the next update instant, which requires the future link degradation to be predicted rather than only observed at the current epoch.

This creates a timing mismatch between the available link margin and the selected telemetry rate. This effect is most visible in Figure 60. Before closest approach, around day 26, the 24 h planning case cannot follow the sequence of increasing rates visible in the unconstrained case, so it remains at lower rates for longer. A similar effect occurs after closest approach, around day 33–34, where the constrained cases again miss part of the high-rate opportunity. Around the closest-approach epoch itself, the link conditions create a short peak of high TM rates in the unconstrained case that is not reproduced by the constrained ones.

The lower panels of Figure 59 and Figure 60 show that the resulting data-volume penalty remains moderate for the 8 h and 12 h update cases, but becomes larger and more timing-dependent for the 24 h case. Around closest approach, the 8 h and 12 h P50 curves fall to approximately 88–90% of the unconstrained cumulative data-volume, while the 24 h P50 curve drops to approximately 70–75%. The min–max envelopes show a wider dependence on update phasing, especially for the 24 h case, where the relative cumulative data-volume spans roughly 60–80% during the most affected interval. After the close-approach interval, the relative curves stabilise as the link evolution becomes slower. By the end of the 100-day window, the 8 h and 12 h cases remain close to approximately 88–90% of the unconstrained reference, while the 24 h case only partially recovers to approximately 65-85%.

Case 2: Direct-To-Earth Link performance under non-nominal geometry

Direct-To-Earth performance depends on the effective spacecraft antenna gain available in the Earth direction. The nominal DTE case relies on boresight antenna performance during the main communication periods, but degraded pointing can reduce the effective gain and therefore the available link performance.

To assess this sensitivity, the DTE link is evaluated using percentile levels of the full RCS1 DTE antenna pattern, with the nominal boresight case retained as the reference. The same DTE telemetry-feasibility and carrier-tracking framework introduced in Equations (3.19) to (3.31) is applied to each antenna-gain case. The resulting link behaviour is shown in Figure

61, where the upper panel reports the achievable telemetry rate and the lower panel reports the corresponding maximum PLL tracking bandwidth for the residual-carrier.

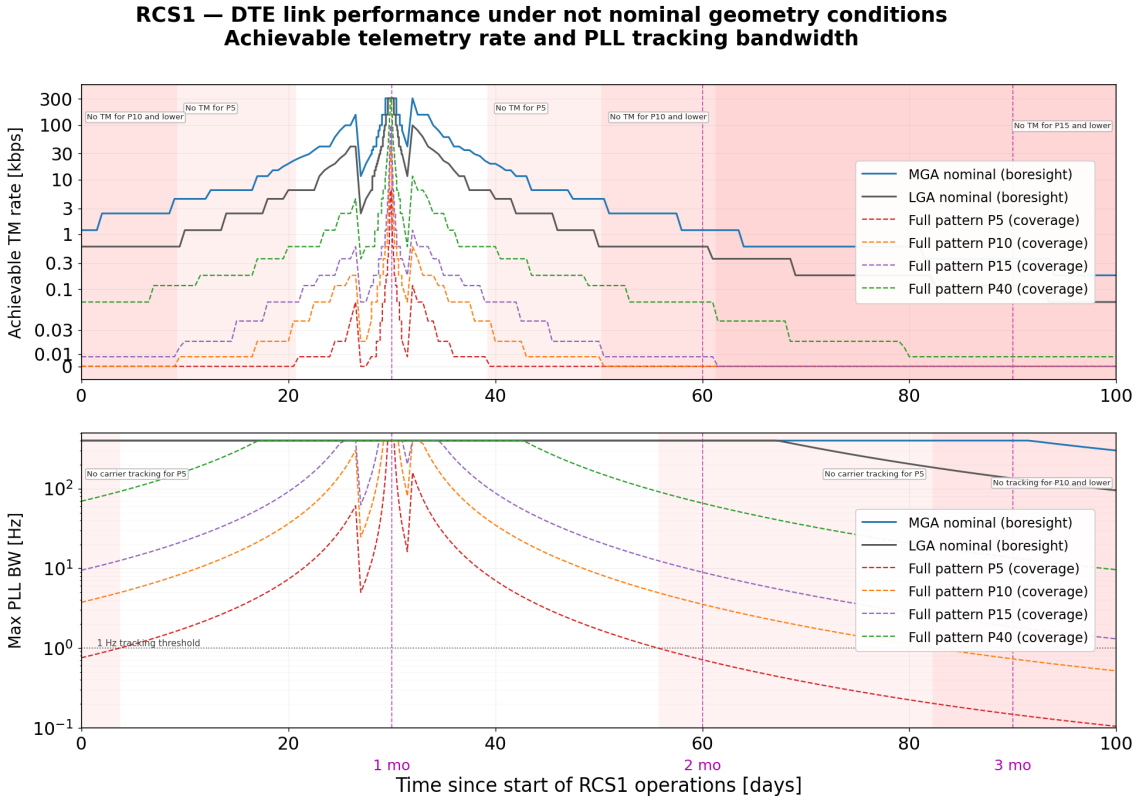


Figure 61: RCS1 DTE Link performance under degraded antenna-gain conditions

Figure 61 shows that RCS1 DTE performance is sensitive to degraded antenna-pointing conditions. The nominal MGA boresight case maintains telemetry capability over the mission timeline, while the LGA boresight case remains capable of sustained communication but with reduced access to higher telemetry rates.

When further antenna-gain degradation is introduced through the lower antenna-coverage percentiles, the achievable telemetry rate is reduced, especially at larger Earth distances.

In more adverse cases, the link becomes unfeasible over extended parts of the mission. This behaviour is most pronounced away from closest approach, where the higher free-space path loss is no longer compensated by favourable antenna gain. Around closest approach, the reduced Earth–spacecraft distance partly mitigates the degraded antenna condition and allows telemetry to be recovered in some cases. Carrier tracking is also affected when antenna-gain degradation becomes severe. The reduction in residual-carrier power lowers the maximum PLL tracking bandwidth, with the most degraded cases falling below the tracking threshold during parts of the mission. In these intervals, the DTE link would not support reliable carrier recovery, and communication would be unavailable.

Overall, the results indicate that the DTE link is not robust to degradation in antenna pointing when operating far from Earth. This represents a critical operational constraint, as attitude or pointing degradation during these phases may lead to a complete loss of communication capability for all the mission except a 20-40 day window centred at the closes approach.

Inter-Satellite-Link cases

Unlike the DTE case, the nominal ISL performance in Section 4.2 is driven mainly by the local RCS1–RAMSES geometry rather than by large-scale range evolution. In the nominal case, the inter-satellite distance remains close to the 15 km reference value, and the RAMSES spacecraft maintains favourable viewing conditions by pointing approximately towards the asteroid. Changes in the RCS1–RAMSES separation distance or relative pointing can alter the ISL path loss and antenna viewing geometry, affecting the achievable data mode and the fraction of time with link availability.

The ISL sensitivity analysis therefore considers two degraded relative-geometry cases. The first increases the inter-satellite distance while preserving a partially controlled geometry, in which RAMSES remains asteroid-pointed. The second removes this pointing capability and treats the mutual viewing geometry as fully unconstrained. These cases assess how ISL coverage and data-mode availability change as the separation and viewing conditions degrade. The same modelling framework used for the nominal RCS1 ISL analysis is applied.

Case 3: ISL robustness to increasing inter-satellite distance under constrained viewing geometry

This case examines whether the RCS1 ISL remains feasible as the inter-satellite distance increases beyond the nominal 15 km reference value, without assuming that RCS1 can keep its ISL antenna boresight directed toward RAMSES. The analysis therefore evaluates the full RCS1 spherical antenna coverage while retaining the nominal RAMSES-side viewing condition, in which RAMSES remains approximately asteroid-pointing.

Figure 62 shows ISL coverage as a function of inter-satellite distance for the available data modes under this constrained RAMSES-side geometry. On the RCS1 side, the link direction is evaluated over the full spherical antenna pattern, covering the possible θ and φ viewing angles. On the RAMSES side, RCS1 is assumed to remain within ±5° of the RAMSES ISL antenna boresight. The nominal 15 km reference distance is also indicated.

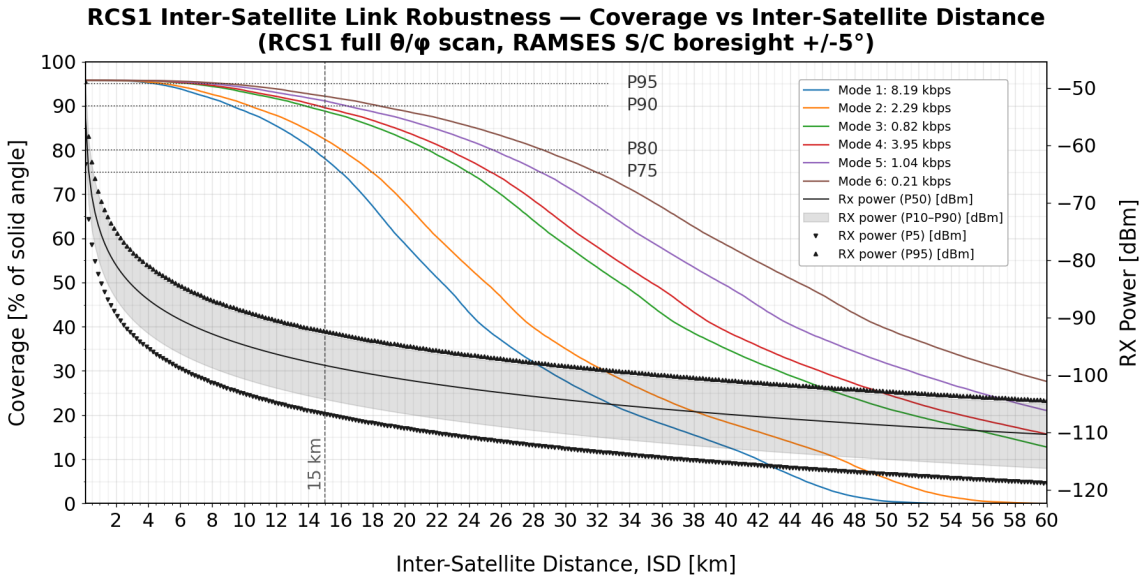


Figure 62: RCS1 ISL robustness under constrained geometry: coverage versus inter-satellite distance

Figure 62 first shows the received-power corridor across the RCS1 antenna-viewing cases. As the inter-satellite distance increases, this corridor shifts downward because of increasing propagation loss. Coverage for each data mode is then determined by the fraction of this received-power distribution that remains above the corresponding modem sensitivity threshold.

At the nominal reference distance of 15 km, the ISL retains substantial coverage for all data modes, with approximately 75–80% for the fastest *Mode 1* and above 90% for the slower modes. This indicates that, even under partially degraded pointing, the nominal proximity-operations geometry remains broadly compatible with maintaining the link. However, the highest data modes are available over a smaller fraction of the antenna-viewing conditions than the lower-rate modes.

Beyond the nominal distance, coverage decreases progressively as the received-power corridor moves downward. The fastest *Mode 1* is the first to degrade, with coverage becoming very limited approaching the 30 km range. Lower-rate modes retain larger coverage over longer distances, with the slowest mode still providing approximately 50% coverage to 45 km.

The constrained RAMSES-side viewing case therefore indicates that the nominal 15 km ISD remains well within a range of high link availability, with useful margin to moderate increases in separation. At larger ISD values, the link progressively shifts towards reduced coverage of the higher-rate modes and greater reliance on the lower-rate modes.

Case 4: ISL robustness under fully unconstrained mutual geometry

A more stringent robustness case is considered by removing the constrained RAMSES-side viewing condition applied in Case 3. In this case, neither spacecraft is assumed to point its ISL antenna boresight toward the other spacecraft. The link is therefore evaluated over the full mutual antenna-viewing range, allowing degraded antenna-gain conditions on both sides of the link. This represents the most adverse relative-geometry sensitivity case considered for RCS1 ISL.

Figure 63 shows the resulting coverage probability as a function of inter-satellite distance for the set of assumed data modes, with the nominal 15 km reference distance indicated.

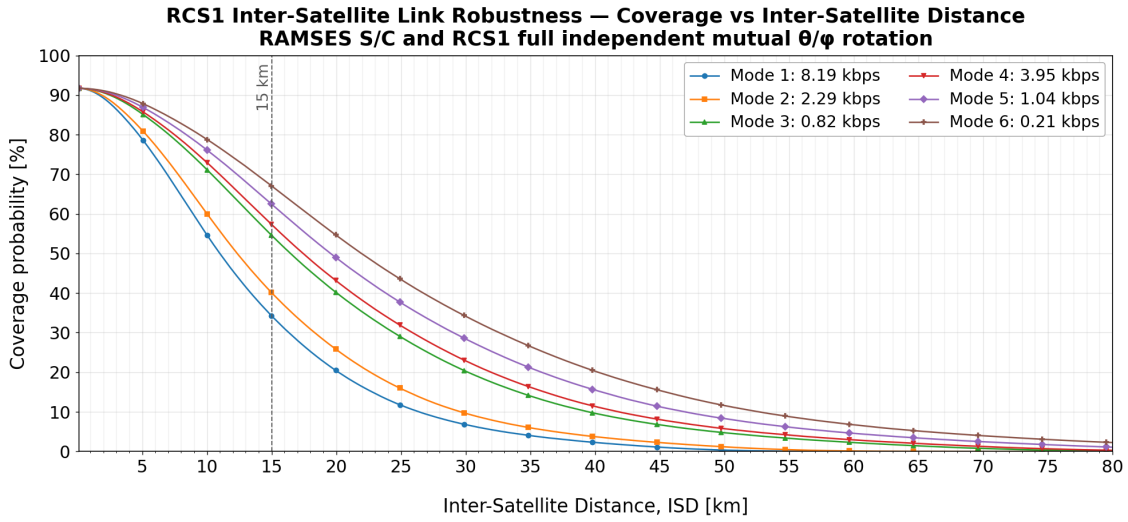


Figure 63: RCS1 ISL robustness under unconstrained geometry: coverage versus inter-satellite distance

Figure 63 shows that, under fully unconstrained mutual geometry, ISL coverage decreases rapidly as inter-satellite distance increases. At the nominal 15 km reference distance, coverage is already strongly mode-dependent, from approximately 35% for *Mode 1*, 40% for *Mode 2*, to 60–70% for the slowest *Modes 5* and *6*.

At the 15 km reference distance, the data-volume performance may therefore be affected by reduced access to the higher-rate modes, but the lower-rate modes still provide a relatively high probability of maintaining communication. This probability decreases with increasing ISD, although the lower modes retain approximately 45% coverage up to 25 km, corresponding to about 1.7 times the reference distance. At approximately 40 km, the coverage probability is below 20% for all modes, and by 50 km it falls below 10%.

This indicates that loss of mutual viewing control would not preclude ISL communication at the nominal separation, but would make the link increasingly dependent on lower-rate modes and shorter-range operating conditions.

RCS1 robustness summary

The sensitivity cases show that the two RCS1 communication architectures have different robustness drivers.

DTE performance is sensitive to Earth-pointing degradation and telemetry-rate planning constraints, especially during the rapidly varying close-approach interval. Degraded onboard antenna-pointing capability reduces telemetry performance and, in severe cases, can compromise carrier recovery when RCS1 is far from Earth. Telemetry-rate planning constraints add a separate limitation by preventing the selected rate from following rapid link budget variations, reducing the achievable data return during the most dynamic parts of the encounter.

For ISL, robustness depends mainly on maintaining the local RCS1–RAMSES relay geometry. Maintaining constrained geometry on RAMSES, the nominal 15 km case retains substantial coverage, although higher-rate modes are lost first as distance increases. With fully unconstrained mutual geometry, coverage is reduced even near the nominal distance.

Table 4.2 consolidates the RCS1 sensitivity and robustness results using the main indicators examined in Section 4.3.

Table 4.2 - RCS1 sensitivity and robustness summary

Sensitivity / robustness indicator	DTE response	ISL response	Main indicator
Mission geometry	Strong range-driven variation; data return concentrated around closest approach.	No sensitivity to Earth range; behaviour governed mainly by local RCS1–RAMSES geometry.	FSPL, selected telemetry rate or data mode, cumulative data-volume.
Antenna pointing / viewing geometry	Strong degradation under loss of MGA Earth-pointing; LGA or full-pattern percentile cases reduce telemetry rate and may compromise carrier recovery.	Robust while the local relay viewing geometry remains controlled; fully unconstrained geometry strongly reduces coverage.	TM-rate feasibility, residual carrier (C/N_0), PLL bandwidth and ISL coverage probability.
Telemetry-rate planning	Update constraints reduce data return mainly during the rapidly varying close-approach interval.	Not sensitive as data-mode selection is adaptive.	Relative cumulative data-volume loss.
Inter-satellite distance	Not relevant.	Coverage decreases with distance; high-rate modes are lost first, while lower-rate modes remain available over larger separations providing margin for operability.	Coverage probability versus ISD and data mode.

These results provide the RCS1 sensitivity basis for the architecture-level comparison in Chapter 6.

4.4Chapter summary

This chapter presented the orbiter RCS1 communication performance analysis for the DTE and ISL architectures using the modelling framework developed in Chapter 3. The analysis covered nominal mission conditions and selected sensitivity cases.

For DTE, nominal performance is driven by the Earth–Apophis range evolution. The free-space path loss decreases from approximately 255 dB at the start of operations to approximately 200 dB around closest approach, before increasing again to 263 dB by the end of the mission. Most of the DTE data-volume is returned within ten days before and after closest approach, with the cumulative return increasing from 4 Gbit to 27 Gbit over this period. The final cumulative return reaches 39.6 Gbit, but the concentration of return around closest approach may affect the achievable scientific return later in the RCS1 mission.

For ISL, nominal performance is driven by the local RCS1–RAMSES geometry. The inter-satellite distance remains close to the 15 km, with only minor FSPL variations. The link budget follows repeated local-geometry cycles of approximately four days, mainly through the RCS1 viewing angles and antenna response. The resulting cumulative data return increases linearly over the entire mission timeline, with a median value of approximately 20 Gbit at close encounter and 62.1 Gbit after three months with only limited dispersion around the median.

The sensitivity analysis shows that the two architectures have different robustness drivers. DTE is most affected by telemetry-rate update constraints during the rapidly varying close-approach interval and by degraded Earth pointing during the farther-range phases. For telemetry planning, 8 h and 12 h update intervals produce moderate losses, while the 24 h case reduces cumulative return to approximately 60–80% of the continuously adaptive reference. For degraded pointing, severe reductions in effective antenna gain can reduce telemetry performance and, in the most adverse cases, compromise carrier recovery.

ISL instead depends on maintaining a favourable local relay geometry. At the 15 km reference distance, coverage remains high when the RAMSES-side viewing direction is kept close to Apophis, with margin for moderate ISD increases. Under fully unconstrained mutual viewing geometry, high-rate access is still possible but limited at the nominal ISD, although lower-rate modes retain relatively high coverage up to 1.7 times the nominal ISD.

The chapter therefore addresses the RCS1 part of Objectives 2, 3 and 4 by establishing the nominal DTE and ISL performance, quantifying the main sensitivity cases, and characterising the dispersion associated with the RCS1 ISL antenna azimuth.

Chapter 5 Landed CubeSat (RCS2) Communication analysis

5.1Introduction

This chapter presents the communication analysis results for the landed RAMSES CubeSat, RCS2, using the modelling framework developed in Chapter 3. In contrast with the orbiting RCS1 case, RCS2 operates from the surface of Apophis after touchdown and its communication window is concentrated around the close-encounter period, within a short and strongly time-constrained mission profile.

As for RCS1, the RCS2 case is assessed for both candidate architectures: direct Earth return through DTE and relay through RAMSES via ISL. Because the post-touchdown state cannot be changed, communication performance depends on the realised landing conditions, including local visibility and surface geometry. This is especially important for RCS2 because the surface mission is short and battery-limited, making timely data return a primary need.

The analysis evaluates both architectures under nominal assumptions and selected sensitivity cases. In line with RCS1 analysis, the main quantities considered are geometric visibility, link feasibility, achievable telemetry rate or ISL data mode, coverage and cumulative data-volume.

The chapter is structured as follows. Section 5.2 presents the nominal communication performance, beginning with a geometric visibility characterisation common to both architectures before examining the nominal DTE and ISL performance separately. Section 5.3 investigates sensitivity to selected geometric, operational and design-related uncertainties. The key observations derived from the analysis are summarised in Section 5.4 and later used in Chapter 6 for the architecture-level evaluation.

5.2RCS2 Nominal Performance across landed population

5.2.1 Geometric visibility characterisation

Before assessing link performance, the geometric visibility of Earth and RAMSES from the RCS2 landing sites is characterised. The analysis uses the TS-axis rotational model and initial phase offset defined in Equations (3.3) and (3.4), with landing-site position and local surface normal obtained from Equation (3.5). Visibility is evaluated through the zenith-angle and minimum-elevation condition of Section 3.3, using Earth as the counterpart for DTE and

RAMSES for ISL. The same landing-site population is used for both architectures, ensuring that differences between the DTE and ISL results are not introduced by different surface-location samples. Because both landing site and initial rotational phase vary across the analysed population, the results are first presented statistically through cumulative zenith-angle distributions and then illustrated through representative time histories.

Direct-To-Earth zenith-angle cumulative distribution

The Direct-To-Earth visibility is first examined through the cumulative distribution of zenith angle over the first 40 h of surface operations. Figure 64 shows these distributions across the analysed landing-site population and initial rotational phases.

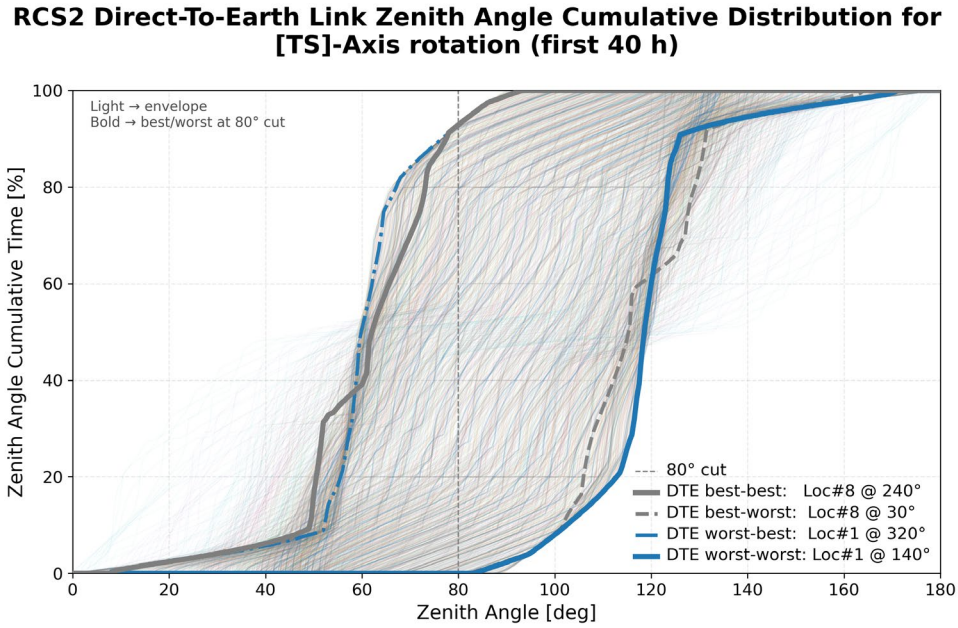


Figure 64: RCS2 DTE zenith-angle cumulative distribution (first 40 h)

Figure 64 shows that DTE visibility varies strongly across landing sites and initial rotational phases. Since the figure is shown as a cumulative distribution, curves that accumulate at lower zenith angles correspond to cases that spend a larger fraction of the time in favourable Earth-viewing conditions. Values further to the left of the $\theta_z = 80°$ limit therefore indicate longer and more favourable Earth visibility, while values close to or beyond this limit indicate near-horizon or non-visible periods. Since the nominal RCS2 antenna boresight is aligned with the local zenith, lower zenith angles also correspond to more favourable antenna geometry.

The four highlighted cases refer to the best- and worst-ranked landing locations, each paired with its most and least favourable initial rotational phase. In each label, the first term identifies the landing-site ranking, while the second term identifies the phase condition for

that same site. The best–best and best–worst cases therefore use the best-ranked landing site, paired respectively with its most and least favourable initial phase. The worst–best and worst–worst cases use the same convention for the worst-ranked landing site.

The relative positions of these cases show that the initial phase can strongly change the visibility outcome. A favourable phase can make an otherwise poor site perform close to the best case, while an unfavourable phase can bring a good site close to the worst case. This confirms that DTE visibility for RCS2 cannot be characterised by landing site alone; the initial rotational phase must also be included in the performance assessment.

The visibility behaviour shown in Figure 64 is therefore used to select representative DTE cases for the following time-domain analysis. These cases illustrate how the statistical visibility spread translates into specific Earth-visibility windows during the short RCS2 surface mission.

Inter-Satellite-Link zenith-angle cumulative distribution

The Inter-Satellite-Link visibility is examined using the same cumulative zenith-angle approach, but with RAMSES as the communication counterpart. Figure 65 shows the zenith-angle distributions over the first 40 h of surface operations across the analysed landing sites and initial rotational phases. The same four representative ranking combinations are highlighted to show how the best and worst landing sites behave under their most and least favourable initial phases. The cases with favourable RAMSES visibility are not necessarily the same as those with favourable Earth visibility, because the relevant direction is now the RAMSES spacecraft rather than Earth.

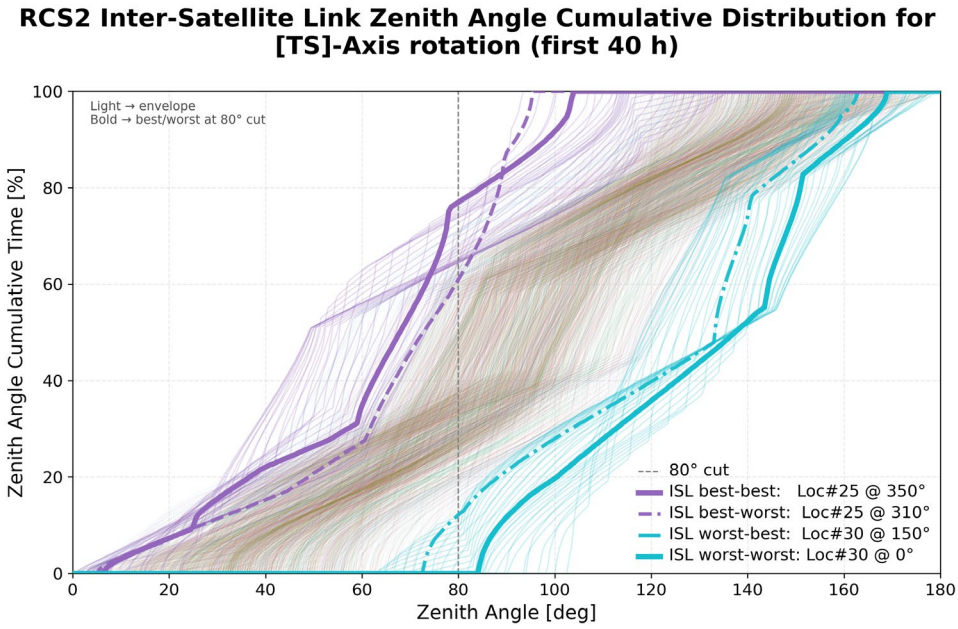


Figure 65: RCS2 ISL zenith-angle cumulative distribution (first 40 h)

Figure 65 shows a clear separation between favourable and unfavourable ISL landing-site cases. The best–best and best–worst cases both remain in the more favourable part of the distribution, while the worst–best case remains close to the worst–worst behaviour. This indicates that, for the ISL cases shown, landing-site geometry has a stronger influence than the initial rotational phase.

This result confirms that RCS2 relay performance must be characterised using ISL-specific landing-site and phase combinations, rather than assuming that the DTE representative cases are also representative for ISL.

Time-dependent DTE geometry: landing-site and phase effects

While the cumulative distributions show the overall spread, they do not explicitly separate the effects of landing-site location and initial rotational phase. Representative visibility histories are therefore used to make these two effects clearer. Figure 66 illustrates this for the DTE case by comparing the best and worst initial rotational phases for the same landing site, and two different landing sites under the same phase condition. Solid segments indicate periods when Earth is above the minimum-elevation limit, $\theta_z = 80°$, while dashed segments indicate periods below the visibility limit. The shaded intervals identify periods when one or both of the compared cases are not visible.

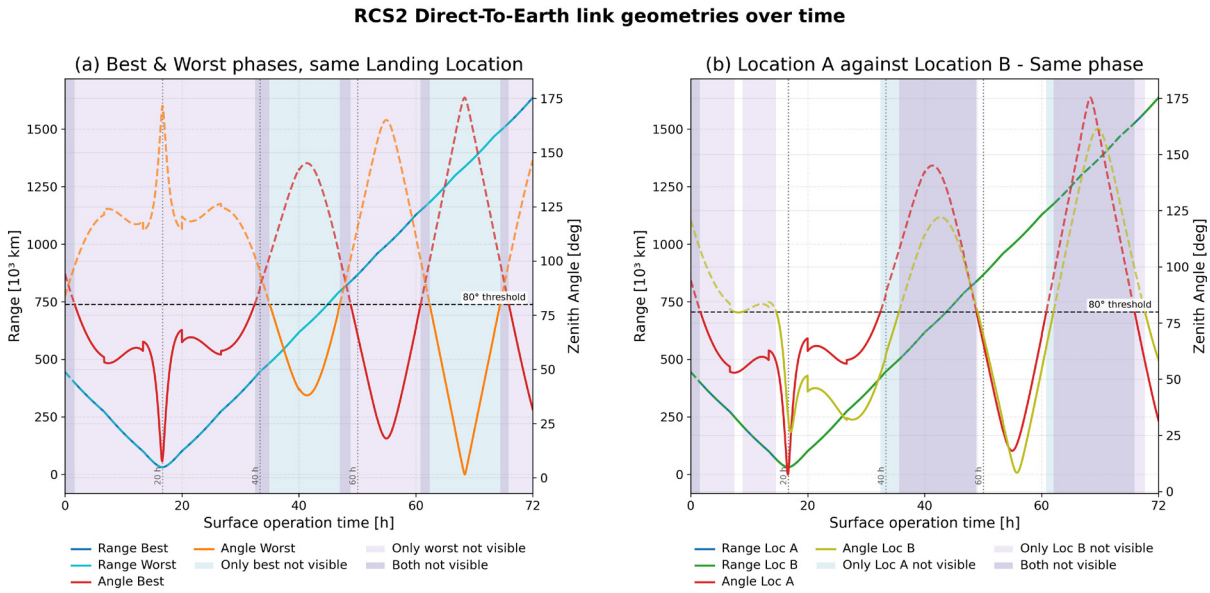


Figure 66: RCS2 DTE time-dependent geometry (phase vs site comparison)

Figure 66 shows that landing site and initial rotational phase both affect RCS2 DTE visibility, but in different ways. Panel (a) compares two initial rotational phases for the same landing site. The different phase shifts the timing of the visibility windows through the phase offset defined in Equation (3.4). The figure shows that the same location can either exploit or miss the most favourable Earth–Apophis range interval: in the favourable initial phase, visibility occurs during most of the low-range interval. In the unfavourable phase, visibility is shifted towards less favourable parts of the surface operations window.

Panel (b) compares two landing sites under the same initial phase. The two sites have different local surface normals, and therefore different zenith-angle histories and visibility windows. The comparison shows that landing-site geometry changes not only when Earth is visible and for how long, but also how favourable the antenna geometry is during those visible intervals.

RCS2 DTE visibility is therefore dictated by the combined effect of landing-site geometry and initial rotational phase.

Time-dependent ISL geometry: landing-site and phase effects

The ISL case is examined using the same comparison approach defined for the DTE geometry. Figure 67 adopts the same visual structure as Figure 66 and shows representative RAMSES-visibility histories, comparing the best and worst initial rotational phases for the same landing site, and two different landing sites under the same phase condition.

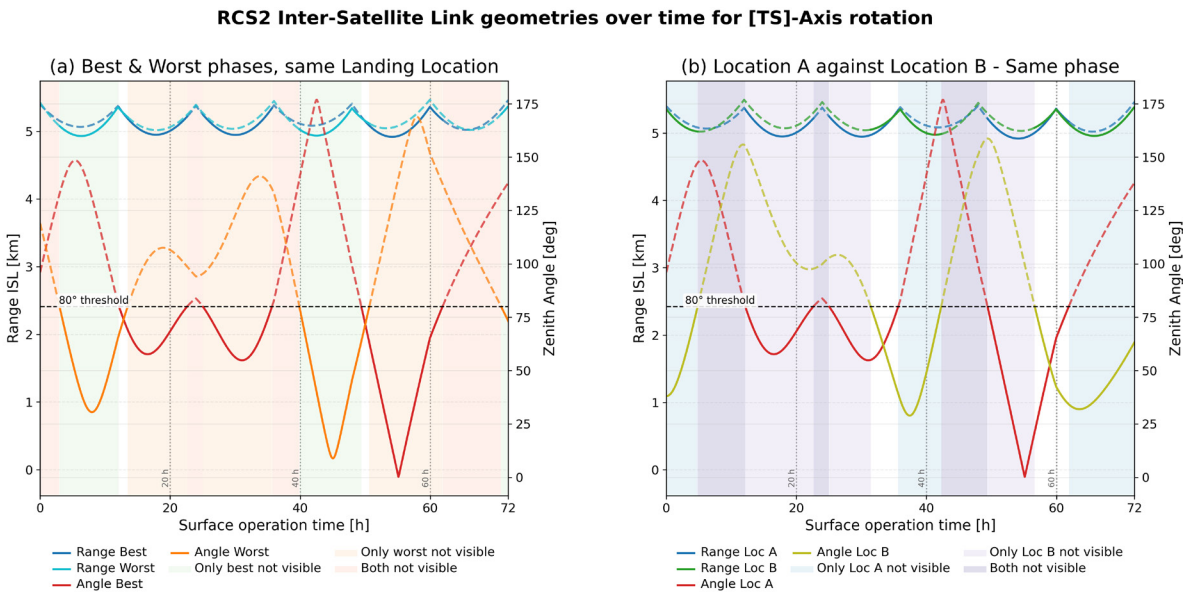


Figure 67: RCS2 ISL time-dependent geometry (phase - site comparison)

Figure 67 shows that landing site and initial rotational phase also affect ISL visibility.

Panel (a) compares two initial rotational phases for the same landing site. Changing the initial phase shifts the timing of the surface-to-RAMSES visibility windows. Unlike the DTE case, this timing is not related to capturing a short favourable communication interval; it reflects the local visibility geometry between the landed CubeSat and RAMSES.

Panel (b) compares two landing sites under the same initial phase. The two sites have different local surface normals, and show different, even though related, zenith-angle histories and visibility windows. The comparison shows that ISL contact opportunities depend on both the realised landing site and the initial rotational phase.

The ISL visibility pattern is therefore set by the direction to RAMSES along its proximity trajectory.

Figure 64 to Figure 67 therefore show that RCS2 visibility is both time-dependent and population-dependent for the DTE and ISL. The statistical distributions describe the spread across the landing-site and phase population, while the time histories show how specific favourable and unfavourable cases arise. These visibility results provide the geometric basis for the DTE and ISL link budget and data-volume analyses in the following subsections.

5.2.2 Direct-To-Earth Communication Performance

Following the visibility analysis, the RCS2 DTE link budget analysis combines the large-scale Earth–Apophis range evolution with the local surface-geometry effects of the landed antenna. The Earth–Apophis distance drives the free-space path loss through Equation (3.20).

At the same time, the changing Earth direction affects the antenna gain and polarisation mismatch loss, with the latter evaluated using Equation (3.21). Figure 68 shows these main link budget contributors over time for one representative landing case, together with the corresponding normalised $(E_b/N_0)_{\text{norm}}$, as defined in Equation (3.23). Shaded regions indicate periods where the DTE link cannot materialize due to minimum elevation constrain violation.

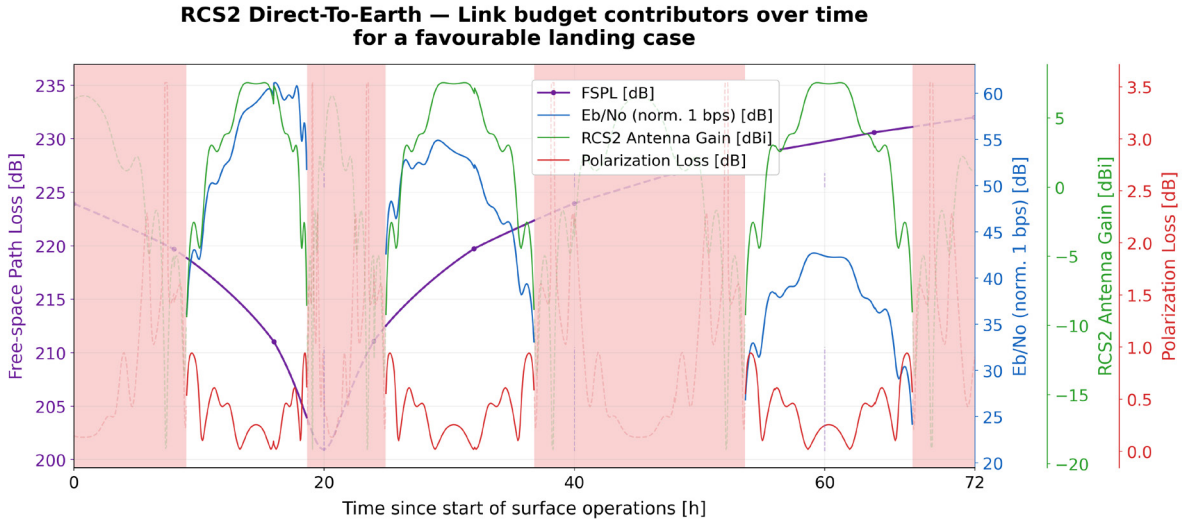


Figure 68: RCS2 DTE Link Budget contributors evolution

Figure 68 shows that the main DTE link budget trend is set by the variation in FSPL. Over the first part of the surface operation timeline, FSPL decreases from approximately 225 dB to approximately 201 dB as RCS2 approaches closest approach. This corresponds to the Earth–Apophis distance decreasing from approximately 500,000 km to approximately 32,000 km. After closest approach, the FSPL increases again, as the Earth–Apophis distance grows, to approximately 230 dB at the 60 h surface operations marker.

This propagation trend is effectively the same Earth–Apophis range effect observed for RCS1 over the close-encounter interval.

The local landed geometry then modulates this range-driven trend. As Apophis rotates, the antenna gain varies by several dB during the visibility intervals, while polarisation loss also changes with geometry and becomes largest close to visibility-window edges.

The normalised $(E_b/N_0)_{\text{norm}}$ therefore follows the favourable range evolution where visibility is available, but the closest-approach advantage is not fully exploited in the shown case because the minimum-range period overlaps with a local-horizon outage. The figure therefore illustrates the main RCS2 DTE constraint: the short Earth-range advantage around

closest approach can strongly improve $(E_b/N_0)_{\mathrm{norm}}$, but the realised landed geometry determines whether that advantage can actually be exploited.

The residual-carrier condition is assessed for the same representative landing case and therefore uses the same geometric visibility conditions as Figure 68.

Figure 69 shows the residual-carrier C/N_0, computed using Equation (3.30), together with the maximum PLL tracking bandwidth estimated from Equation (3.31) and visibility caps with shaded regions.

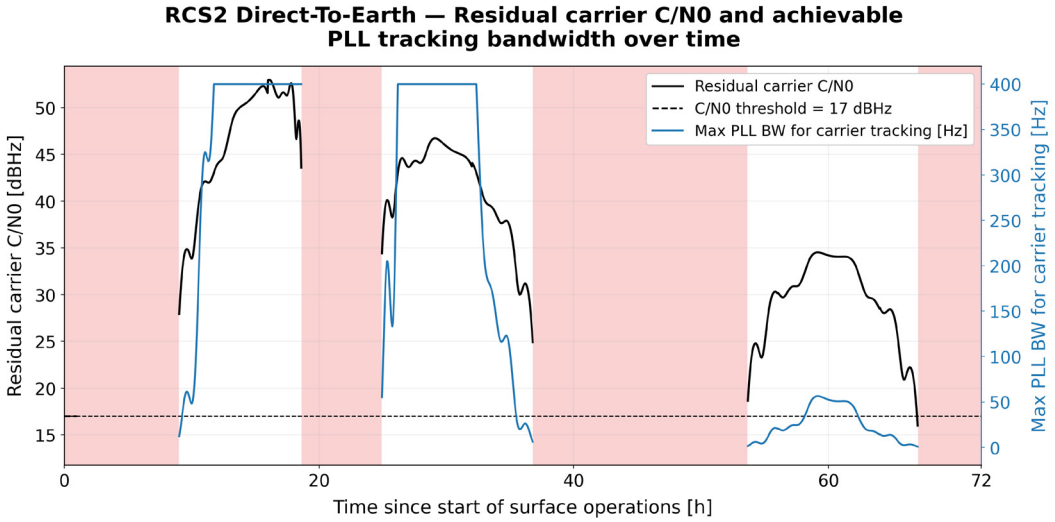


Figure 69: RCS2 DTE Residual carrier C/No and PLL tracking bandwidth

Figure 69 shows that the residual-carrier behaviour follows the same broad trend observed in the main DTE link contributors. During the first two visible intervals, the residual-carrier C/N_0 remains well above the adopted carrier-tracking threshold, reaching values above approximately 50 dBHz in the most favourable parts. The corresponding maximum PLL bandwidth is therefore comparatively large, indicating that carrier tracking is not a limiting condition during these main usable intervals.

The range-driven improvement is visible in the rising residual-carrier level before the outage, but the best propagation condition cannot be exploited continuously. The final visible interval around 60 h of surface operations is less favourable, with lower residual-carrier C/N_0 and reduced PLL tracking bandwidth because the Earth–RCS2 range has already increased. Carrier recovery is not a limiting factor during the main visible intervals, but the later reduction in PLL bandwidth shows that carrier-tracking robustness also decreases after the closest-approach opportunity has passed.

Figure 70 shows the telemetry feasibility margins for the available Direct-To-Earth modes still for the same representative landing case. The selectable telemetry-mode set is shown over the RCS2 mission timeline, with the Earth–RCS2 range to highlight the closest-approach region.

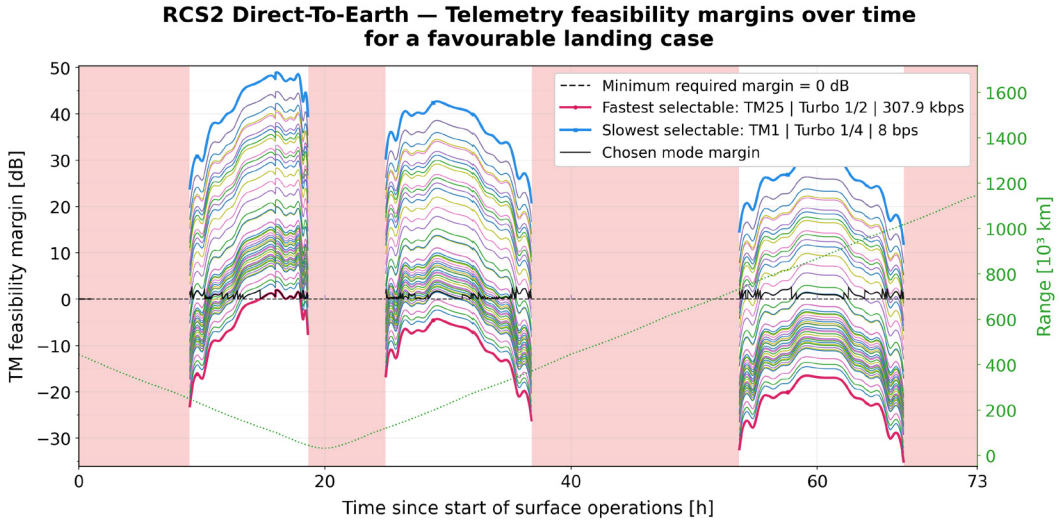


Figure 70: RCS2 DTE TM feasibility margins over mission time

The telemetry feasibility margins in Figure 70 are derived from the DTE formulation defined in Chapter 3. The available E_b/N_0 is obtained using Equations (3.23) and (3.24), and compared with the required threshold including the adopted design margin from Equations (3.25) to (3.27). The figure shows how the time-varying DTE link budget translates into the set of selectable telemetry modes over the RCS2 mission timeline.

The highest telemetry margins occur during the visible intervals adjacent to closest approach, where the reduced Earth–RCS2 range provides the most favourable propagation conditions. However, the local-visibility outage around closest approach prevents the minimum-range condition from being fully exploited: shortly before the 20 h mark, the faster TM rate has positive feasibility margin, while after the visibility gap a portion of the TM-rate set is already unfeasible. Within each visibility interval, the margin variation also reflects the changing antenna geometry as Apophis rotates, with less favourable antenna gain occurring near line-of-sight limits.

After closest approach, the margin corridor shifts downward as the range increases, and the selected mode progressively degrades.

To complement the representative DTE case shown so far, the best and worst landing locations in terms of achievable data-volume at the 40 h marker are analysed.

Figure 71 shows the best-ranked landing location, while Figure 72 shows the worst-ranked landing location. In each figure, the top panel shows the most favourable initial rotational-phase realisation for that location, and the bottom panel shows the least favourable realisation.

The time histories show the key DTE communication quantities for each case: instantaneous selected telemetry rate, cumulative returned data-volume and maximum PLL tracking bandwidth.

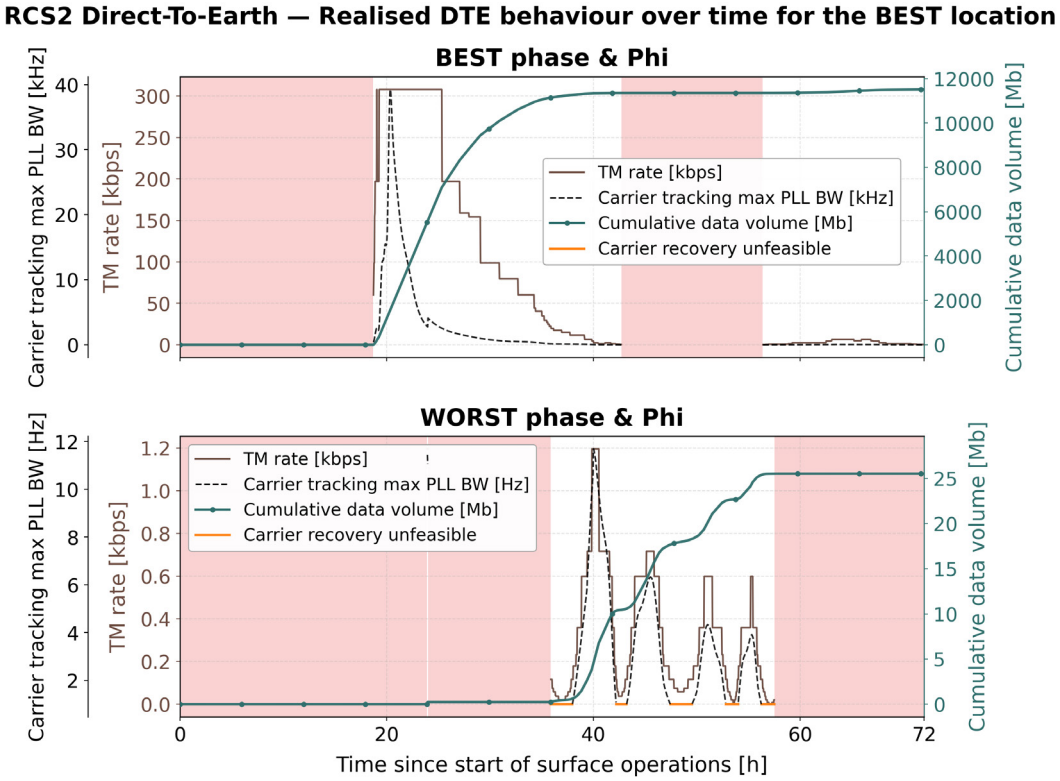


Figure 71: RCS2 DTE realised behaviour over time for the best landing location

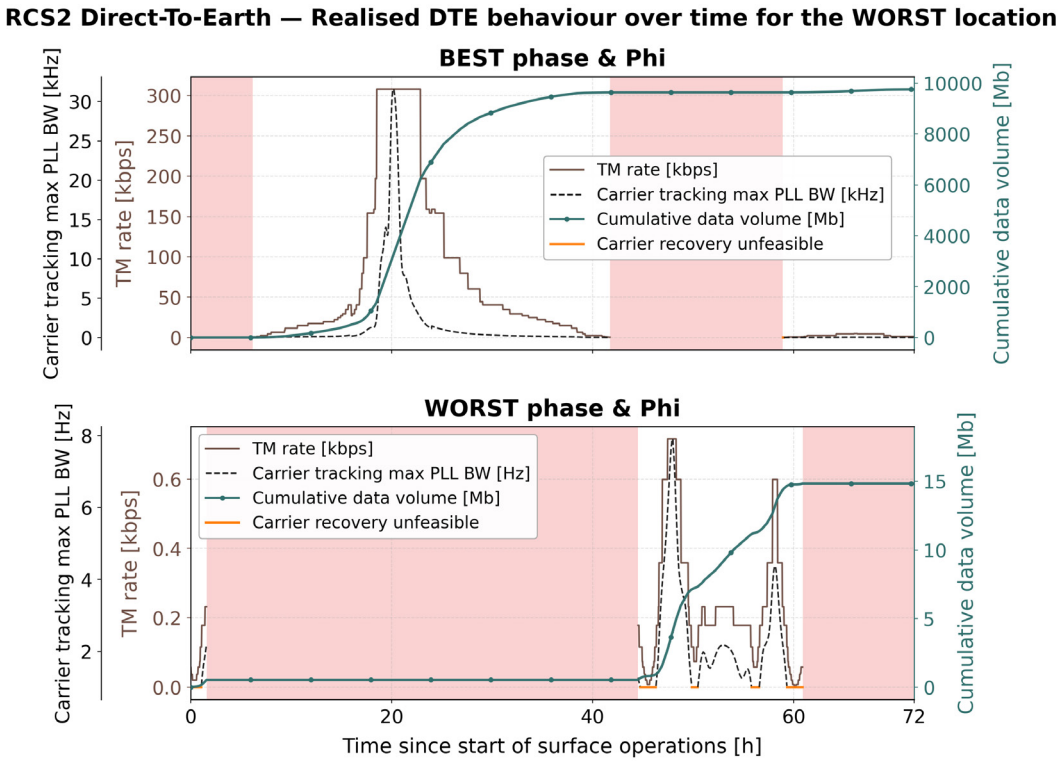


Figure 72: RCS2 DTE realised behaviour over time for the worst landing location

Figure 71 and Figure 72 show how the realised DTE visibility and link budget translate into telemetry rate, cumulative data return and carrier-tracking margin. For each time step, the maximum feasible telemetry rate is selected from the available TM-rate set using the selection criterion defined in Equation (3.28). The selected rate is then integrated using Equations (3.32) and (3.33) to obtain cumulative data return, while the maximum PLL tracking bandwidth provides an indication of residual-carrier tracking capability.

For the best landing location shown in Figure 71, the realised DTE behaviour remains strongly dependent on the initial asteroid phase. In the top panel, the best initial phase realisation places the main visibility interval close to the most favourable part of the Earth–Apophis range evolution. The selected telemetry rate can rise to the highest available mode of 300 kbps, and the cumulative data return reaches approximately 11 Gbit after 60 h. Carrier recovery also remains feasible during the main data-return interval.

The bottom panel shows the worst initial phase realisation for the same landing location. In this case, the main visibility opportunities occur later and under less favourable link conditions. The selected telemetry rate remains much lower, generally around 1 kbps or less, carrier recovery is not always feasible, and the cumulative data return reaches only approximately 25 Mbit. Even for the most favourable landing site, the initial phase therefore has a first-order effect on both the timing and the effectiveness of the DTE return.

Figure 72 shows the corresponding realised DTE behaviour for the worst landing location. The same general pattern is observed: the best phase realisation places the main visibility interval near the most favourable part of the Earth–Apophis range evolution, while the worst realisation shifts the useful visibility later, when link conditions are less favourable and carrier recovery is intermittent. The best realisation still allows substantial data return, with telemetry rates also reaching 300 kbps and cumulative return rising to approximately 10 Gbit. In the worst realisation, the selected telemetry rate remains generally below approximately 1 kbps, carrier recovery is not always feasible, and the cumulative return reaches only approximately 15 Mbit.

The cumulative data-volumes in Figure 71 and Figure 72 show that the two best realisations differ by approximately 20%, while both worst realisations remain limited to only a few tens of Mbit. Across these representative cases, the larger contrast is therefore associated with the realised visibility phase, rather than with the selected best or worst landing location alone.

Figure 73 compares the cumulative data return over time for three landing locations: the best site, the worst site and the representative median site. The site ranking is based on cumulative data-volume at the 40 h time marker, using for each site the median value. The representative median location is defined using the minimum Kolmogorov–Smirnov distance criterion introduced in Section 3.5. Figure 73 also shows the best and worst realisation for the best and worst sites.

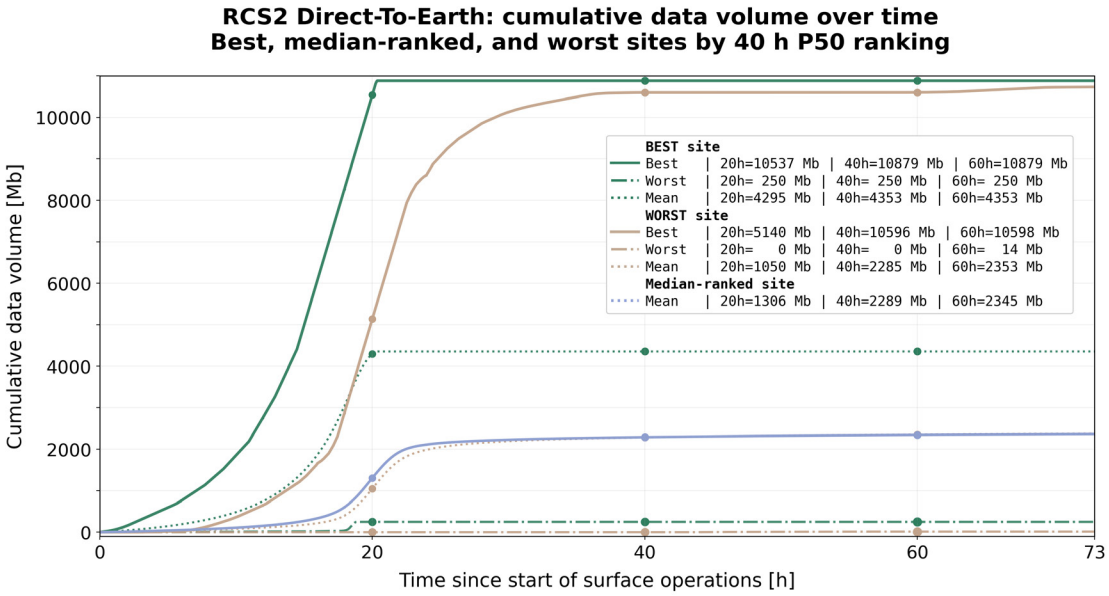


Figure 73: RCS2 DTE cumulative data-volume over time for selected landing sites

Figure 73 shows that the cumulative DTE data return is strongly concentrated in the early part of the RCS2 operational timeline. For all three representative landing sites, most of the final data-volume is accumulated within approximately the first 20–30 h. After this interval, the cumulative curves flatten and only limited additional data-volume is accumulated.
Figure 73 also confirms the behaviour illustrated in Figure 71 and Figure 72: performance can vary strongly even for the same landing site. For both the best and worst sites, favourable realisations can exceed 10 Gbit by 40 h, whereas the worst realisations remain much lower, reaching as little as approximately 14 Mbit for the worst site. This shows that the outcome for each site remains strongly affected by the specific realisation.

The difference between landing locations is clearer in the mean behaviour. For the best site, the mean cumulative return reaches approximately 4.3 Gbit by 20 h, with only limited growth afterwards. For the worst site, the mean return is approximately 1 Gbit at 20 h and increases to approximately 2.3 Gbit at 40 h, again with only limited additional growth later. The representative median site remains only slightly above the worst-site mean, indicating that most of the landing-site population is closer to the lower-performing behaviour than to the best-site case.

Figure 74 shows the cumulative distribution functions of achievable RCS2 DTE data-volume at 20 h, 40 h and 60 h. For each time marker, the best, worst and representative median landing sites are selected from the corresponding cumulative data-volume ranking. This allows the distribution of outcomes to be compared at each mission time using the site ranking relevant to that specific time marker.

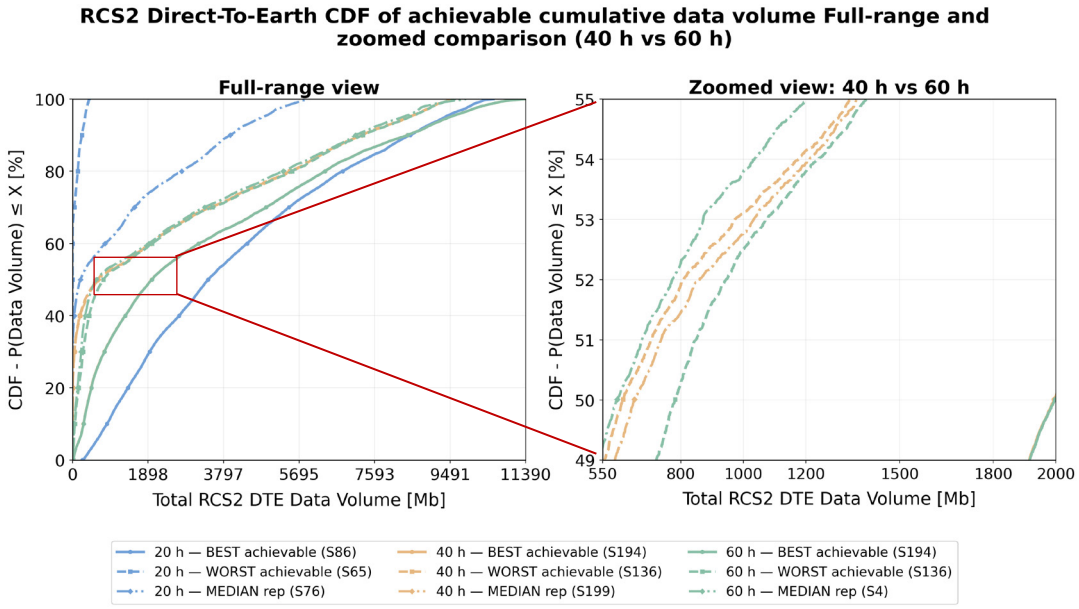


Figure 74: RCS2 DTE CDF of achievable cumulative data-volume at key mission times

Figure 74 shows that the RCS2 DTE cumulative-data distributions remain very broad at all three time markers. Even at 40 h and 60 h, the lower part of the distributions remains close to zero, indicating that a longer surface-operation window does not guarantee significant data-volume restitution. Conversely, some favourable 20 h realisations already achieve data-volumes comparable with the upper part of the later distributions. The time marker alone is therefore not sufficient to characterise DTE performance and again the realised landing location, initial phase and antenna heading remain the main sources of dispersion.

Between 20 h and 40 h, the lower and median parts of the distributions move towards higher data-volumes, showing that some additional realisations benefit from the longer available window. The best-site distribution instead becomes more concentrated: most cases move into a narrower intermediate range, while only the upper tail retains the highest returned data-volumes. However, the distributions remain broad overall, and the poor-performing tail is still present.

Between 40 h and 60 h, the change becomes much smaller. This is particularly visible in the zoomed panel on the right, where the 40 h and 60 h curves nearly overlap over the central part of the distribution. Extending the surface-operation window beyond 40 h therefore provides little additional DTE data return for most realisations.

Figure 75 shows the global percentiles of achievable RCS2 DTE cumulative data-volume at key observation times 20 h, 40 h, 60 h and at further extended 72 h. The percentiles are computed across the full set of landing sites, initial phases and landed antenna headings,

providing a compact summary of the nominal DTE cumulative data-volume return capability over the RCS2 surface operations window.

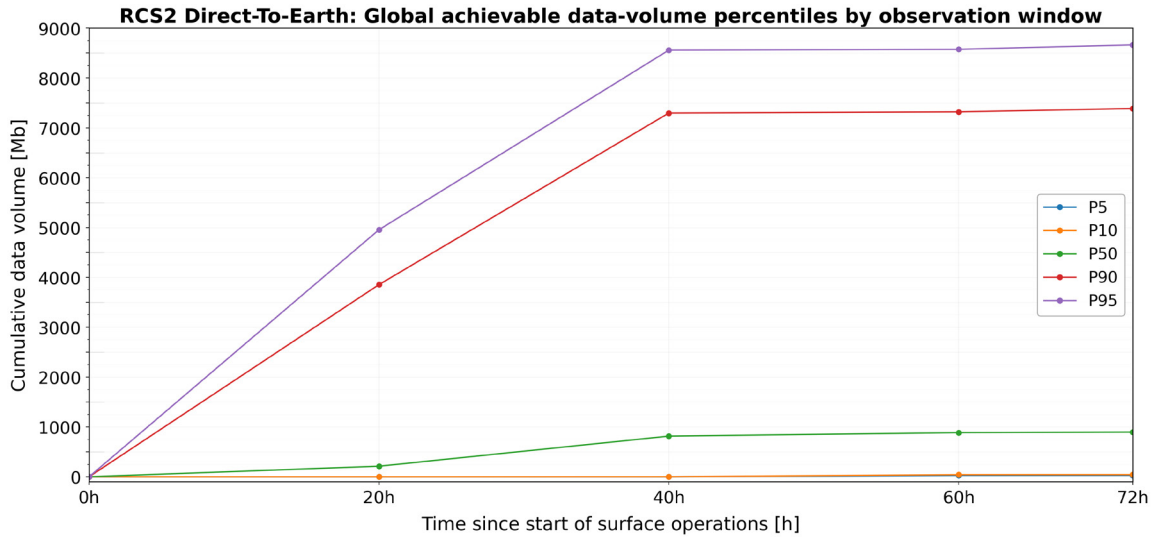


Figure 75: RCS2 DTE global achievable cumulative-data-volume percentiles

Figure 75 shows that the main accumulation of RCS2 DTE data-volume occurs before approximately 40 h, consistent with the behaviour shown in Figure 74. However, the achievable return differs strongly across the percentile range. The upper tail accumulates data rapidly: P90 increases from approximately 3.8 Gbit at 20 h to approximately 7.3 Gbit at 40 h, while P95 increases from approximately 5.0 Gbit to approximately 8.5 Gbit over the same interval. The median case remains much less, rising from approximately 0.2 Gbit at 20 h to approximately 0.8 Gbit at 40 h with no further significant increase later in the timeline.

The lower tail remains strongly constrained throughout the analysed window. P5 stays close to zero, while P10 remains below approximately 0.1 Gbit even at 72 h. This indicates that a significant fraction of the realised cases returns little or no DTE data, despite the extension of the surface-operation window.

After 40 h, all percentiles change only marginally. The P90 and P95 curves remain close to their 40 h values, and the median increases only slightly. Extending the operational window to 60 h or 72 h therefore provides little additional data return, confirming that the nominal RCS2 DTE data-volume is largely determined during the first 40 h of surface operations, around the Earth–Apophis closest-approach period.

Figure 75 show the wide spread of achievable RCS2 DTE return across the complete population of landing sites and landed conditions.

However, those global statistics do not identify which landing sites provide robust performance. It is therefore useful to rank the landing-site population directly, so that

favourable locations can be distinguished from sites that perform well only under specific realised conditions or time window.

The ranking uses the low-tail P10 cumulative data-volume at the three main mission time markers of 20 h, 40 h and 60 h. P10 is used to provide a conservative measure of site performance under relatively adverse, but not extreme, realised conditions. For each landing site, this value is computed across all combinations of initial rotational phase and landed antenna heading.

At each time marker, the sites are ranked according to their P10 cumulative data-volume and expressed as a normalised goodness rank, where 0 represents the lowest-ranked site and 1 the highest-ranked site. The overall landing-site goodness metric is then taken as the minimum of the three normalised ranks. This makes the metric conservative across mission time: a site receives a high score only if it remains favourable at 20 h, 40 h and 60 h.

Figure 76 presents the resulting landing-site ranking in four vertical panels. The first three panels show the site-ranked cumulative-data-volume percentiles at 20 h, 40 h and 60 h, respectively. These panels retain the percentile spread for each landing site, showing how the realised rotational phase and antenna heading affect the achievable return. The bottom panel shows the normalised site-goodness metric. In all panels, the landing sites are ordered according to the overall goodness ranking.

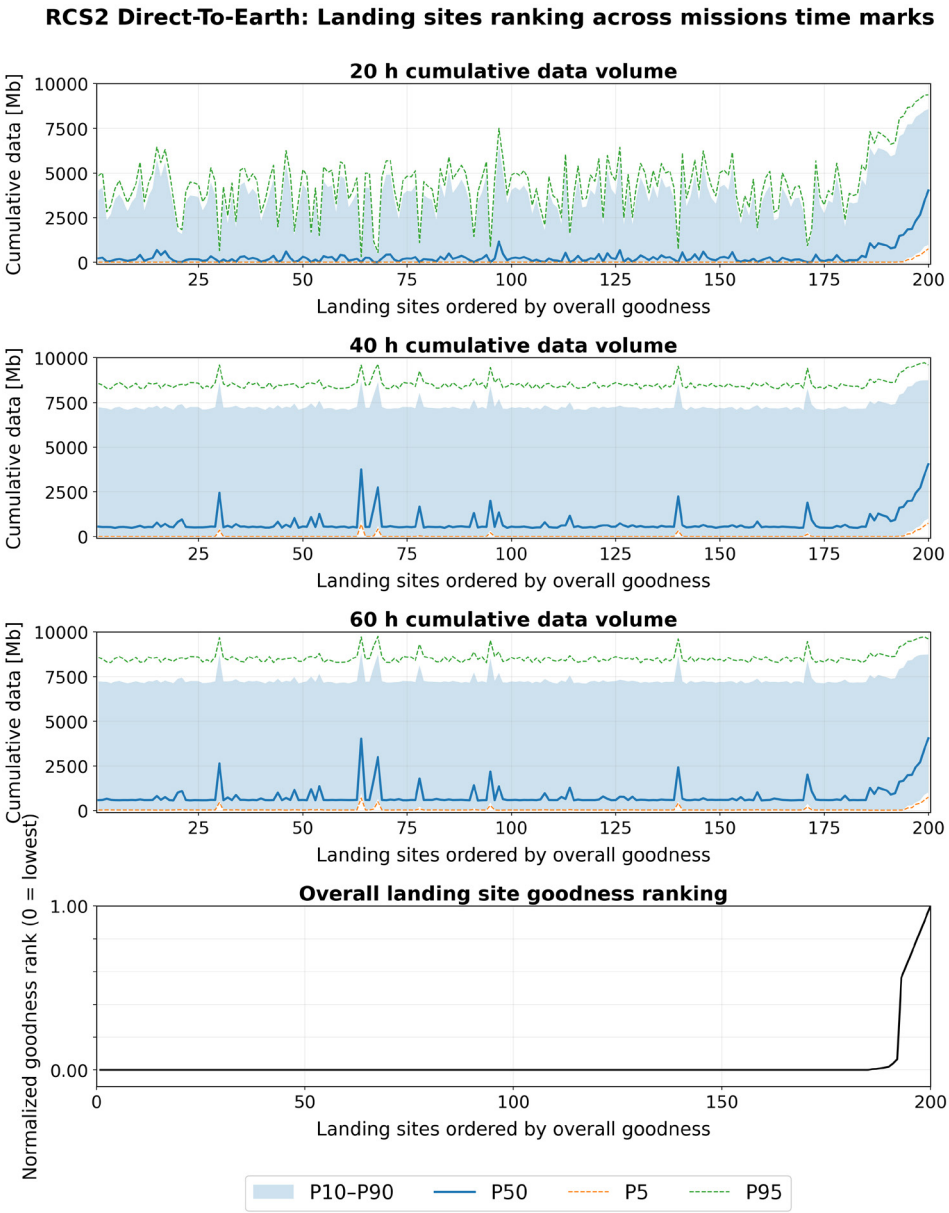


Figure 76: RCS2 DTE landing-site ranking across mission time marks

Figure 76 shows that the nominal RCS2 DTE landing-site performance is highly uneven across the ranked population. At all three time markers, a large fraction of landing sites remains close to zero in the lower percentiles, while high cumulative data-volumes are confined mainly to favourable realised conditions. This indicates that strong DTE performance is possible, but only for a limited subset of site–phase–heading combinations.

The 20 h panel already shows this separation, with only the highest-ranked sites beginning to develop non-zero median and lower-percentile return. By 40 h, the favourable upper percentiles have increased substantially, and the P95 curve reaches high values across a much larger part of the population. However, the median and lower percentiles remain low for most sites, showing that the improvement mainly affects favourable realisations rather than the full

landed population. The 60 h panel is similar to the 40 h case, indicating that the main DTE data return opportunity is largely established by approximately 40 h.

The bottom panel synthesises this behaviour through the conservative site-goodness metric. The normalised goodness remains close to zero for most landing locations and increases sharply only in the high-performing tail. This indicates that favourable RCS2 DTE performance exists, but is consistently retained only by a small subset of landing sites across the analysed phase and heading variations at 20 h, 40 h and 60 h.

Figure 77 and Figure 78 show the same landing-site performance quantities introduced in Figure 76, mapped onto the asteroid surface in 2D and 3D form. The 2D maps provide a clearer overview of the full surface distribution, while the 3D maps retain the asteroid geometry and spatial context more directly, although some regions are less visible because of the viewing projection.

Both Figure 77 and Figure 78 contain four panels: the first three show the conservative P10 cumulative data-volume respectively at 20 h, 40 h and 60 h, while the fourth shows the normalised overall site-goodness ranking. Their purpose is to indicate where favourable and unfavourable DTE landing sites are located on the asteroid.

A colour scale is used to represent either the P10 cumulative data-volume or the ranking value, depending on the panel. The best and worst sites for the quantity shown in each panel are highlighted with dedicated markers.

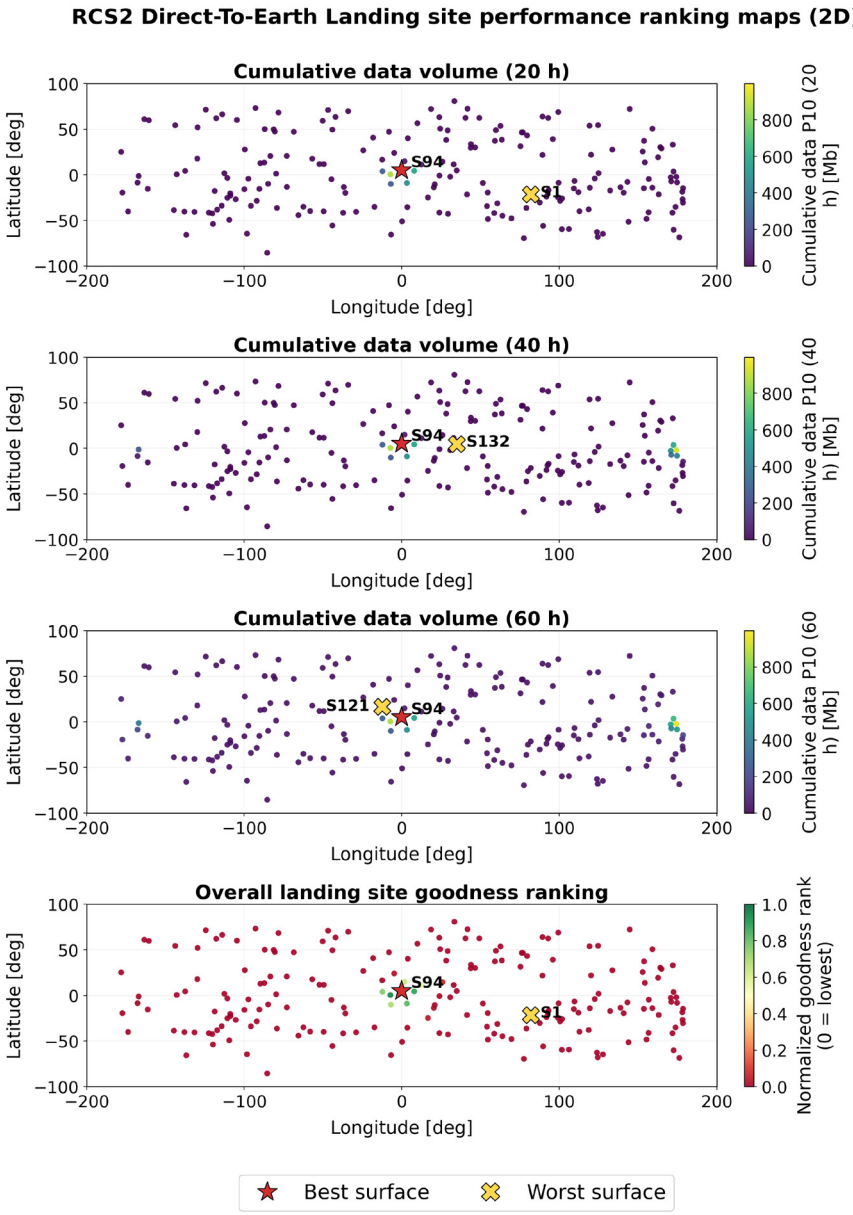


Figure 77: RCS2 DTE spatial distribution of landing-site performance and ranking (2D)

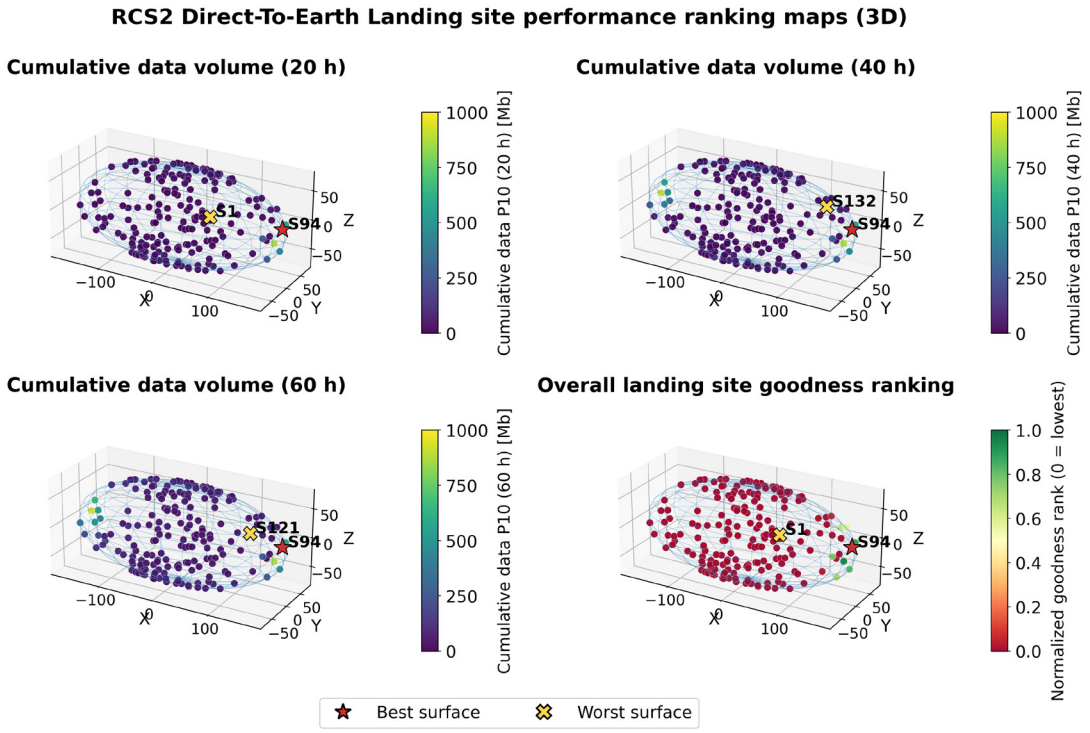


Figure 78: RCS2 DTE spatial distribution of landing-site performance and ranking (3D)

Figure 77 and Figure 78 show that favourable landing-site performance is not uniformly distributed over the asteroid surface, but is instead spatially concentrated. At the 20 h time marker, the best-performing landing locations are clustered within one relatively limited region, while most of the remaining surface is associated with generally lower performance. At the 40 h and 60 h time markers, two small clusters of higher-performing sites become visible.

The overall goodness ranking remains concentrated in a single localised region near one edge of the ellipsoid, while most of the remaining surface is characterised by poorly ranked sites. While the best-site marker is consistently located in the same location across the panels, the worst-site marker changes between the individual time markers and the overall goodness map. This indicates that poor DTE performance is widespread across the surface.

5.2.3 Inter-Satellite-Link Communication Performance

The RCS2 ISL case is evaluated using the link budget and data-volume formulation defined in Section 3.4. The received power is computed from the ISL link budget of Equation (3.34) and compared with the data-mode feasibility thresholds of Equation (3.36). The highest feasible data mode is then selected at each time step and integrated over the visible intervals to obtain the instantaneous and cumulative data-volume, following Equations (3.37) and (3.38).

Figure 79 shows the evolution of the main ISL link budget contributors and the resulting received power for an illustrative nominal realised case over an extended period of 72 h of RCS2 surface operations.

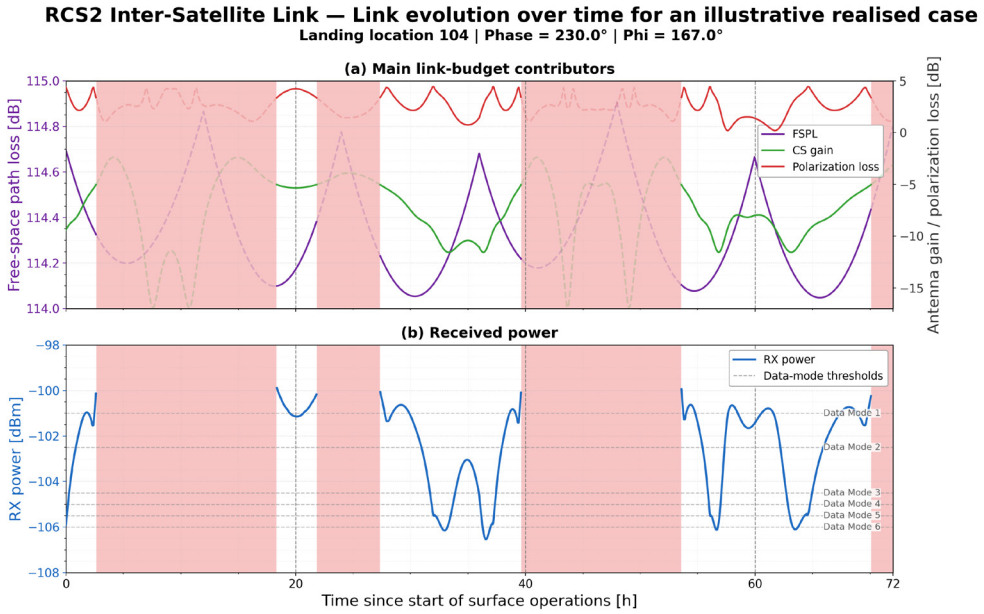


Figure 79: RCS2 ISL link budget contributors evolution

Figure 79 has two vertical panels: the upper one shows the main received-power contributors, including FSPL, RCS2 antenna gain and polarisation loss. The lower panel shows the resulting received power. Visibility interruptions due to the link geometry are indicated by shaded regions.

The FSPL remains nearly constant over the analysed interval, with a value of approximately 114.5 dB and a variation of less than 1 dB. This reflects the local nature of the RCS2–RAMSES link, with lower and less varying propagation losses than in the DTE case.

The received-power variation is driven by local visibility and by the realised antenna geometry during the visible intervals. During periods of visibility, the received power ranges from approximately −100 dBm to approximately −107 dBm. This variation is caused mainly by changes in antenna gain and polarisation loss. The RCS2 antenna gain decreases from slightly negative values to roughly −12 dBi in the less favourable parts of the visible intervals, while the polarisation loss ranges from approximately 0 to 5 dB.

Figure 80 links the received-power behaviour of Figure 79 to the resulting ISL telemetry performance. The upper panel of Figure 80 shows the received-power evolution against the

ISL data mode thresholds, while the lower panel shows the corresponding selected data mode and cumulative data-volume returned.

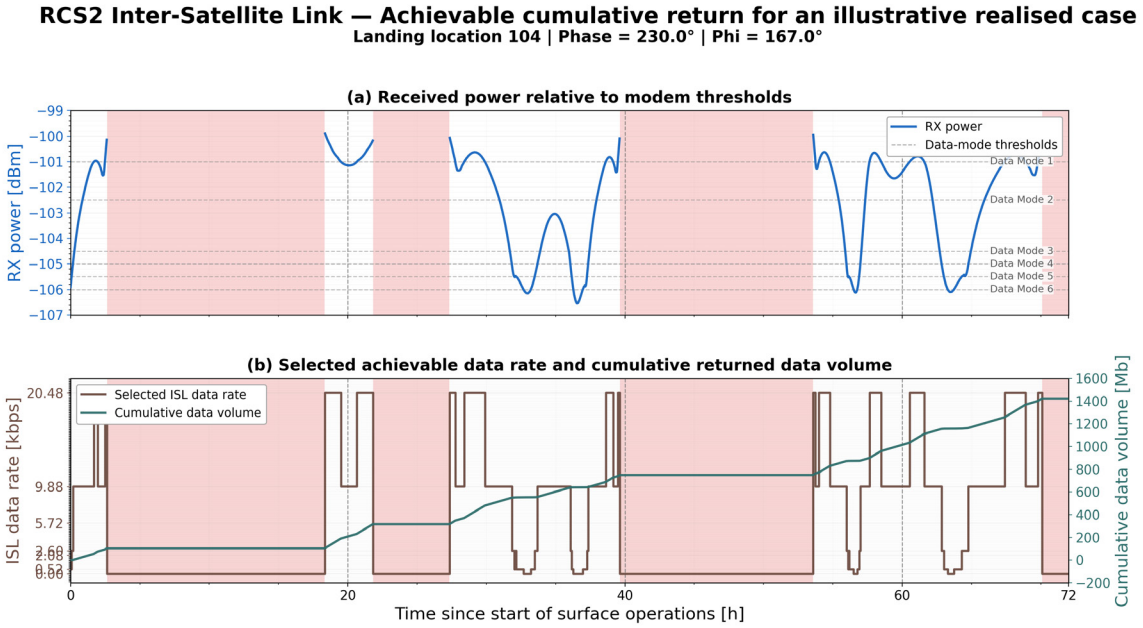


Figure 80: RCS2 ISL illustrative achievable data mode and cumulative return

Figure 80 shows that the nominal RCS2 ISL data return is accumulated during repeated visibility intervals rather than through continuous communication. Within each interval, the received power varies with the local link geometry and may cross several modem thresholds, leading to changes in the selected data mode.

This produces a cumulative data-volume profile with a stepped behaviour, where increases occur during visibility intervals and flat sections correspond to periods without line of sight. The achievable return in this representative case is primarily driven by the duration and quality of the communication windows.

Figure 81 and Figure 82 show the variability behind the RCS2 ISL performance presenting two representative realisations from the favourable and adverse parts of the population respectively. These examples are not intended to define absolute best- and worst-case, instead, they show how distinct sequences of geometrical visibility produce different cumulative data-volume profiles.

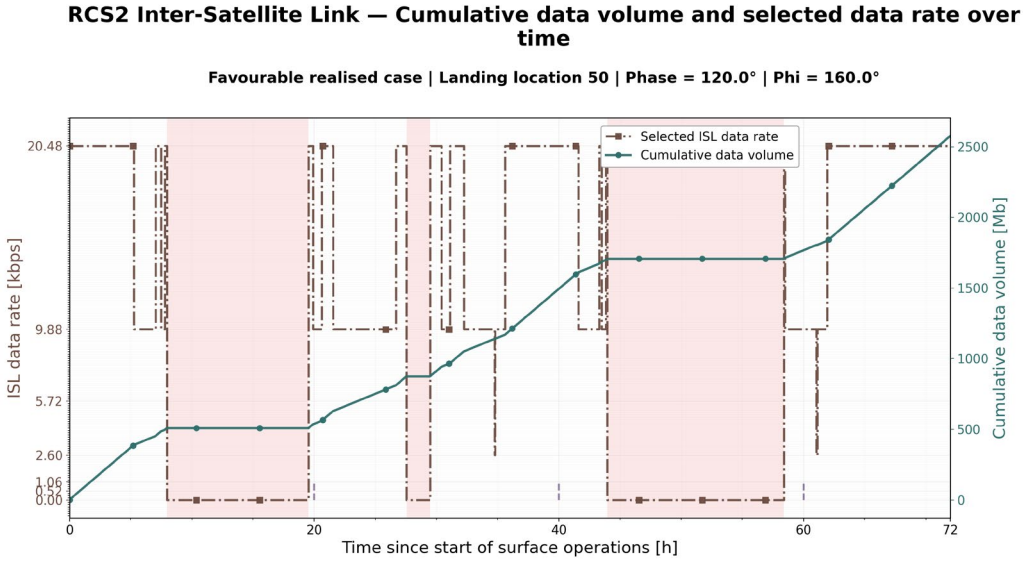


Figure 81: RCS2 ISL selected data rate and cumulative volume for a favourable case

Figure 81 shows a favourable ISL case. In this realisation, a large fraction of the surface operation window is available for communication, and much of the available contact time supports the two highest ISL data modes. The cumulative data-volume therefore increases rapidly during the usable contact periods, reaching approximately 1.8 Gbit at the 60 h marker and further increasing to approximately 2.5 Gbit by the end of the extended 72 h interval.

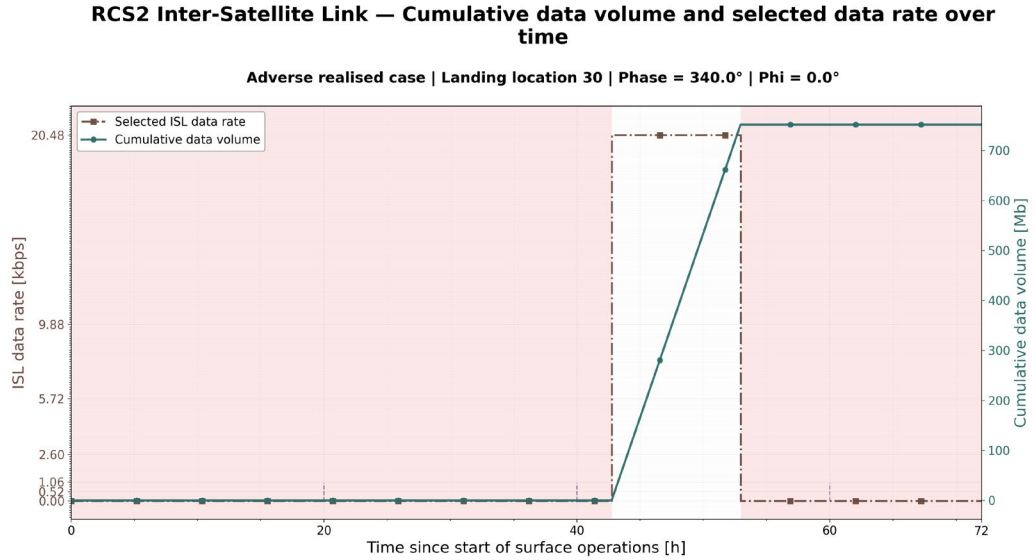


Figure 82: RCS2 ISL selected data rate and cumulative volume for adverse case

Figure 82 shows an adverse ISL case. In this realisation, the lower cumulative return is driven mainly by the reduced availability of line-of-sight visibility time. When visibility is available, the received power allows constant selection of the highest ISL data mode. However, only one short usable interval is present in this adverse case. The cumulative data-volume

therefore reaches only approximately 750 Mb at the 60 h marker, despite increasing at the maximum rate. No further data-volume is accumulated afterwards, as no additional visibility occurs. The adverse outcome is therefore mainly associated with limited contact availability, rather than with weak instantaneous ISL link performance.

Figure 83 shows the nominal cumulative ISL data return over the surface operation window for the absolute best, median representative and absolute worst performance cases in the analysed landing-site population. The median representative case is selected using the Kolmogorov–Smirnov distance criterion introduced in Section 3.5. For the best and worst landing sites, the figure also shows the best, mean and worst realisations, allowing both the site-to-site spread and the within-site dispersion to be compared.

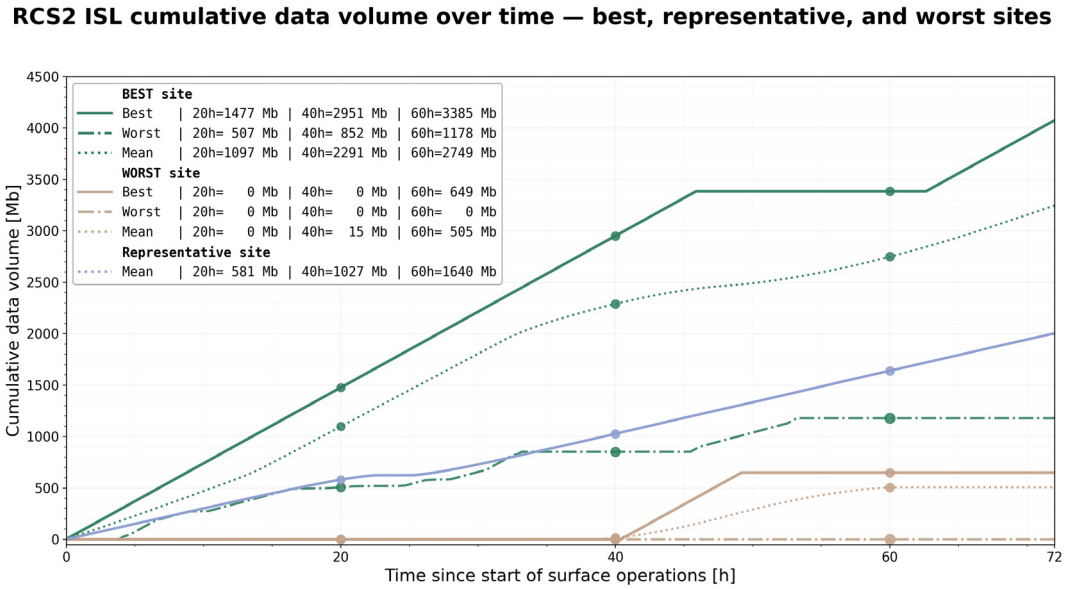


Figure 83: RCS2 ISL cumulative data-volume over time for selected landing sites

Figure 83 shows a clear separation between the best, median representative and worst ISL performance cases. The best site accumulates data early and reaches approximately 2.7 Gbit in the mean case at 60 h, with an absolute best realisation of approximately 3.4 Gbit and an absolute worst realisation still above 1.1 Gbit. This indicates a relatively wide within-site dispersion, but also shows that even the least favourable realisation for the best site remains well separated from the worst-site performance. The worst site accumulates data much later and remains close to zero until around 40 h; even its absolute best realisation reaches only about 0.65 Gbit at 60 h, while its worst realisation remains at zero.

The median representative case lies between the two limiting behaviours, reaching approximately 1.6 Gbit at 60 h. The figure therefore shows that the nominal ISL population

has a clear site-to-site spread, with the median representative case positioned between the favourable and adverse limits.

Figure 84 shows the CDF of cumulative ISL data-volume for the same best, median-representative and worst landing-site cases used in Figure 83, at 20 h, 40 h and 60 h time markers.

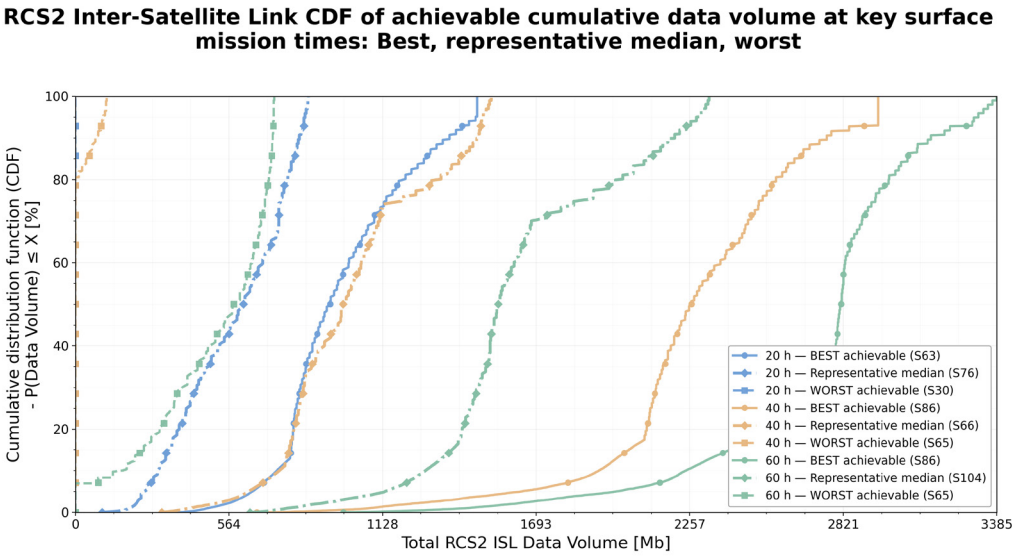


Figure 84: RCS2 ISL CDF of achievable cumulative data-volume at key mission times

The separation between the curves in Figure 84 shows that landing-site geometry has a persistent effect on ISL performance across the surface-operation window.

The shape of the curves shows substantial differences in the spread within each selected site. The worst-site case remains concentrated close to zero at 20 h and 40 h for most realisations, but shifts to a finite, although still limited, data-volume by 60 h. The representative median case moves progressively to higher data-volumes between 20 h, 40 h and 60 h, while the best-site case reaches the highest values and retains a broader spread, especially at 60 h.

Figure 84 shows that RCS2 ISL performance is not defined only by the selected landing site, but also by the distribution of realised initial phase and antenna azimuth conditions within each site. The selected sites remain clearly ordered, but each site still contains internal dispersion that affects the achievable cumulative data-volume.

Figure 85 shows the global nominal ISL performance across the full landing-site population using selected cumulative-data-volume percentiles at the main observation times.

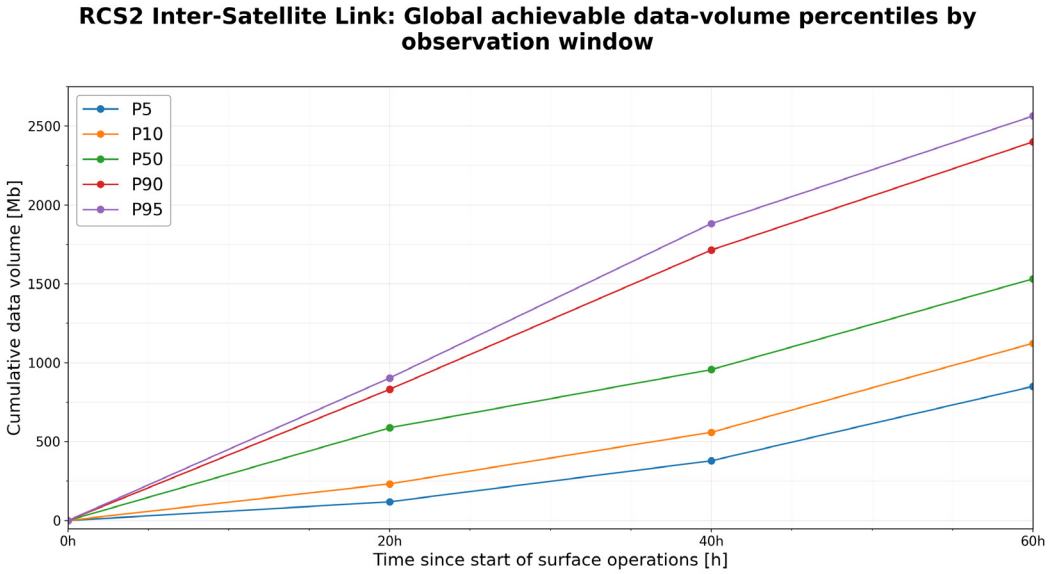


Figure 85: RCS2 ISL global achievable cumulative-data-volume percentiles

Figure 85 shows that the global ISL cumulative data-volume percentiles increase across the 20 h, 40 h and 60 h observation windows, indicating continued data accumulation throughout the analysed surface-operation period. The upper percentiles show stronger growth between 20 h and 40 h than between 40 h and 60 h, while the lower percentiles continue to increase more noticeably in the later interval. P95 increases by approximately 980 Mb between 20 h and 40 h, but by about 680 Mb between 40 h and 60 h. By contrast, P10 increases by approximately 320 Mb over the first interval and by approximately 560 Mb over the second. This indicates that all parts of the landed population continue to accumulate ISL data-volume over the analysed window, but with different time profiles. The upper percentiles already achieve a large part of their return by 40 h and continue to increase afterwards at a reduced rate. The lower percentiles also continue to grow, particularly between 40 h and 60 h, but remain well separated from the upper tail. The median remains between the lower and upper percentiles throughout, showing that the central population behaviour is broadly consistent with the overall spread.

Figure 86 shows the ISL landing-site ranking across the full analysed population. Following the same ranking method introduced for the DTE case in Figure 76, the first three panels show the cumulative-data-volume percentiles at 20 h, 40 h and 60 h, while the fourth panel shows the normalised overall site-goodness metric used to order the landing sites.

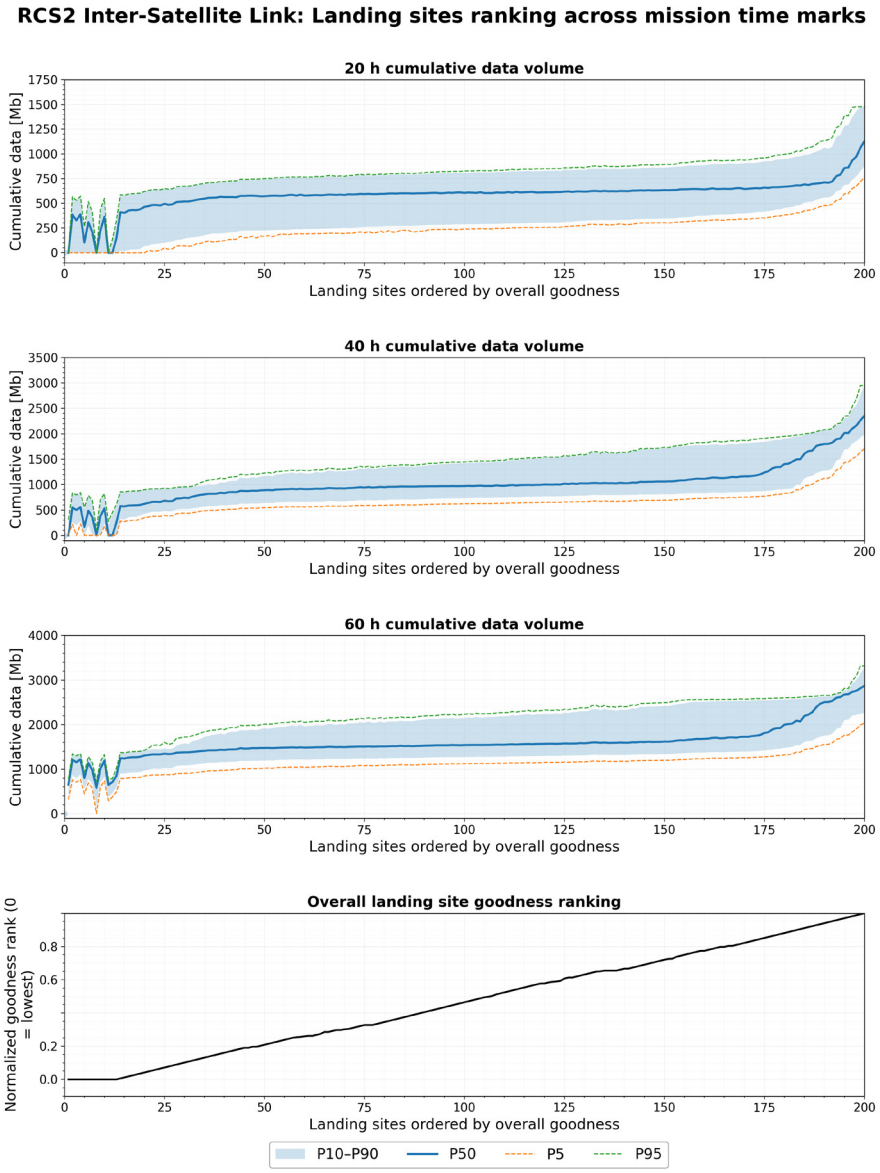


Figure 86: RCS2 ISL landing-site ranking across mission time marks

Figure 86 shows that ISL landing-site performance is spread across the ranked population, rather than being concentrated in a few isolated favourable cases. In the 20 h panel, most landing sites already form a broad central group with similar, non-zero cumulative return even at P5 level, with lower-performing and higher-performing tails. The same pattern is preserved at 40 h and 60 h, with all percentiles shifting upward as more data are accumulated over the longer surface-operation windows. For most landing sites, the percentile curves remain grouped around a non-zero return. Phase and antenna-heading conditions still affect the achievable return, but across much of the ranked population the lower percentiles remain close to the median rather than falling towards zero.

This behaviour is reflected in the overall goodness panel. The normalised ranking increases progressively across most of the population, with only approximately 10% of landing sites remaining at zero goodness.

Following the same spatial mapping adopted for the DTE case, Figure 87 and Figure 88 map the ranked ISL landing-site performance from Figure 86 onto the asteroid surface in 2D and 3D form. Consistently with the DTE case, the first three panels show the cumulative data-volume at 20 h, 40 h and 60 h, while the final panel shows the normalised overall site-goodness metric. This spatial representation shows how the ranked landing sites are distributed across the asteroid surface.

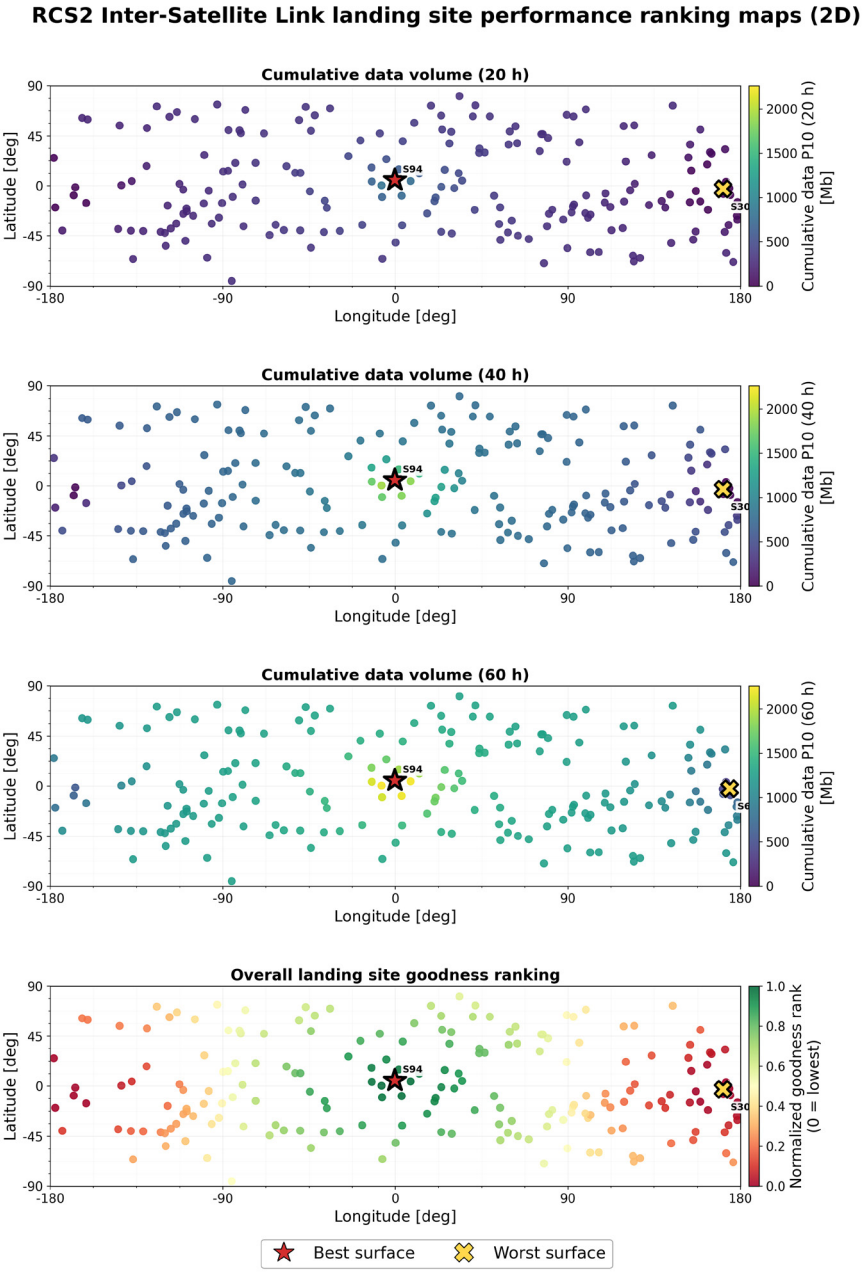


Figure 87: RCS2 ISL landing site performance ranking maps (2D)

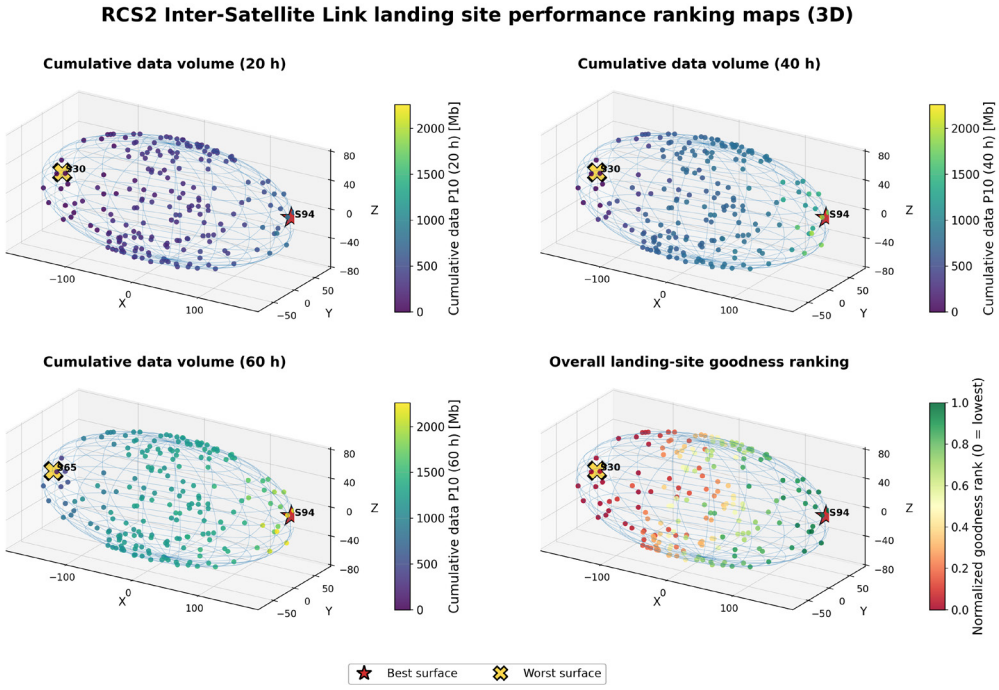


Figure 88: RCS2 ISL landing site performance ranking maps (3D)

Figure 87 and Figure 88 show that the nominal ISL landing-site performance is spatially structured rather than randomly scattered across the candidate surface set. At 20 h, 40 h and 60 h, as well as in the overall goodness ranking, sites with similar performance tend to cluster in broad regions, with a progressive transition from more favourable to more adverse areas across the asteroid. The same best and worst landing sites are identified across the time-marker panels and the overall ranking, indicating that the main spatial performance pattern is persistent over the analysed surface operation windows. The best and worst sites lie on broadly opposite sides of the body, indicating an almost edge-to-edge performance gradient rather than a pattern dominated by a small number of isolated cases or local clusters.

This pattern is visible in both the 2D and 3D representations. The 2D maps make the regional grouping and gradual transition easier to identify in longitude–latitude space, while the 3D maps show that the same structure is distributed across the body surface and is not an artefact of the map projection.

Figure 87 and Figure 88 show that nominal ISL performance depends on broad landing-location geometry, producing coherent favourable and adverse surface regions rather than a scattered distribution of isolated sites.

5.2.4 Consolidated nominal RCS2 communication behaviour

The nominal DTE and ISL results presented in Sections 5.2.2 and 5.2.3 are consolidated in Table 5.1 using common performance indicators. The table summarises the main nominal-performance characteristics of the two candidate architectures.

Table 5.1 - Consolidated nominal RCS2 communication behaviour

Metric / status	DTE nominal behaviour	ISL nominal behaviour	Remarks
Dominant timescale of useful return	Early-window dominated; main return by ~20–30 h.	Data-volume return continues to accumulate linearly across the entire timeline.	Captures the temporal distribution of achievable return.
Main physical/geometric driver	Driven by Earth–Apophis range evolution, local Earth visibility, and landed antenna orientation.	Driven by local RCS2–RAMSES visibility, antenna geometry, and received power relative to modem thresholds.	Identifies the main mechanism governing nominal performance.
Favourable realised capability	Favourable realised cases reach ~10–11 Gbit, with peak selected rates of order 300 kbps.	Favourable realised cases reach Gbit magnitude return, with the best selected case reaching ~3.4 Gbit by 60 h.	Captures the upper realised nominal capability.
Expected cumulative return	Median return is ~0.17 / 0.53 / 0.59 Gbit at 20 / 40 / 60 h, with limited growth after 40 h.	Median return is ~0.5 / 1.0 / 1.5 Gbit at 20 / 40 / 60 h, with progressive growth over the entire time window.	Captures the central population outcome.
Population dispersion and lower-tail status	Very high dispersion: adverse realised cases fall to ~15–25 Mbit and P5/P10 remain very low.	Moderate dispersion; lower and upper percentiles separated from the median typically by a factor 2, with the lower percentiles generally well above zero.	Captures the spread between favourable and adverse realised outcomes.
Post-40 h incremental return	Limited; 40 h and 60 h distributions are close for most realisations.	Non-negligible; global percentiles continue increasing from 40 h to 60 h.	Indicates whether extending surface operation can benefits the result.
Landing-site selectivity / spatial structure	High selectivity: favourable locations are concentrated in two limited surface regions.	Structured regional dependence: favourable and adverse regions form broader spatial patterns across the asteroid surface.	Captures the spatial dependence of nominal performance.

These nominal results provide the baseline for the sensitivity and robustness analysis in the following section and as input to Chapter 6.

5.3 RCS2 Sensitivity analysis

Section 5.2 showed that the nominal RCS2 results are strongly affected by landing site, initial rotational phase and landed antenna heading. This section examines whether those conclusions remain valid when the landed geometry and communication assumptions are modified through selected constraints and degraded conditions.

5.3.1 Direct-To-Earth sensitivity and robustness

The DTE sensitivity analysis considers four effects. The first is the degree of control over the realised landing conditions. It is assessed by varying whether landing timing, landed heading, or both are controlled, to estimate how much DTE performance can improve when favourable initial conditions are targeted rather than left uncontrolled. The second effect is the local horizon constraint, assessed by varying the assumed minimum elevation condition. The third is non-ideal surface geometry, represented through tilt and slope perturbations relative to the nominal ideal surface. The fourth, consistent with the RCS1 case, is telemetry-planning operability, assessed by limiting how often the selected telemetry rate can be updated during the short RCS2 operational window.

Case 1: DTE sensitivity to landing timing and heading control authority

In the nominal RCS2 analysis, the initial rotational phase of Apophis and the landed spacecraft heading are treated as uncontrolled variables contributing to the statistical distribution of achievable DTE performance, as described in Section 5.2. The RCS2 mission may, however, possess some degree of authority over the landing timing and/or the final landed heading, thereby improving the alignment between local visibility, antenna orientation, and the short initial interval in which the link is most favourable. This sensitivity case therefore evaluates the achievable cumulative data return for four levels of control authority over the initial conditions: no control, heading-only control, initial-phase-only control and full control. The objective is to quantify the potential performance benefit of such control authorities, while its practical implementation is outside the scope of the analysis.

Figure 89 shows the effect of landing-control authority on RCS2 DTE cumulative data return. The first three panels show the achievable data-volume across the landing-site population at 20 h, 40 h and 60 h respectively, for the four control-authority cases. In these panels, landing sites are ordered by their cumulative data-volume in the no-control case. The same ordering is then retained for the three control cases. This allows the improvement introduced by each level of control authority to be compared against the same uncontrolled baseline. The fourth

panel provides a compact comparison of the same cases, plotting the median data-volume at 20 h against the median data-volume at 40 h, with the vertical whisker indicating the corresponding 60 h value. In this panel, cases farther towards the upper-right indicate higher median data return at both 20 h and 40 h, and therefore better overall performance, while the vertical whiskers show the additional growth achieved by 60 h.

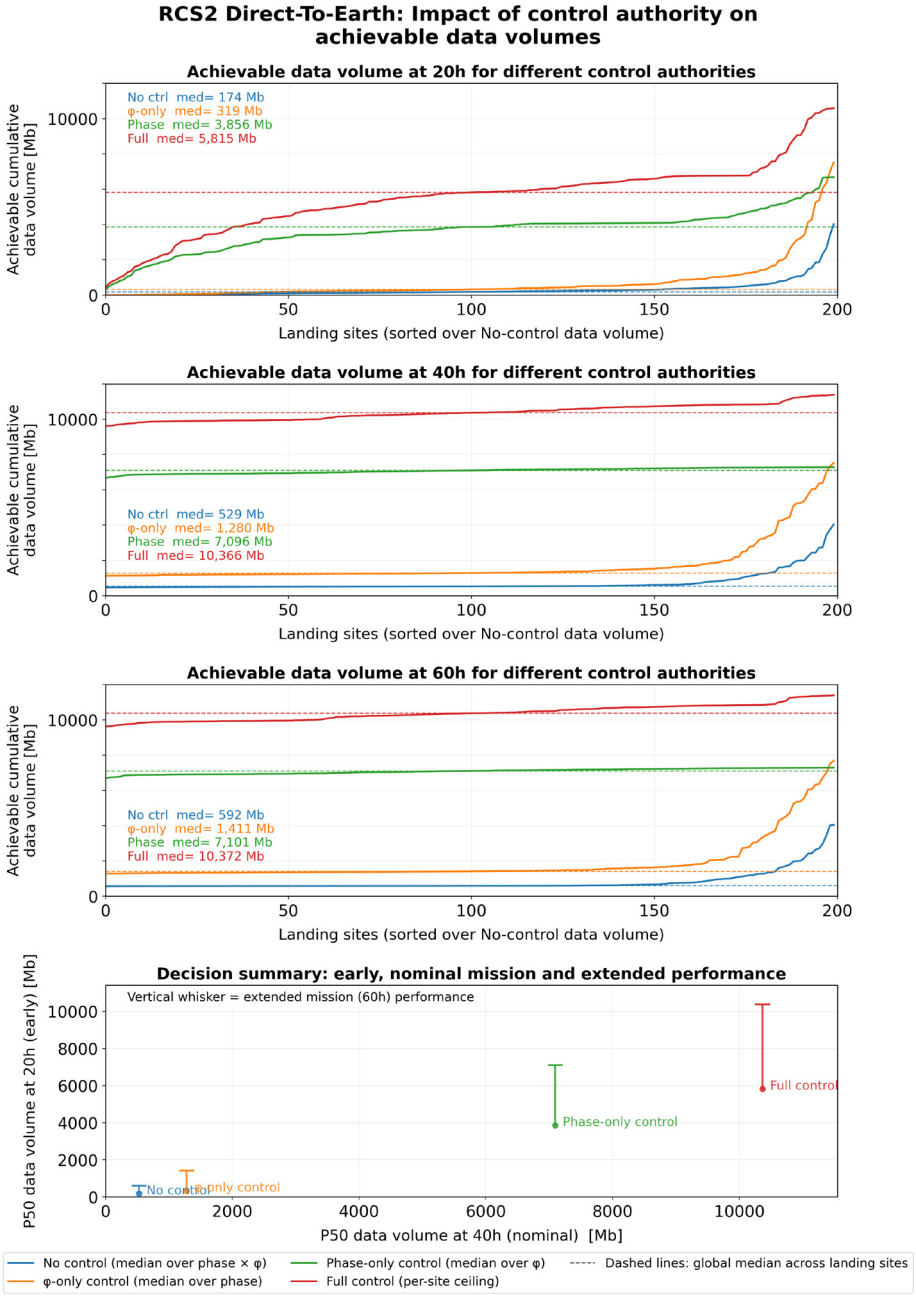


Figure 89: RCS2 DTE sensitivity to landing control authority

Figure 89 shows that the achievable DTE performance is highly sensitive to control authority over the landed initial conditions. Heading-only control provides only limited improvement over the uncontrolled case, with the benefit mainly visible in the upper tail of the population.

Control over the initial phase produces a much larger increase in achievable return, while full control gives the highest performance and makes the for 40 h and 60 h the outcomes nearly equal across the landing site population.

The median cumulative data-volumes and relative improvement against the uncontrolled case are summarised in Table 5.2.

Table 5.2 - RCS2 DTE performance under different landing-control authorities

Control authority	20 h [Mb]	40 h [Mb]	60 h [Mb]	Gain relative to no control at 20 h	Gain relative to no control at 40 h	Gain relative to no control at 60 h
No control	174	529	592	1.0x	1.0x	1.0x
Heading-only control	319	1,280	1,410	1.8x	2.4x	2.4x
Phase-only control	3,860	7,100	7,100	22.2x	13.4x	12.0x
Full control	5,820	10,370	10,370	33.4x	19.6x	17.5x

Table 5.2 shows that the performance gain is highly uneven across the different control cases. Heading-only control provides only a limited improvement over the uncontrolled case, whereas control over the initial phase produces a step change in performance and already captures most of the achievable benefit. This indicates that the key factor is the timing alignment between the landed configuration and the short early interval of favourable Earth visibility and reduced Earth–Apophis range. Once this alignment is achieved, additional heading optimisation further improves the result, but only as a secondary contribution.

In the bottom panel of Figure 89, the higher-authority cases lie clearly further towards the favourable upper-right region, while the 60 h whisker endpoints remain comparatively close to the corresponding 40 h values. This indicates that most of the DTE return has already been secured by approximately 40 h, with only limited additional gain by 60 h.

Case 2: sensitivity to the local horizon assumption

In the nominal RCS2 DTE analysis, geometric visibility is evaluated using a baseline minimum-elevation mask of 10°. As introduced in Section 3.3, this provides a simplified representation of local horizon masking and near-horizon degradation, imposed through the zenith-angle visibility condition of Equation (3.14) and illustrated conceptually in Figure 26. Since the baseline value represents an analysis assumption rather than a detailed terrain-derived horizon model, its influence on the DTE results is assessed through sensitivity analysis. This sensitivity case first quantifies the impairment introduced by the nominal 10° mask relative to an ideal free-local-horizon. It then examines how the achievable DTE

performance is influenced by modestly more favourable and more adverse minimum-elevation assumptions.

Figure 90 shows the impairment introduced by the nominal minimum elevation $\varepsilon_{\min} = 10^{\circ}$ mask in the no-control DTE case. The figure reports the percentage reduction in achievable cumulative data-volume relative to an ideal free-local-horizon case with a minimum elevation $\varepsilon_{\min} = 0^{\circ}$, using the P5, P50 and P95 statistics of the global no-control population at 20 h, 40 h and 60 h.

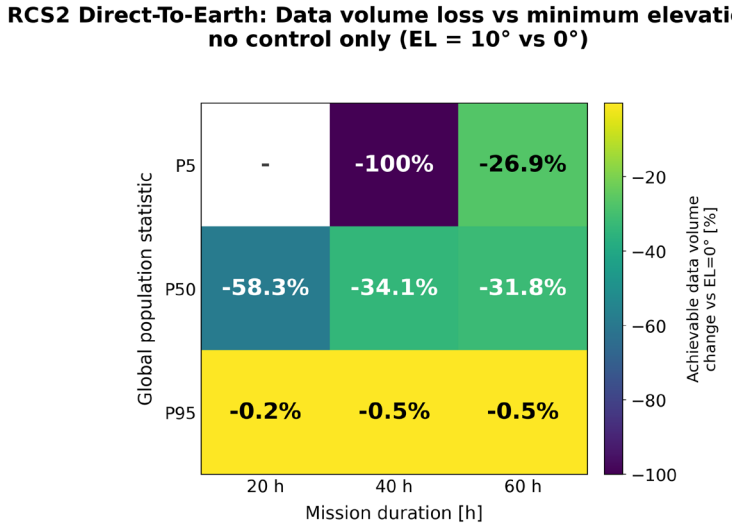


Figure 90: RCS2 DTE impairment due to the nominal minimum-elevation mask

Figure 90 shows that the nominal minimum-elevation assumption has a measurable impact, but that this impact is not uniform across the realised population. The lower tail is the most sensitive part of the distribution. At P5, a dash is shown at 20 h because the $\varepsilon_{\min} = 0^{\circ}$ reference percentile is already zero, making the relative change undefined. At 40 h, the P5 value reaches -100%, indicating a complete loss of data return relative to the 0° reference case: the selected percentile is greater than zero with the 0° mask, but falls to zero when the 10° mask is applied.

The median return is also reduced substantially by the baseline 10° mask, with losses of approximately 58% at 20 h and approximately 32–34% at 40–60 h. The upper tail is much less affected, with P95 losses close to zero at all three time markers. The adopted assumption $\varepsilon_{\min} = 10^{\circ}$ therefore mainly affects median and lower-tail outcomes, while the most favourable cases remain comparatively insensitive.

Having shown that the nominal minimum-elevation mask introduces a non-negligible impairment, Figure 91 examines the sensitivity of the no-control DTE results to modestly

more favourable and more adverse assumptions around the adopted baseline. The figure shows the percentage variation in achievable cumulative data-volume for $\varepsilon_{\min} = 7^\circ$ and $\varepsilon_{\min} = 13^\circ$ relative to the nominal $\varepsilon_{\min} = 10^\circ$ case. The results are again summarised using the P5, P50 and P95 statistics of the global no-control population at 20 h, 40 h and 60 h.

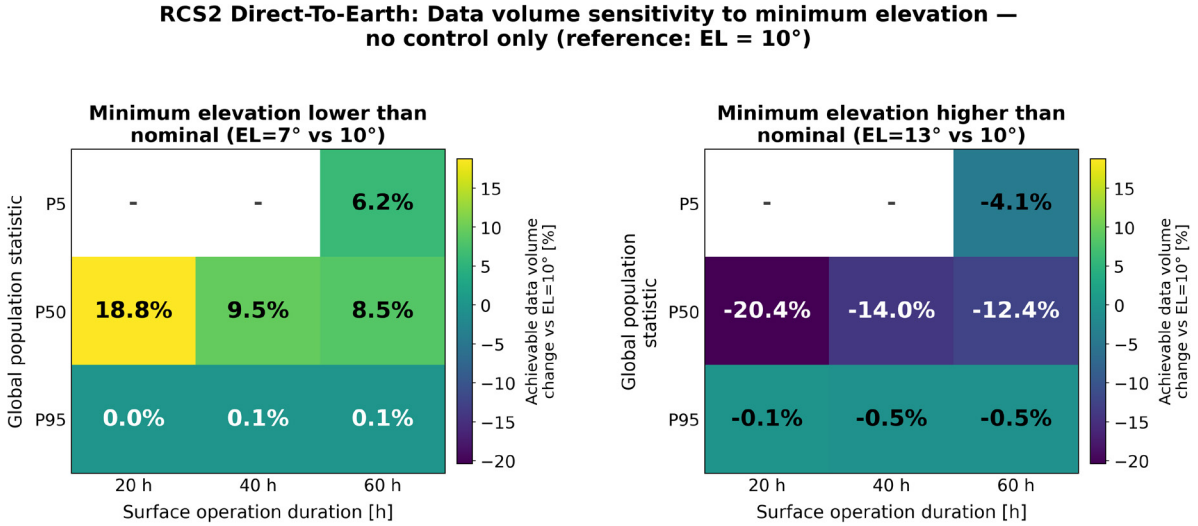


Figure 91: RCS2 DTE sensitivity to the minimum-elevation assumption

Figure 91 shows that moderate variations around the nominal minimum-elevation assumption affect the no-control DTE case mainly at the median level. Reducing the mask to 7° increases the median cumulative return by approximately 19% at 20 h and by approximately 9–10% at 40–60 h. Increasing the mask to 13° reduces the median return by approximately 20% at 20 h and by approximately 12–14% at 40–60 h.

The upper tail remains almost unaffected, with P95 changes close to zero for both favourable and adverse variations. The lower-tail statistic is undefined at 20 h and 40 h because the nominal reference percentile is zero, and becomes defined only at 60 h, where the change remains modest.

Case 3: sensitivity to local surface geometry perturbations (tilt / slope)

Beyond the local-horizon assumption, the landed geometry may also depart from the idealised nominal ellipsoidal shape because of local tilt or slope. In Section 3.3, the nominal case was defined using the local surface frame and zenith-angle visibility condition, while the corresponding tilt/slope perturbation was introduced through Equation (3.15) and illustrated in Figure 27. This sensitivity case evaluates how mild (15°) and severe (30°) tilt/slope conditions affect the achievable DTE return relative to the nominal landed geometry.

Figure 92 shows the effect of a mild perturbation on the achievable DTE return. The figure reports the percentage variation in achievable data-volume relative to the nominal landed configuration for a δ = 15° tilt/slope, summarizing also in this case the results through the P5, P50 and P95 statistics of the global no-control population at 20 h, 40 h and 60 h. Negative values indicate a reduction of achievable data-volumes in presence of the tilt/slope with respect to the reference.

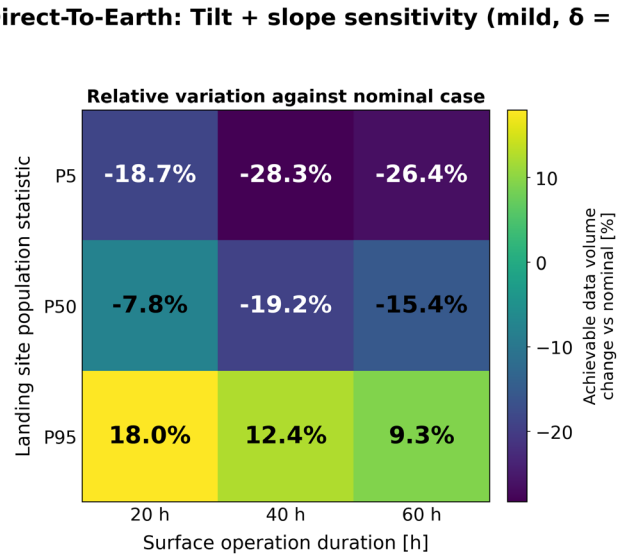


Figure 92: RCS2 DTE sensitivity to mild local tilt and slope

Figure 92 shows that even a mild tilt/slope perturbation affects the DTE performance distribution, although the impact is not uniform across the population. The lower and median parts of the distribution degrade at all three mission-duration markers. The P5 statistic decreases by approximately 19–28%, while the P50 statistic decreases by approximately 8–15%. This indicates that mild local surface-orientation perturbations can already reduce the achievable return for adverse and median realised cases.

By contrast, the upper tail shows a positive variation relative to the nominal case, with the P95 statistic increasing by approximately 18% at 20 h, 12% at 40 h and 9% at 60 h, indicating that, for a limited subset of favourable realised geometries, the tilted boresight can improve the instantaneous alignment with the Earth direction relative to the nominal local-zenith pointing assumption.

Figure 93 extends the tilt/slope assessment to δ = 30°, representing a severe perturbation, to assess whether the moderate degradation observed for the mild case becomes a first-order limitation under more adverse landed conditions.

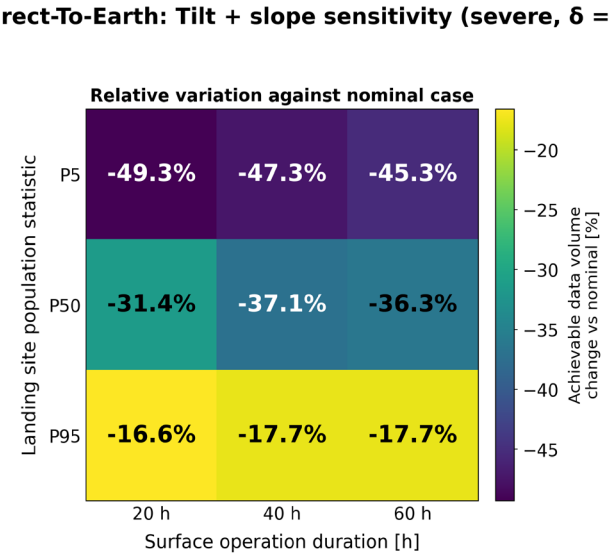


Figure 93: RCS2 DTE sensitivity to severe local tilt and slope

Figure 93 shows that, under the no-control case, a severe tilt/slope perturbation has a consistently adverse effect on DTE performance across the analysed statistics and mission durations. The P50 degradation is approximately 31–37%, indicating that the median achievable data-volume is reduced by about one third relative to the nominal landed configuration. The lower tail is affected more strongly, with P5 reductions of approximately 45–49%. Unlike in the mild case, the upper tail also degrades at all three mission-duration markers, with P95 reductions of approximately 17–18%. This indicates that the benefits observed for a limited subset of favourable mild-perturbation cases are no longer present under severe perturbations.

The severe case therefore acts as a first-order limitation on DTE performance under no-control conditions, affecting not only conservative outcomes but also the more favourable realised geometries.

Figure 92 and Figure 93 show that, under no-control conditions, RCS2 DTE performance is moderately sensitive to mild local tilt/slope perturbations, mainly through degradation of the lower and median parts of the distribution. Severe departures from the nominal landed configuration instead degrade the full distribution and therefore represent a stronger limitation for DTE performance.

Case 4: sensitivity to telemetry planning and operability constraints

The preceding DTE sensitivity cases focus on local visibility and landed attitude. The final DTE sensitivity considered here addresses telemetry-rate planning, following the approach already applied to RCS1 in Section 4.3. In the unconstrained DTE model, the

selected telemetry rate can change whenever the instantaneous link conditions allow a different feasible mode. In practice, however, rate changes would be planned and executed only at discrete opportunities, because reconfiguration requires interruption, update and re-establishment of the link.

Consistent with the corresponding RCS1 case, the limitation is modelled by allowing telemetry-rate updates only at fixed intervals. For RCS2, update intervals of $K = 2$, $K = 4$, and $K = 8$ are considered as sensitivity cases rather than as operational requirements. The shorter intervals reflect the substantially shorter RCS2 operational timescale and the more rapid variation of the landed DTE link. Different update-window timings are also considered, allowing the effect of update frequency to be separated from the timing alignment between the update windows and the short favourable DTE interval. This TM-rate planning constraint is potentially more severe for RCS2 than for RCS1 because the landed mission is concentrated within this rapidly varying link interval.

Figure 94 shows the effect of telemetry-rate update constraints on the RCS2 DTE case. The upper panel shows the selected telemetry-rate evolution over time for the unconstrained case and for the three K update intervals. The middle panel shows the corresponding cumulative data-volumes evolution, while the lower panel shows the cumulative data-volumes as a percentage of the unconstrained case. For each constrained update interval, the median response with respect to the timing of the update sensitivity and the min–max band are shown.

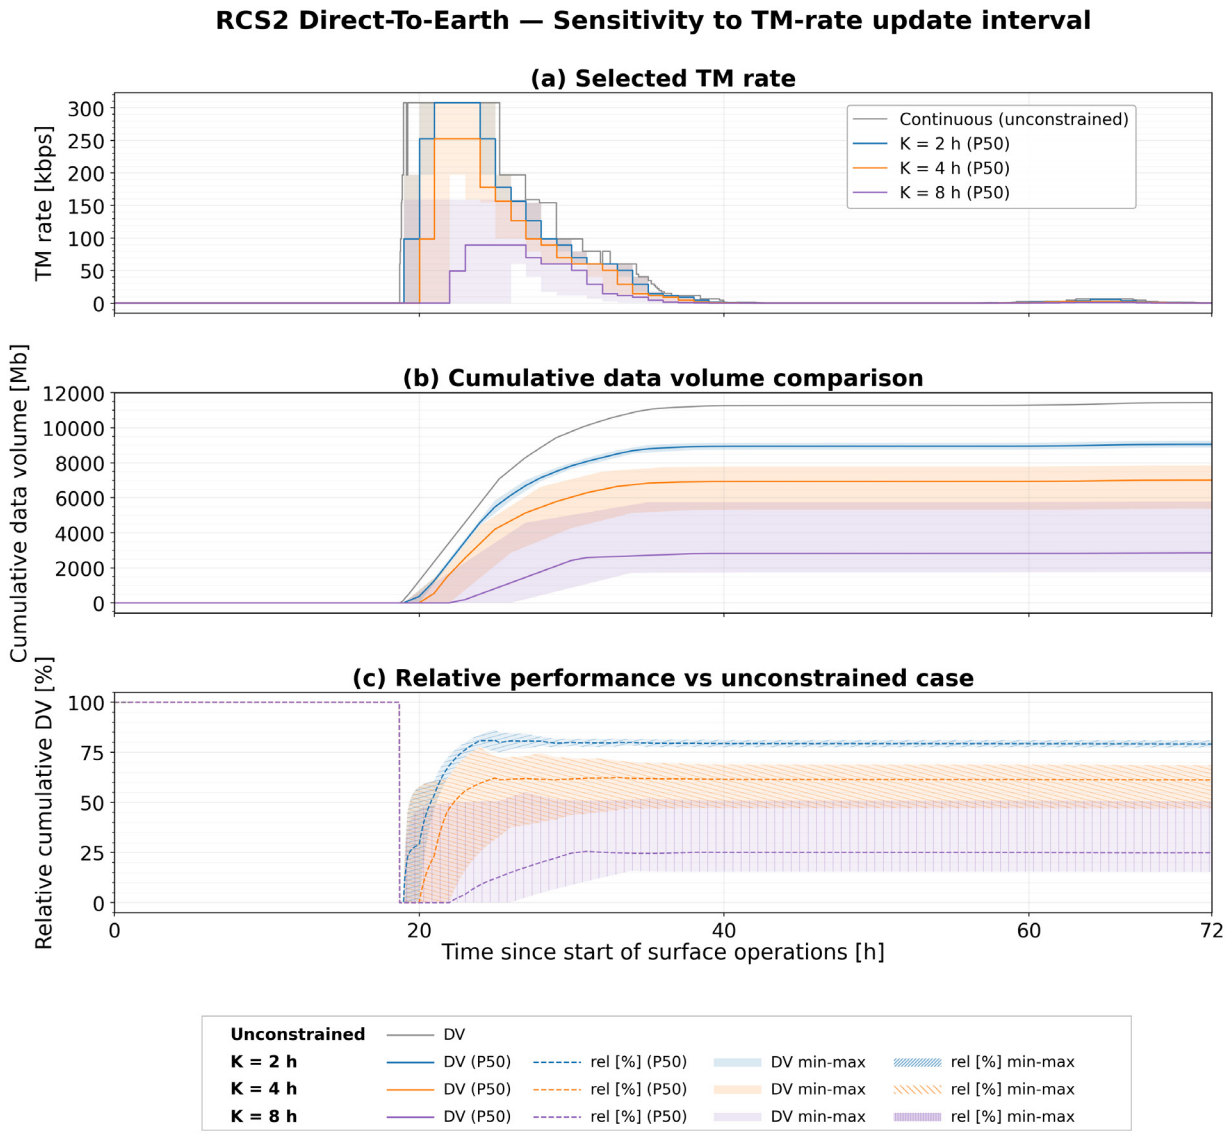


Figure 94: RCS2 DTE sensitivity to telemetry planning and operability constraints

Figure 94 shows that the nominal RCS2 DTE return is highly sensitive to telemetry-rate update limitations. In the unconstrained case, the selected rate can follow the instantaneous link conditions and exploit the short interval of favourable geometry around the main communication opportunity. When updates are restricted to discrete K-hour intervals, the selected rate becomes progressively less able to track the available link margin, and the achievable cumulative return is reduced. This behaviour is visible in panel (a), where the constrained profiles increasingly smooth and lower the selected telemetry-rate evolution as K increases.

The effect on cumulative return is shown in panels (b) and (c). For $K = 2\,h$, the median cumulative data-volume remains relatively close to the unconstrained case, reaching approximately 80% of the nominal return. For $= 4\,h$, the median relative performance

decreases to approximately 60%, while for $= 8\ h$ it falls to about 25%. The hatched min–max bands show that the result also depends on the update-window timing: even for a fixed K, different alignments with the short favourable DTE interval can produce a non-negligible spread in outcome for several percentage points.

The main result is that RCS2 DTE performance is sensitive not only to how frequently telemetry-rate updates are allowed, but also to when those updates occur relative to the narrow early communication opportunity. When the link is degrading, the selected rate must also remain feasible until the next update instant, which requires some prediction of the future link condition. Telemetry planning and update frequency can therefore become a first-order practical limitation for the DTE architecture, even when the underlying geometry and link conditions remain nominal.

DTE robustness summary

The DTE sensitivity cases show that RCS2 DTE performance is strongly dependent on the realised landing and operational conditions. The strongest sensitivity is associated with control over the initial rotational phase, because this determines whether the short early interval of favourable Earth visibility and reduced Earth–Apophis range can be exploited. Heading control alone provides only limited benefit.

The minimum elevation mask mainly affects adverse and median realised cases. Mild tilt/slope surface perturbations reduce the lower and median parts of the distribution, whereas severe perturbations degrade the full distribution, including the favourable upper tail.

Telemetry scheduling is also a significant practical driver, as restrictive update schemes can prevent the link from following the rapid variation in link conditions and cause substantial data-volume loss.

The DTE architecture can therefore provide high return under favourable realised conditions, but its performance is not broadly robust across the analysed RCS2 uncertainty space. This finding is used in the Chapter 6 architecture evaluation.

5.3.2 Inter-Satellite-Link sensitivity and robustness

The ISL sensitivity analysis examines how RCS2 performance changes when the landed geometry and antenna assumptions depart from the nominal baseline.

Case 1: sensitivity to landing timing and heading control authority

As in the corresponding DTE sensitivity *Case 1* in Section 5.3.1, this case examines how the achievable RCS2 ISL return changes when different levels of authority are assumed over the landed initial conditions. The nominal ISL analysis treats both the initial rotational phase of Apophis and the landed antenna heading as uncontrolled variables within the analysed population. Some authority over landing timing and/or final landed heading could allow the lander to achieve more favourable ISL visibility and antenna geometry.

This sensitivity case therefore evaluates the achievable cumulative data return for the same four levels of control authority investigated for DTE: no control, heading-only control, initial-phase-only control and full control. The objective is to quantify the potential ISL performance benefit of such control authorities, while their potential implementations are outside the scope of the analysis.

Figure 95 shows how landing-control authority affects the RCS2 ISL cumulative data return. The first three panels compare the landing-site population for the four control cases, at 20 h, 40 h and 60 h, respectively. The sites are ordered according to the no-control case.

The bottom panel provides a compact median-performance summary. It plots the median data-volume at 20 h against the median data-volume at 40 h, while the vertical whisker indicates the corresponding 60 h value. Best performing control cases are therefore located toward the upper-right corner.

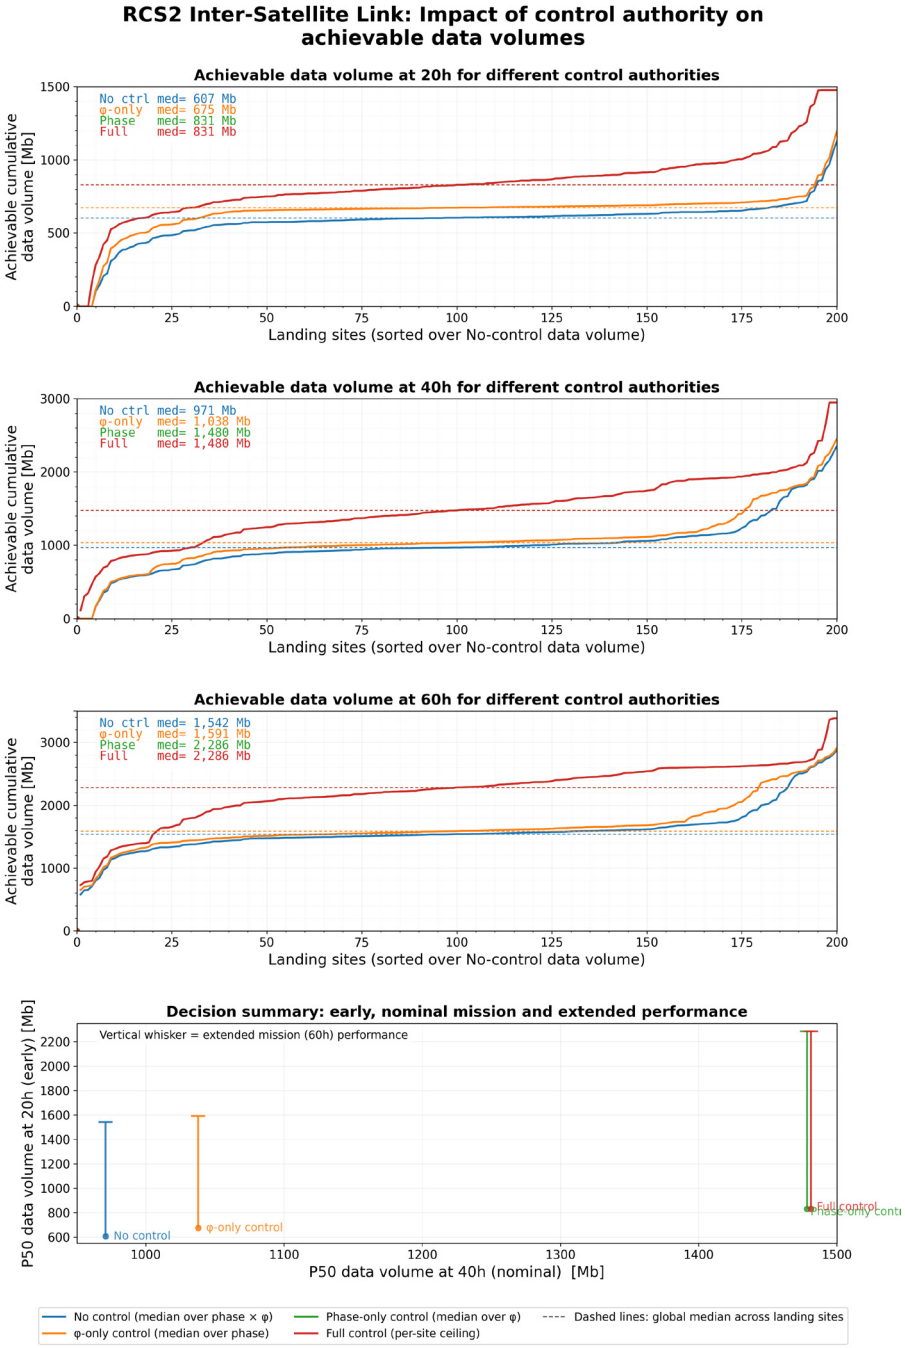


Figure 95: RCS2 ISL impact of control authority on achievable cumulative data-volume

Figure 95 shows that nominal ISL performance varies with the degree of control authority over the landed initial conditions, but that the benefit is uneven across the two control variables. The no-control case gives the lowest achievable return across the landing-site population at all three mission times, while φ-only control produces only a small improvement. This limited change is visible mainly as a slight upward shift of the central part of the population, with the lower and upper tails remaining broadly similar to the uncontrolled case.

The phase-only and full-control cases are effectively coincident at 20 h, 40 h and 60 h. They shift much of the population towards higher cumulative data-volumes, especially across the central and upper parts of the ranking, while the lowest-ranked sites remain comparatively weak. The overlap between these two cases indicates that the main control benefit comes from selecting the initial rotational phase, while additional landed-heading optimisation provides no visible further improvement.

The median cumulative data-volumes and relative improvement against the uncontrolled case are summarised in Table 5.3.

Table 5.3 - RCS2 ISL performance under different landing-control authorities

Control authority	20 h [Mb]	40 h [Mb]	60 h [Mb]	Gain relative to no control at 20 h	Gain relative to no control at 40 h	Gain relative to no control at 60 h
No control	607	971	1,542	1.0x	1.0x	1.0x
Heading-only control	675	1,038	1,591	1.1x	1.1x	1.0x
Phase-only control	831	1,480	2,286	1.4x	1.5x	1.5x
Full control	831	1,480	2,286	1.4x	1.5x	1.5x

Table 5.3 confirms that the main gain is associated with control over the landing phase rather than with heading optimisation. Relative to the uncontrolled case, φ-only control yields only a small increase in median achievable return, whereas phase-only control raises the median more substantially at all three mission times. However, the improvement remains limited in magnitude, reaching at most a factor of approximately 1.5 in the median values reported here. Full control gives the same median values as phase-only control.

Figure 95 also shows that control authority does not remove the weakest part of the population even with phase-only or full control.

This behaviour is also visible in the bottom summary panel of Figure 95, which provides a compact visual representation of the Table 5.3 results. The phase-only and full-control cases occupy the same position and remain clearly separated from the no-control and φ-only cases. The 60 h whiskers preserve this ordering, showing that the control advantage remains as the operational window is extended.

Case 2: sensitivity to the local horizon assumption

As in the corresponding RCS2 DTE *Case 2* in Section 5.3.1, the nominal ISL analysis adopts a 10° minimum-elevation mask to represent local horizon masking and near-horizon degradation. This assumption is introduced in Section 3.3 and applied through the zenith-angle visibility condition of Equation (3.14). Since this assumption is not derived from a

detailed terrain model of the local landing environment, it is appropriate to assess how strongly the ISL results depend on its value. The present case therefore considers both the penalty associated with the nominal baseline mask and the sensitivity of the achievable ISL performance to modestly lower or higher minimum-elevation assumptions.

Figure 96 shows the reduction in achievable ISL cumulative data-volume introduced by the nominal minimum-elevation $\varepsilon_{\min} = 10°$ assumption relative to an ideal free-local-horizon reference with $\varepsilon_{\min} = 0°$. Consistent with the corresponding DTE assessment, the reduction is evaluated using the P5, P50 and P95 statistics of the global no-control population at 20 h, 40 h and 60 h.

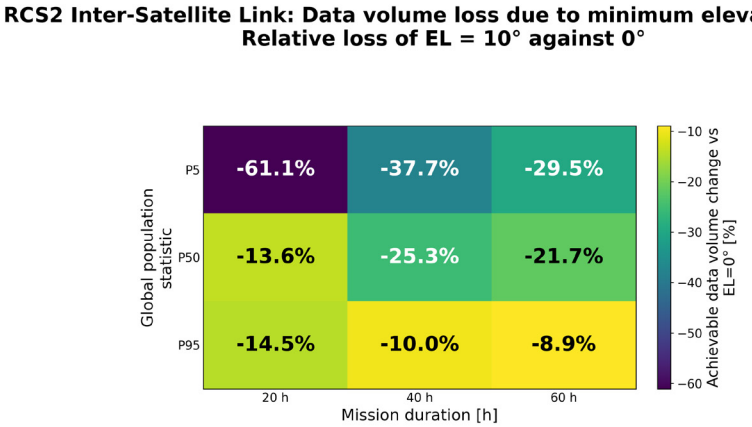


Figure 96: RCS2 ISL impairment due to the nominal minimum-elevation mask

Figure 96 shows that the adopted 10° horizon mask produces a consistent reduction in achievable ISL data-volume, although the magnitude of the reduction depends strongly on the part of the population considered. The strongest relative impact occurs at the P5 level, where the loss reaches approximately 61% at 20 h and remains close to 30% at 60 h, substantially degrading an already limited population portion.

At the median level, the reduction is smaller but still significant, ranging from approximately 14% at 20 h to approximately 22–25% over the later time windows. By contrast, the P95 degradation remains comparatively limited, at approximately 9–15%. The minimum-elevation mask therefore mainly affects the lower and central parts of the ISL population, while the most favourable realised cases retain most of their achievable return.

Having shown that the nominal minimum-elevation mask introduces a non-negligible impairment, Figure 97 examines how modest variations around the adopted 10° minimum-elevation baseline affect the no-control ISL return. The figure reports the percentage variation in achievable cumulative data-volume for $\varepsilon_{\min} = 7°$ and $\varepsilon_{\min} = 13°$ relative to the nominal

$\varepsilon_{\min} = 10°$ case, showing how those changes in the local visibility threshold affect the P5, P50 and P95 statistics at 20 h, 40 h and 60 h.

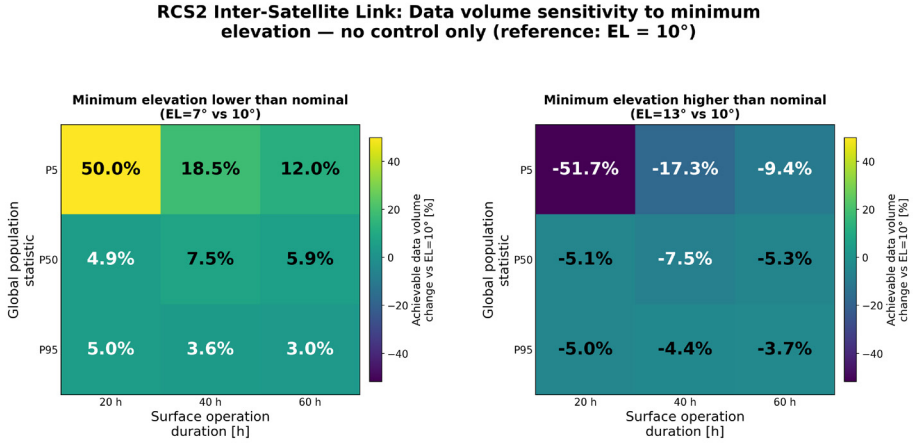


Figure 97: RCS2 ISL sensitivity to the minimum-elevation assumption

Figure 97 shows that modest variations around the nominal 10° minimum-elevation mask affect the no-control ISL results mainly in the lower tail of the population. Reducing the mask to 7° increases the P5 return by approximately 50% at 20 h and 12% at 60 h. Increasing the mask to 13° produces the opposite effect, as expected, with P5 reductions of approximately 52% at 20 h and 9% at 60 h.

The median and upper-tail responses are much less affected. For the 7° case, P50 increases by approximately 5–8% and P95 by approximately 3–5%. For the 13° case, the corresponding reductions remain approximately 5–8% for P50 and 4–5% for P95. The minimum-elevation uncertainty therefore mainly affects adverse ISL realisations, while the central and favourable parts of the population remain comparatively stable around the nominal 10° assumption.

This result should be interpreted as a first-order geometric sensitivity rather than as a detailed RF-channel assessment or as a model of the waveform-level robustness of the reference terminal. However, the GAMALINK RF air interface includes frequency hopping spread spectrum (FHSS), which may provide additional resilience to some narrow-band or local propagation impairments [34].

Case 3: sensitivity to local surface geometry perturbations (tilt / slope)

As for the DTE sensitivity analysis *Case 3* presented in Section 5.3.1, the ISL sensitivity analysis also considers how local surface tilt or slope modifies the nominal landed geometry. The perturbation is defined in Section 3.3 through the tilted-geometry formulation of Equation (3.15) and illustrated in Figure 27. This case evaluates how mild (15º) and severe (30º) local surface perturbations affect the achievable ISL return relative to the nominal landed configuration.

Figure 98 shows the effect of a mild local surface-geometry perturbation, reporting the percentage variation in achievable cumulative data-volume relative to the nominal landed configuration for a $\delta = 15^{\circ}$ tilt/slope case. The P5, P50 and P95 statistics of the no-control landing-site population at 20 h, 40 h and 60 h are reported.

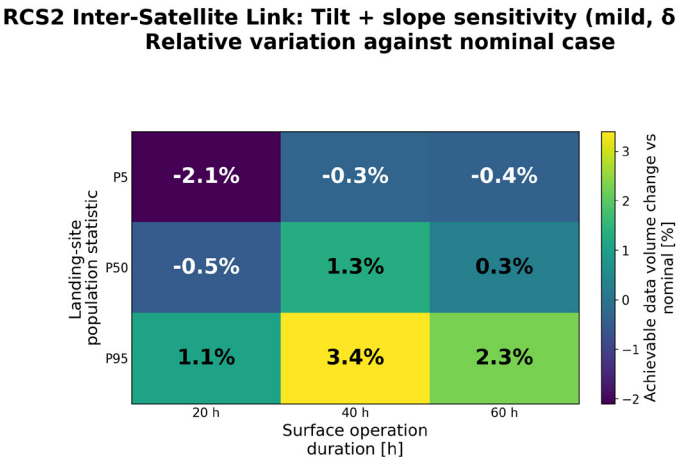


Figure 98: RCS2 ISL sensitivity to mild local tilt and slope

Figure 98 shows that a mild tilt-and-slope perturbation has only a limited effect on ISL performance. Across all three mission times, the variations remain small, generally within a few percent of the nominal case. The P5 values show a slight degradation, reaching -2.1% at 20 h, while the P50 values remain close to zero and become slightly positive at 40 h and 60 h. The P95 values are also positive, with gains of approximately 1–3%. This indicates that mild local surface-geometry perturbations have little effect on ISL performance at population level. The small positive changes occur in cases where the perturbed antenna orientation improves alignment with the RAMSES spacecraft.

Figure 99 extends the tilt/slope assessment to the severe perturbation case, with $\delta = 30^{\circ}$, assessing whether stronger departures from the nominal landed geometry produce a more significant degradation of ISL performance.

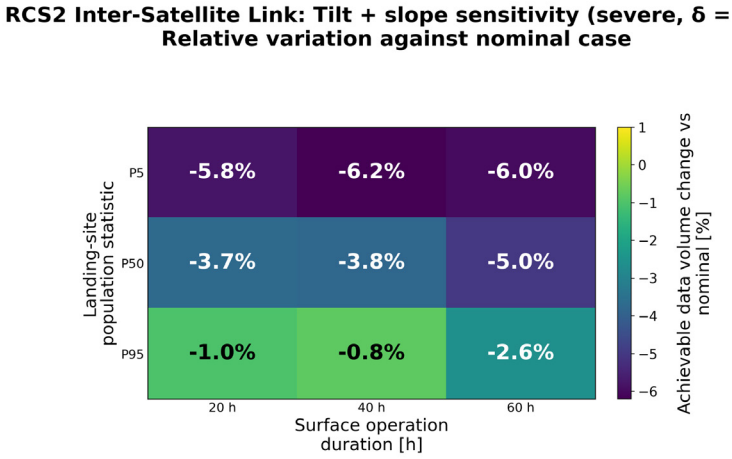


Figure 99: RCS2 ISL sensitivity to severe local tilt and slope

Figure 99 shows that, under the severe tilt/slope perturbation, the impact on ISL performance becomes consistently adverse across the landing-site population, although its magnitude remains limited across the full population. The P5 degradation is approximately 6% at all three mission times, while the P50 values are reduced by approximately 4–5%. The P95 values also become negative, with reductions of approximately 1% at 20 h to approximately 3% at 60 h.

Figure 98 and Figure 99 show that the RCS2 ISL architecture remains comparatively robust to local tilt-and-slope perturbations, with only modest degradation even in the severe case.

Case 4: sensitivity to antenna-surface coupling or landed antenna-pattern effects

The final RCS2 ISL sensitivity case investigates whether the communication results are affected by changes in the effective CubeSat antenna radiation pattern after touchdown. The baseline configuration defined in Section 3.6.2 assumes two opposite low-profile antennas, with the nominal simulated pattern shown in Figure 38. In this sensitivity case, the nadir-facing antenna is assumed to be shielded by the asteroid surface. The CubeSat is expected to rest in direct contact with the surface after touchdown, and the radiated contribution from the nadir-facing antenna is treated as completely blocked. The effective landed antenna pattern therefore retains only the zenith-facing antenna contribution influenced by the RCS2 structure. This case is not intended to model antenna–surface interaction in detail. Instead, it provides a bounding sensitivity assessment of whether RCS2 ISL performance is sensitive to the post-landing antenna-pattern assumption. The resulting ISL performance is compared with the nominal two-antenna baseline using the same received-power, mode-selection and data-volume framework defined in Section 3.4.

Figure 100 shows the antenna-gain envelopes for the nominal two-antenna configuration and for the assumed landed configuration in which only the top antenna contributes. The shaded region indicates the angular range excluded by the nominal minimum-elevation mask $\varepsilon_{\min} = 10°$.

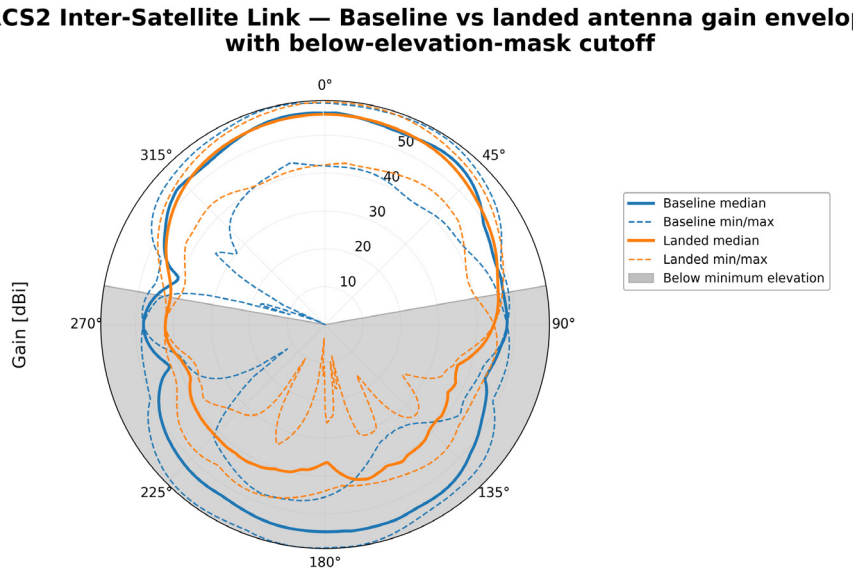


Figure 100: RCS2 ISL baseline and landed antenna-pattern gain envelopes

Figure 100 shows that the landed-pattern assumption differs most strongly from the nominal two-antenna pattern in the angular region below the minimum-elevation mask (80º to 280º). Within the ISL-visible angular range, the gain envelopes remain broadly similar, with only modest local differences. This indicates that removing the nadir-facing antenna changes the overall antenna pattern, but has only limited effect on the part of the pattern that contributes to ISL cumulative data-volume performance.

Figure 101 shows the achievable cumulative data-volume distributions obtained with the nominal two-antenna baseline and with the assumed landed antenna configuration at 20 h, 40 h and 60 h of surface operations time. The median, adverse P5, favourable P95 and P10–P90 interval are shown for each antenna-pattern assumption.

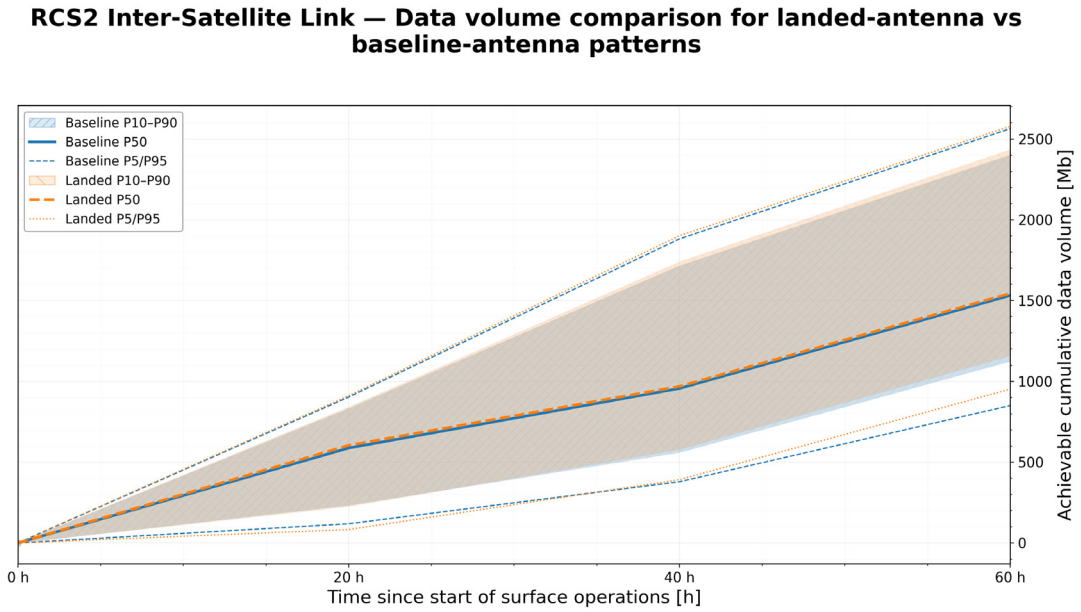


Figure 101: RCS2 ISL data-volume comparison for baseline and assumed landed antenna

Figure 101 shows that the landed-pattern assumption has only a marginal effect on the global ISL data-volume distribution. The median and the upper tail remain almost unchanged, with only a very small increase relative to the baseline case across the surface operation window. The P10–P90 interval is also largely preserved: the landed-pattern envelope remains close to the baseline envelope at all three time markers, with only a slight improvement towards the 40 h to 60 h window. This indicates that the central spread of the population is not strongly altered by removing the nadir-facing antenna contribution, although minor improvements are visible. The lower tail behaves less uniformly: at earlier times, the poorer cases are slightly degraded relative to the baseline, whereas at later times this adverse effect is reduced and becomes an improvement of approximately 10%.

The assumed landed pattern produces a small improvement in the global ISL distribution, while introducing a minor early-time penalty in the lower tail.

Building on the nominal ISL landing-site ranking introduced in Figure 86 and its surface distribution, Figure 102 shows how the landed antenna assumption affects the robustness ranking of the RCS2 ISL landing-site population. The top panel shows the robust-goodness 2D map for the nominal baseline antenna case, while the second panel shows the corresponding map for the landed single-antenna case. The third panel overlays the global robust-goodness ranking for the two antenna assumptions, with sites ordered according to the baseline-antenna ranking on a common normalised scale. The final panel shows on a 2D map

the difference in robust-goodness score across the candidate surface sites, also on the normalised scale.

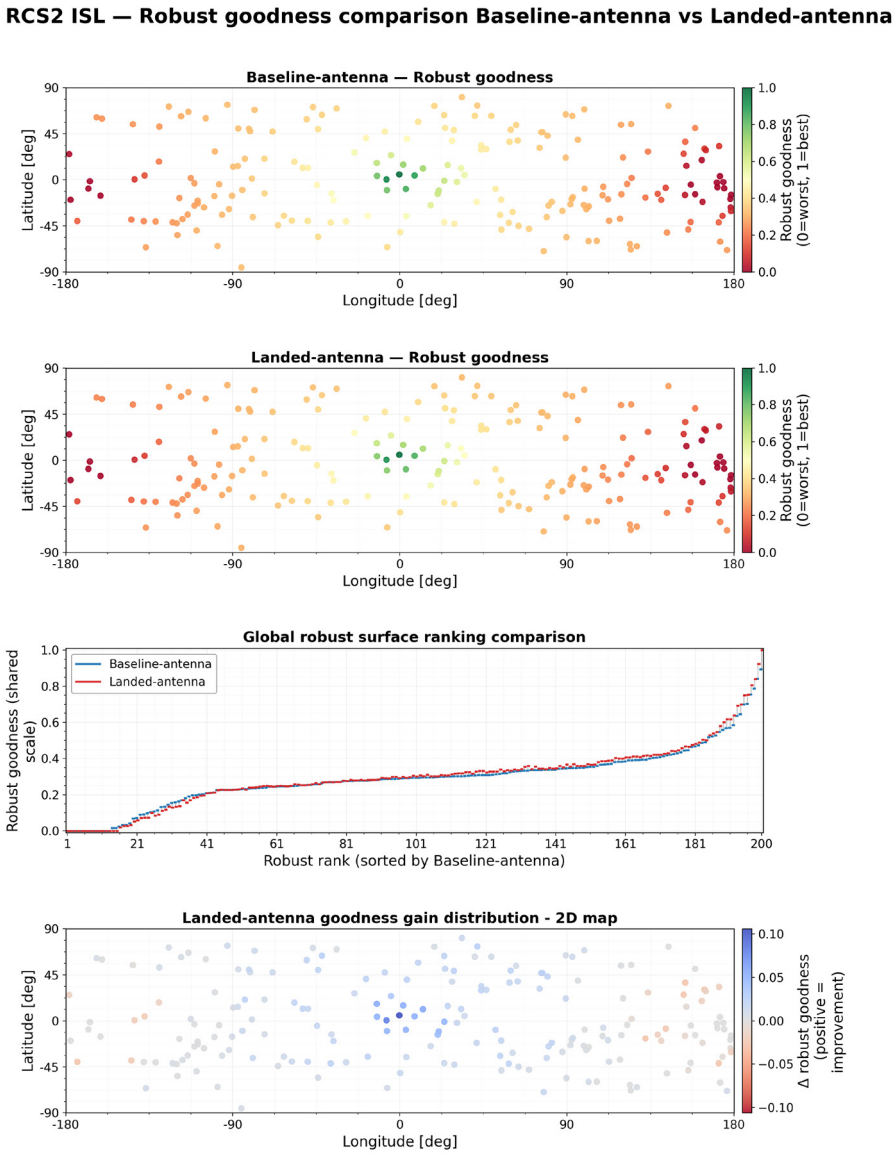


Figure 102: RCS2 ISL robust goodness comparison for baseline and assumed landed antenna

Figure 102 shows that the landed-pattern assumption does not substantially alter the global spatial distribution of ISL performance. The same broad regions remain favourable or unfavourable, with the robust ranking only weakly altered within the ±10% range, mainly in the lower and upper tails. The goodness-difference map shows mostly small local changes where the most favourable site clusters become slightly more pronounced, while some weaker regions deteriorate slightly further.

Figure 101 to Figure 102 show that the assumed post-landing antenna-pattern change has only a limited effect on RCS2 ISL performance. The landed-pattern assumption does not

redefine the overall ISL site ranking, although it can slightly increase the contrast between poorer and more favourable realised cases and produces a small improvement in cumulative data return. The nominal two-antenna baseline analysis therefore remains valid for the level of fidelity considered in this thesis.

ISL robustness summary

The ISL sensitivity cases show that RCS2 ISL performance is driven mainly by the realised surface-to-RAMSES geometry and by the degree of control authority over the landed initial conditions. Control over the initial phase provides the main opportunity for improvement, while heading control adds little in the analysed cases.

The sensitivity to the minimum-elevation assumption is concentrated mainly in the lower tail of the population. Reducing the mask from the nominal 10º to 7º improves the achievable return, especially for adverse realisations, while increasing it to 13º produces the corresponding loss. The median and upper-tail returns change only modestly in both cases. Mild tilt/slope perturbations have little impact on the ISL performance, while severe perturbations produce a consistent but still moderate degradation across the population. The landed-pattern analysis produces the smallest changes without altering the spatial ranking of favourable and unfavourable sites.

The main result is that RCS2 ISL performance is driven primarily by the landing location and the resulting geometry with RAMSES, while the other evaluated parameters have smaller, or only marginal, effects.

Table 5.4 consolidates the DTE and ISL robustness findings discussed in Sections 5.3.1 and 5.3.2.

Table 5.4 - RCS2 sensitivity and robustness status summary

Sensitivity / robustness indicator	DTE response	ISL response	Reference metric / remarks
Landing timing / heading control authority	Very high impact. Phase control dominates; full control gives approximately 20× gain at 40 h relative to no control.	Moderate impact. Phase control dominates; phase/full control gives approximately 1.5× gain at 40–60 h.	Median cumulative return gain vs no control.
Minimum-elevation / local-horizon assumption	Moderate impact; 7°–13° variation mainly affects median/lower-tail return.	Moderate impact; 7°–13° variation mainly affects lower-tail return.	Percentile variations against 10° baseline.
Mild local tilt/slope, δ = 15°	Mixed / moderate impact. Lower and median cases degrade, while the upper tail can improve through favourable pointing changes.	Very low impact. Variations remain within a few percent at population level.	Percentile data-volume change relative to nominal landed geometry.
Severe local tilt/slope, δ = 30°	High impact. Median and lower-tail performance degrade substantially; upper-tail cases also degrade.	Low impact. Degradation is consistent but remains limited to a few percent.	Percentile data-volume change relative to nominal landed geometry.
Telemetry-rate planning / update interval	Very high operational impact. Retained median return falls from approx. 80% for 2 h updates to approx. 25% for 8 h updates.	Not applicable. ISL automatic adaptive data-mode selection.	Architecture-specific operational constraint.
Landed-Antenna pattern	Not assessed as a DTE sensitivity case.	Low impact. Median and spatial ranking remain broadly unchanged.	ISL-only case; nadir-facing antenna blocked by local surface.

These results provide the RCS2 sensitivity and robustness basis for the architecture-level comparison in Chapter 6.

5.4Chapter summary

This chapter presented the lander RCS2 communication performance analysis for the DTE and ISL architectures, using the modelling framework developed in Chapter 3. The analysis covered nominal landed-operation conditions and selected sensitivity cases.

For DTE, performance is driven by the combination of Earth–Apophis range evolution and local landed geometry. High data-volume return is possible when the short favourable Earth-range interval coincides with Earth visibility from the landing site. Under nominal assumptions, favourable realised cases can reach approximately 10–11 Gbit, with telemetry rates reaching the highest available. However, the median data-volume return is much lower,

approximately 0.17, 0.53 and 0.59 Gbit after 20 h, 40 h and 60 h, respectively, while the lower tail remains close to zero.

For ISL, performance is also driven mainly by geometry, but the relevant driver is the local surface-to-RAMSES geometry. Data return is accumulated during each available RCS2–RAMSES contact opportunity without degrading trends. The best case reaches approximately 3.4 Gbit by 60 h. The median return increases progressively, from approximately 0.5 Gbit at 20 h to 1.0 Gbit at 40 h and 1.5 Gbit at 60 h. Although the most favourable ISL return is lower than the highest DTE cases, its median and lower-tail behaviour are stronger.

The sensitivity analysis showed that the two architectures have different robustness drivers. DTE is strongly affected by landing conditions, local surface and operational constraints. Full control over the landing conditions produces a large median gain, while severe tilt/slope perturbations and restricted telemetry-rate updates substantially reduce the achievable return. ISL is also affected by control authority and local-horizon assumptions, but the remaining sensitivity cases showed marginal effects. Its robustness is driven primarily by landing location and the resulting geometries, while tilt/slope and landed-antenna pattern have a smaller, or only marginal, effect.

These results provide the RCS2 performance basis for the architecture-level evaluation in Chapter 6.

The analysis therefore addresses the RCS2 contribution to Objectives 2, 3 and 4 by providing the nominal performance assessment, sensitivity and robustness assessment, and statistical characterisation of performance for the landed CubeSat case.

Chapter 6 Communication architecture evaluation and recommendations

6.1Introduction

This chapter evaluates the Direct-To-Earth and Inter-Satellite-Link architectures for the two RAMSES CubeSat cases. The evaluation uses the performance results from Chapters 4 and 5 and interprets them against the analysis requirements defined in Section 2.5 and detailed in Appendix I, providing an architecture-level assessment of the relative suitability of DTE and ISL for each CubeSat.

The evaluation therefore considers the architectures beyond nominal data return, including the operational and implementation conditions under which that return can be achieved.

6.2Evaluation framework and criteria

The evaluation framework used in this chapter is derived from the research aim and objectives in Chapter 1, and from the analysis requirements and implementation constraints defined in Chapter 2. Its purpose is to compare the DTE and ISL options at architecture level, considering their nominal capability, robustness, implementation implications and operational dependence within the RAMSES CubeSat mission context.

Following the structure adopted in Chapters 4 and 5, and reflecting the specific characteristics of each mission case, the evaluation is organised separately for RCS1 and RCS2. The results are then considered together to assess their implications at RAMSES mission level and to identify a possible common RAMSES CubeSat communication architecture. Where used in the assessment, “Goal” identifies a desirable target value beyond the basic requirement.

Table 6.1 defines the evaluation structure adopted in this chapter. The first column identifies the evaluation criteria, while the second states the purpose of each criterion in the assessment. The third column identifies the relevant results from Chapters 4 and 5 for the evaluation, such as cumulative data-volume, percentile spread, sensitivity behaviour and resource implications. The final column relates each criterion to the analysis requirements and constraints defined in Table Ap.3.

Table 6.1 - Evaluation criteria for the architecture-level assessment

Evaluation criteria	Assessment focus	Assessment evidence	Table Ap.3 assessment check
Baseline communication performance	Nominal link feasibility and data-return capability.	Link feasibility, rate/mode, availability and cumulative data-volume.	**R4:** ISL data return assessed against 9.2 Gb/CubeSat allocation. **R5**: DTE return ≥ 12 Gb. **R6:** link feasibility, rate/mode and data-volume.
Nominal performance dispersion	Spread and lower-tail compliance of nominal performance across analysed cases.	Percentile spread, lower-tail data-volume and compliance of adverse cases against the relevant data-return check.	**R4/R5:** data return checked for median, favourable and lower-tail cases where applicable. **R6:** dispersion across analysed cases.
Sensitivity to degraded or uncertain conditions	Effect of sensitivity cases on architecture performance.	Median stability, percentile-envelope changes and lower-tail shifts	**R6:** robustness under degraded or uncertain conditions.
CubeSat resource and accommodation impact	Compatibility with CubeSat resources and accommodation.	Power demand, mass, volume, antenna accommodation and implementation complexity.	**R1:** ≥ 3U payload retained; comms goal < 1U excl. antenna; mass ≤ 1.33 kg/U. **R2:** 28 V input; CubeSat-compatible power envelope. **R3:** existing or near-existing designs.
Operational dependence and architectural risk	Architecture dependencies and their operational consequences.	Conditionality of data return and impact on flexibility/robustness.	**R3**: implementation credibility. **R6**: architectural dependence and operational flexibility.

The last column of Table 6.1 provides traceability to the requirements and constraints defined in Table Ap.3. R4–R6 support the data-return and robustness assessment, while R1–R3 are used where accommodation, power or implementation implications are considered. For R2, the RCS1 assessment uses the communication-system DC power envelope estimated in Equation (3.44), while the RCS2 assessment considers the peak DC power demand of the candidate communication subsystems against the available platform information discussed in Section 3.6.1.

A common set of assessment outcomes is used in Sections 6.3 and 6.4 to keep the architecture evaluation consistent across the RCS1 and RCS2 cases. *Pass* denotes a check satisfied under the adopted baseline assumptions. *Assumption-dependent pass* denotes a check satisfied only while a specific operational or modelling assumption is maintained. *Partial* denotes a case where some elements of the check are satisfied, but at least one important limitation remains. *Fail* denotes a check that is not satisfied under the analysed baseline.

The RCS1 and RCS2 evaluations later in this chapter include compact summary tables that use qualitative marks defined in Table 6.2. These marks provide a subjective assessment of relative advantage or limitation based on the preceding analysis and the applicable requirements identified in Table 6.1, and are not summed into a total score.

Table 6.2 - Comparative assessment marks

Mark	Meaning
++	Strong advantage
+	Advantage
/	Mixed or neutral result
–	Limitation
––	Major limitation

For each CubeSat mission, the assessment is carried out in two steps. First, the key results from the corresponding analysis chapter are summarised, focusing on the factors that affect the architectural decision. Second, the DTE and ISL options are assessed against the evaluation criteria in Table 6.1 to identify the relative strengths, limitations and conditions of applicability.

The purpose of the evaluation is not to assign a single universal score to DTE or ISL. Instead, it identifies the conditions under which each option is credible for RCS1 and RCS2, and how those results support the selection of a common RAMSES CubeSat communication architecture.

6.3RCS1 architecture evaluation

This section evaluates the DTE and ISL communication architectures for the orbiting CubeSat, RCS1, using the Chapter 4 results and the criteria defined in Section 6.2.

6.3.1 RCS1 architecture observations

The key RCS1 observations relevant to the architecture-level assessment are summarised below.

Observation 1: RCS1 DTE link feasibility is continuous, while most data return is concentrated around the close-encounter interval.

The nominal analysis in Section 4.2.1 shows that the RCS1 direct link to Earth remains feasible throughout the analysed mission timeline, but that most cumulative data return is

concentrated within approximately ten days before and after closest approach as a result of the Earth–Apophis range evolution and the adopted antenna and ground-segment profiles. From an architecture perspective, this concentration may affect operations and scientific activities, because data generated earlier would be returned with delay, while data generated after the favourable interval may have limited opportunity for high-volume transmission.

Observation 2: RCS1 DTE robustness is conditional on maintaining Earth-pointing capability.

The DTE sensitivity cases show that the nominal capability depends on preserving the assumed Earth-pointing conditions and antenna-use profile. When RCS1 is forced to rely on the LGA rather than the nominal MGA configuration, the achievable telemetry rate is strongly reduced, especially away from closest approach where the Earth–Apophis range is larger. In adverse LGA coverage cases, telemetry feasibility can be lost for extended parts of the mission, and the residual carrier-level can fall below the adopted threshold which would be a direct operability concern for RCS1.

This affects not only cumulative data-volume, but also the ability to maintain the communication path for basic operational purposes. Since the analysis already assumes a high-performance ground segment, loss of effective MGA Earth-pointed performance cannot be readily compensated. RCS1 DTE therefore remains viable under the nominal assumptions, but its robustness depends on maintaining effective Earth-pointed MGA operation, especially away from closest approach.

Observation 3: RCS1 DTE data return is sensitive to telemetry-rate planning around closest approach.

The telemetry-rate planning sensitivity analysis in Section 4.3 shows that RCS1 DTE performance depends on the ability to follow the rapidly varying link conditions around closest approach. Under the analysed update constraints, the achievable cumulative return is reduced by approximately 10% for the shorter update intervals and by up to about 15–35% for the most restrictive 24 h planning case at the end of the mission, with peaks up to 40% during the most dynamic period. RCS1 DTE therefore introduces an operational dependence on telemetry-rate planning around the close-encounter interval.

Observation 4: RCS1 ISL provides sustained nominal data return over the full mission timeline.

For the ISL architecture, nominal performance is governed by the local RCS1–RAMSES geometry rather than by Earth–Apophis range. During the entire mission timeline, the inter-satellite distance remains close to the 15 km reference value, and free-space path loss does not introduce a large long-term trend. The total link losses instead vary within repeated local-geometry cycles.

The cumulative data return builds up steadily throughout the RCS1 mission, with the median reaching approximately 62 Gbit after three mission months and only moderate dispersion between favourable and adverse onboard antenna realisations.

From an architecture perspective, this sustained return profile is a major advantage because it supports both science data return and operational telemetry over the full RCS1 timeline.

Observation 5: RCS1 ISL retains partial communication robustness under degraded relative geometry.

The RCS1 ISL sensitivity analysis shows that the relay link remains feasible under degraded relative-geometry assumptions, although with reduced access to the higher data modes. When the RAMSES–RCS1 geometry remains partly constrained, the link tolerates moderate increases in inter-satellite distance retaining high coverage up to 1.5 times the nominal value. The fully unconstrained mutual-geometry case produces a stronger reduction, because neither spacecraft is assumed to maintain favourable ISL antenna orientation. Even in this case, distances above the nominal 15 km do not lead to an immediate loss of coverage with the link becoming progressively more opportunistic and limited to slowest data modes as range and angular uncertainty increase, especially beyond approximately 1.5 times the nominal spacecraft separation.

Degraded geometry and increased relative distance reduce science-return capability by shifting the link towards lower ISL modes. However, because the nominal RCS1 ISL case shows substantial spare data-return capacity relative to the reference requirement, continued access to lower ISL modes can still maintain a functional RCS1 relay capability under moderately degraded conditions.

6.3.2 RCS1 DTE and ISL evaluation

Table 6.3 applies the evaluation criteria defined in Table 6.1 to the RCS1 DTE and ISL architectures. The assessment draws on the Chapter 4 results and on the observations

identified in Section 6.3.1 to compare the two options at architecture level, focusing on their relative strengths, limitations and conditions of applicability.

Table 6.3 - RCS1 DTE and ISL architecture evaluation

Evaluation criteria	RCS1 DTE assessment	RCS1 ISL assessment	RCS1 architecture implication
Baseline communication performance	**Pass**. Nominal return ≈ 39.6 Gbit at 3 months, above R5 ≥ 12 Gbit. Link feasibility and carrier recovery are maintained.	**Pass**. Median nominal return ≈ 62.1 Gbit at 3 months, above R4 ≥ 9.2 Gbit/CubeSat.	Both options have large margin against the data-volume target.
Nominal performance dispersion	**Assumption-dependent pass**. Nominal return exceeds R5, but most useful return is generated around close encounter.	**Pass**. Small dispersion across favourable-adverse azimuth cases. Steady data-volume return over the entire mission.	DTE compliance depends on accepting a time-concentrated return profile.
Sensitivity to degraded or uncertain conditions	**Partial**. Degraded Earth pointing can reduce DTE to a few bit/s or to no-link intervals; the TM rate-planning limitation can reduce the data-volume restitution.	**Assumption-dependent pass**. Inter-satellite distance of up to 1.5 times the nominal value can be tolerated, at the cost of lower data-volume restitution.	DTE robustness is limited by pointing and rate planning. ISL robustness depends on accepting lower data-volumes.
CubeSat resource and accommodation impact	**Partial**. Mass/Volume allocations R1 met, but not the <1U goal; DC power allocation R2 is not met.	**Pass**. Mass/Volume allocations R1 met, including the <1U goal; DC power allocation R2 is met.	DTE is accommodation-feasible but more demanding and not compliant for DC power demand.
Operational dependence and architectural risk	**Assumption-dependent pass**. DTE depends on RAMSES-compatible G/S support with MUPA capability, contact availability and on TM-rate planning, MGA Earth pointing, which may interrupt asteroid-pointing operations.	**Assumption-dependent pass**. ISL return assumes RAMSES availability and ISD/geometry envelope; the link is feasible during asteroid-pointing operations.	DTE feasibility is tied to the availability of RAMSES ground station support. ISL feasibility depends on the assumed relative separation and relay geometry.

Table 6.3 shows that both RCS1 communication options are credible, but with different architectural implications. In particular, DTE provides independent Earth return and satisfies the R5 data-return reference with margin, but its return is concentrated around close encounter, exceeds the DC power allocation and depends strongly on Earth-pointing capability and ground-station availability as defined for RAMSES spacecraft. ISL provides a return profile well spread over the full RCS1 mission timeline, satisfies the R4 data-return reference with margin and has lower onboard resource demand, while remaining dependent on RCS1–RAMSES relay geometry.

Table 6.4 summarises the comparison using the qualitative marks defined in Table 6.2, providing a compact view of the relative strengths and limitations of the two architectures.

Table 6.4 - RCS1 comparative assessment summary

Evaluation criteria	DTE mark	ISL mark	Rationale
Baseline communication performance	**+**	**++**	Both pass the data-return checks with ISL providing higher spare capacity.
Nominal performance dispersion	**/**	**+**	DTE compliance depends on accepting time-concentrated return; ISL adverse nominal return remains well above the R4 criterion and spread over the entire mission timeline.
Sensitivity to degraded or uncertain conditions	**–**	**+**	DTE is sensitive to Earth-pointing and rate-planning degradation; ISL retains coverage under moderately degraded link conditions.
CubeSat resource and accommodation impact	**–**	**++**	DTE exceeds the preferred volume goal and not meeting the power allocation; ISL meets also the accommodation goal.
Operational dependence and architectural risk	**–**	**+**	DTE requires ground support, rate planning and Earth-pointing operations; ISL depends on relay geometry and can run during asteroid-pointing operations.

The RCS1 assessment identifies ISL as the preferred primary communication architecture. Its main advantages are sustained data return over the entire mission timeline, stronger lower-tail behaviour, reduced CubeSat resource demand and compatibility with asteroid-pointing operations.

DTE remains technically viable and relevant where independent Earth return is required, but it is less attractive as the primary RCS1 baseline because it depends more strongly on Earth-pointing, ground support, onboard resources and rate-planning capability.

This RCS1 conclusion is carried forward into the combined architecture assessment in Section 6.5.

6.4RCS2 architecture evaluation

This section evaluates the DTE and ISL communication architectures for the landed RAMSES CubeSat, RCS2, using the Chapter 5 results and the criteria defined in Section 6.2.

6.4.1 RCS2 architecture observations

The key RCS2 observations relevant to the architecture-level assessment are summarised below.

Observation 1: RCS2 DTE provides high return only in favourable realised cases.

For the RCS2 DTE architecture, the nominal analysis shows a wide dispersion of achievable cumulative data-volume across the analysed landed population. Favourable cases can reach approximately 10–11 Gbit at the 60 h operational time marker, but the median and lower-tail outcomes are much weaker, at approximately 0.59 Gbit and close to zero, respectively.

Without favourable conditions, there is therefore a credible risk that RCS2 would return very little data through the direct Earth link.

Observation 2: RCS2 DTE favourable cases depend on combined landed conditions.

High RCS2 DTE return is possible only when the realised landed state provides favourable Earth visibility and antenna geometry during the short useful DTE interval. The opportunity is therefore not broadly available across the landing population, but depends on a favourable combination of landing location, initial rotational phase and landed antenna heading.

Once touchdown has occurred, these conditions are fixed for the remainder of the surface mission. Unless the landing site and landed conditions can be constrained before touchdown, RCS2 DTE carries a substantial risk of low data return.

Observation 3: RCS2 DTE has limited robustness to surface perturbations and rate-planning constraints.

The sensitivity analysis in Section 5.3.1 shows that the broad RCS2 DTE performance distribution is further affected by adverse surface conditions and telemetry-rate planning constraints. Severe tilt/slope perturbations degrade the full distribution, including favourable cases, while restrictive rate-update intervals can produce substantial data-volume losses. RCS2 DTE therefore remains sensitive not only to the realised landing conditions, but also to local surface uncertainty and rate-planning assumptions, increasing the risk of limited data return through the direct Earth link.

Observation 4: RCS2 ISL provides consistent sustained data return across the surface-operation window.

For the RCS2 ISL architecture, nominal performance is driven mainly by local RCS2–RAMSES visibility and relative geometry, rather than by the changing Earth–Apophis range. Data return is accumulated during available relay intervals and continues to build from the beginning of operations throughout the RCS2 mission timeline.

Although the achievable data-volume varies across the population, the ISL architecture provides a relatively sustained return profile and reduces dependence on a single short communication opportunity. The median cumulative return reaches approximately 1.5 Gbit after 60 h, while the best realised case reaches approximately 3.4 Gbit. These values remain below the assumed RCS2 data-volume allocation, but the more sustained profile reduces the risk of very limited data return.

Observation 5: RCS2 ISL is landing-site dependent but robust to surface perturbations.
The RCS2 ISL sensitivity cases show that performance remains mainly governed by the realised landing location and the resulting visibility geometry with RAMSES. Once this landed geometry is defined, the analysed surface-condition variations do not substantially alter the architecture-level behaviour. Minimum-elevation variations, mild tilt/slope perturbations and landed-antenna performance produce limited changes, while severe tilt/slope perturbations introduce a moderate degradation.
The main result is therefore that RCS2 ISL remains landing-site dependent but is comparatively robust to the investigated surface condition variations.

6.4.2 RCS2 DTE vs ISL evaluation

Using the same evaluation framework applied to RCS1, Table 6.5 assesses the DTE and ISL options for the lander RCS2 case. The assessment draws on the Chapter 5 nominal and sensitivity results, together with the architecture observations identified in Section 6.4.1, and reflects the short surface-operation window, the dependence on realised landing conditions and the lower-tail behaviour of each option.

Table 6.5 - RCS2 DTE and ISL architecture evaluation

Evaluation criteria	RCS2 DTE assessment	RCS2 ISL assessment	RCS2 architecture implication
Baseline communication performance	**Fail.** Favourable cases reach 10–11 Gbit, but the 60 h median is only 0.59 Gbit, well below R5 ≥ 12 Gbit.	**Partial**. Median nominal return is about 1.5 Gbit after 60 h, below R4 ≥ 9.2 Gbit, but relatively consistently across the landed population.	Neither architecture satisfies the nominal data-volume criterion at median population level.
Nominal performance dispersion	**Fail**. DTE dispersion is large: adverse-case data return is approximately 15 times lower than in favourable cases.	**Pass**. Low-to-favourable-case dispersion within a factor of 3 and median nearly equally distant from the extremes.	DTE favourable cases remain isolated with weak median and lower-tail return. ISL provides stronger population-level consistency and better lower-tail return.
Sensitivity to degraded or uncertain conditions	**Fail**. DTE is highly sensitive to landing phase, local terrain and rate planning	**Pass**. Moderate geometry perturbations preserve the main ISL behaviour; Controlled landing conditions improve median by 1.5 times	DTE robustness is poor while ISL robustness remains stronger under the perturbations.
CubeSat resource and accommodation impact	**Partial**. Volume/Mass accommodation R1 met, but not the <1U goal; 12 W peak DC power and 720 Wh over 60 h.	**Pass**. ISL met volume/mass including the <1U goal. 5.5 W peak DC power and 330 Wh over 60 h.	DTE is accommodation-feasible but exceeds the preferred volume. ISL has lower DC power and energy demand.

Operational dependence and architectural risk	**Partial**. DTE return depends on favourable Earth visibility, RAMSES-compatible G/S support with MUPA capability and rapid rate planning during the short communication opportunities.	**Assumption-dependent pass**. ISL return assumes RAMSES availability and visibility during the short timeline.	DTE carries high realised-geometry and operational risks.

Table 6.5 shows that the RCS2 communication choice remains constrained under the current assumptions. Neither DTE nor ISL satisfies the adopted data-volume criterion at median population level, but each architecture retains a different advantage. DTE can provide high return in favourable realised cases, but its median and lower-tail behaviour are weak, and its robustness is strongly affected by landing conditions, surface perturbations and TM-rate planning. ISL remains below the R4 data-volume allocation, but provides a more consistent return profile across the population, lower onboard resource demand and stronger robustness across the analysed conditions.

Table 6.6 summarises the assessment using the qualitative marks defined in Table 6.2, providing a compact view of the relative strengths and limitations of the two architectures for RCS2.

Table 6.6 - RCS2 comparative assessment summary

Evaluation criteria	RCS2 DTE assessment	RCS2 ISL assessment	RCS2 architecture implication
Baseline communication performance	−−	−	DTE median return is far below R5 and even below the R4 reference; ISL provides the stronger baseline but remains below R4 at median level.
Nominal performance dispersion	−−	+	DTE favourable cases remain isolated relative to weak median and lower-tail return; ISL lower-percentile cases continue to accumulate data.
Sensitivity to degraded or uncertain conditions	−−	+	DTE is highly sensitive to landing geometry and short-window rate planning; ISL behaviour remains broadly preserved under moderate perturbations.
CubeSat resource and accommodation impact	−	+	DTE exceeds the volume goal and DC power envelope; ISL meets the accommodation goal and DC power check.
Operational dependence and architectural risk	−	+	DTE depends on favourable Earth visibility and rapid planning during a short opportunity; ISL depends on RAMSES relay availability and landed-window visibility.

The RCS2 assessment identifies ISL as the stronger landed communication option. Its main advantages are the more sustained return profile, better median and lower-tail behaviour, lower CubeSat onboard resource demand and greater robustness to the investigated surface-condition variations.

The conclusion remains conditional, because the median ISL data-volume remains below the R4 data-volume allocation under the adopted assumptions. The ISL option should therefore undergo further refinement if the R4 quantity allocation is retained. DTE remains capable of high return in favourable realised cases, but its weak median and lower-tail behaviour create a credible risk that RCS2 could return very little data through DTE.

6.5 Combined assessment and recommendations

Sections 6.3 and 6.4 identified ISL as the preferred communication option for both RAMSES CubeSat cases, although with different levels of confidence.

Since ISL is preferred for both CubeSat cases, and the adopted reference ISL terminal supports two parallel links, a common ISL architecture is the most logical baseline to carry forward at combined architecture level. While a common DTE architecture is not retained, a mixed DTE/ISL architecture could remain relevant in two specific cases: if RCS1 were required to operate more independently from RAMSES after deployment because of its longer mission duration, or if future RCS2 DTE assumptions make the favourable landed cases more representative and the higher onboard resource demand acceptable.

The recommended baseline is therefore a common ISL architecture, using RAMSES as the Earth-return path for both CubeSats and allowing the deployed spacecraft to share the same relay-based communication concept. This avoids the need to mature two distinct CubeSat communication implementations under the RAMSES schedule constraints.

Table 6.7 summarises the final communication-architecture recommendation for each CubeSat case and for the combined RAMSES CubeSat architecture.

Table 6.7 - Combined communication architecture recommendation

CubeSat case	Recommended baseline	Recommendation status	Remarks / follow-up
RCS1	**ISL**	Strong	Closes the data-return with large margin, provides sustained return during the entire RCS1 mission and reduces CubeSat onboard resource demand.
RCS2	**ISL**	Preferred	Provides more consistent and sustained data-volume return but does not fully comply with the target volume. Reduces CubeSat onboard resource demand.
Combined RAMSES architecture	**Common ISL**	Recommended baseline	Derived from the selected ISL baseline for both RCS1 and RCS2. RAMSES as the Earth-return path for both CubeSats avoiding duplicated communication implementations.

The resulting recommendation establishes a common ISL baseline for both deployed spacecraft, while identifying RCS2 data-volume closure against the R4 allocation as the remaining point requiring further consolidation. At system level, this makes RAMSES relay availability, ISL contact geometry and relay-operations planning central to the combined CubeSat communication concept.

6.6Chapter summary

This chapter evaluated the candidate DTE and ISL communication architectures for RCS1 and RCS2 using the performance results from Chapters 4 and 5 and the analysis requirements defined earlier in the thesis.

For RCS1, the assessment identified ISL as the preferred primary communication architecture. RCS1 ISL satisfies the adopted data-return reference with spare capacity, provides sustained return over the extended mission timeline and reduces CubeSat onboard resource demand relative to DTE. RCS1 DTE remains technically viable and provides independent Earth return, but its stronger dependence on Earth-pointing capability, ground support including availability of the MUPA technique, onboard DC power demand and telemetry-rate planning makes it less attractive as a baseline candidate.

For RCS2, the assessment also identified ISL as the stronger landed communication option. RCS2 DTE can provide high return in favourable realised cases, but its weak median and lower-tail performance creates a credible risk of very limited data return. RCS2 ISL provides a more consistent landed-return profile, lower onboard resource demand and greater robustness to the investigated surface-condition variations. However, under the adopted assumptions, the median ISL data-volume remains below the data-volume R4 target. The RCS2 case therefore requires further consolidation if full closure of the landed data-volume allocation remains mandatory.

The combined assessment therefore favours a common ISL architecture for the RAMSES CubeSats. This conclusion establishes a clear and efficient baseline for both deployed spacecraft, using RAMSES as the Earth-return path and avoiding distinct CubeSat communication implementations.

This chapter therefore completes Objective 5 by identifying the preferred communication baseline, the conditions attached to that recommendation and the remaining RCS2 data-volume point requiring further consolidation.

Chapter 7 Conclusion and Future Work

7.1Research summary

This research assessed candidate communication architectures for the two RAMSES CubeSats operating in the Apophis 2029 close Earth-encounter mission context. The study compared Direct-To-Earth (DTE) communication, where each CubeSat communicates directly with Earth ground stations, and Inter-Satellite-Link (ISL) communication, where each CubeSat transmits locally to the RAMSES spacecraft for onward Earth return.

The comparison was framed as a mission-specific architecture problem. RCS1 and RCS2 have different operational roles, timescales and geometry drivers. RCS1 is an orbiting CubeSat operating in proximity to the asteroid over an extended timeline of approximately three months, while RCS2 is a landed CubeSat operating over a short surface mission of less than three days around close encounter. These differences affect geometric visibility, operational constraints, onboard resource demands and the robustness drivers of each architecture.

A common modelling framework was developed and then tailored to the DTE and ISL cases for both CubeSats. For RCS1, the analysis represented the extended mission timeline, the Earth–Apophis range evolution relevant to DTE and the local RCS1–RAMSES geometry relevant to ISL. For RCS2, it represented the short landed-operation window, the fixed surface state after touchdown and the visibility conditions affecting direct Earth return and relay communication with RAMSES.

The framework was applied under nominal assumptions and selected sensitivity cases to characterise communication performance and robustness for each CubeSat.

7.2Main findings

Three main findings emerge from this work, ranging from the individual CubeSat cases to the combined RAMSES CubeSat communication architecture.

The first key finding is that communication-architecture suitability depends strongly on the CubeSat mission case. RCS1 and RCS2 operate over different timescales and under different geometric and operational constraints. The same DTE or ISL architecture therefore leads to different implications for data return, operational dependence, onboard resources and robustness.

For RCS1, both architectures are technically capable under the baseline assumptions. The DTE case returns approximately 39.6 Gbit over the three-month operational window, while the ISL case returns approximately 62.1 Gbit over the same period. Both are above the applicable data-return reference targets of 12 Gbit and 9.2 Gbit, but the two outcomes are not equivalent at architecture level. DTE return is concentrated around the close-encounter interval, depends on ground-station performance and availability, is sensitive to RCS1 Earth-pointing and telemetry-rate planning, and exceeds the target DC power allocation. ISL provides a more sustained return profile over the entire mission timeline, with higher cumulative return, substantially lower onboard resource demand and greater tolerance to the degraded-geometry cases. ISL was therefore identified as the stronger option for RCS1.
For RCS2, communication performance is driven by the short operation window and by the realised landing conditions. DTE can provide high return in favourable cases, reaching approximately 10–11 Gbit, but this upper-tail behaviour is not representative of the full population: the median return is only approximately 0.59 Gbit, with the lower tail remaining close to zero. In contrast, the ISL median return is approximately 1.5 Gbit, more than 2.5 times the DTE median. The lower tail is also stronger than for DTE: the ISL P10 return remains approximately 0.7–0.8 Gbit at 60 h, whereas the corresponding DTE lower-tail return remains close to zero. ISL therefore provides the stronger RCS2 option in terms of median and lower-tail data return, although it remains below the assumed 9.2 Gbit data-volume target. Further consolidation is required if full closure of the data-volume remains mandatory.

The second finding is that architecture evaluation must consider the full performance distribution, not only the best or nominal cases. This is shown most clearly by the RCS2 DTE case, where favourable realised cases can provide high data return, but median and lower-tail outcomes remain weak. For RCS2, with fixed touchdown geometry, high upper-tail return cannot be treated as representative unless the landing conditions that produce it can be made sufficiently probable.

The final finding is that the combined result favours a common ISL architecture for the RAMSES CubeSats. This recommendation is strongly supported by the RCS1 result and remains the preferred option for RCS2 under the adopted assumptions. At architecture level, the common ISL baseline provides a clear and efficient solution: RAMSES acts as the Earth-return path for both CubeSats, using a shared relay communication concept rather than requiring independent implementations for each deployed spacecraft.

7.3 Achievement of research objectives

The research objectives defined in Chapter 1 were addressed through the modelling framework, performance analyses, sensitivity assessments and architecture evaluation developed in the thesis. Their achievement is summarised below, linking each objective to the corresponding part of the research and to its contribution to the final architecture recommendation.

Objective 1 was to develop a consistent modelling framework for evaluating DTE and ISL communication architectures for the RAMSES CubeSats. This was achieved in Chapter 3 by defining the mission geometries, visibility modelling and link-performance calculations used throughout the analysis. The framework includes the Earth–Apophis trajectory, the asteroid TS-precession rotational model, local visibility constraints, and the distinct orbiting and landed geometries of RCS1 and RCS2. It also defines the DTE and ISL link-performance models, feasible-rate selection and cumulative data-volume calculation. The resulting framework provided a common basis for comparing DTE and ISL communication, while preserving the different geometry drivers of the RCS1 and RCS2 cases.

Objective 2 was to evaluate the nominal DTE and ISL communication performance of RCS1 and RCS2. This was addressed in Chapter 4 for RCS1 and Chapter 5 for RCS2. For RCS1, both DTE and ISL provide data return above the applicable target values, with ISL providing the higher cumulative return over the three-month operational window. For RCS2, the nominal analysis showed a stronger dependence on the realised landed state: DTE can provide high return only under favourable realised conditions, while ISL provides a more consistent return profile across the analysed landed cases, although its median value remains below the assumed data-volume target.

Objective 3 was to quantify sensitivity and robustness under representative geometric, operational and design variations. The relevant modelling capability was defined in Chapter 3 and applied in Chapters 4 and 5. For RCS1 DTE, the results show sensitivity to Earth-pointing capability and telemetry-rate planning. RCS1 ISL is instead moderately robust to degraded inter-satellite distance and viewing angles. For RCS2 DTE, sensitivity is dominated by the realised landed conditions, with additional sensitivity to asteroid surface-property perturbations. RCS2 ISL also depends on landed and surface conditions but shows lower sensitivity than DTE to the investigated surface-condition variations.

Objective 4 was to assess statistical variation across representative scenario populations. This was addressed primarily through the RCS2 landed analysis in Chapter 5, where Monte Carlo landing populations were used to evaluate percentile behaviour, lower-tail outcomes and the representativeness of favourable cases. It was also reflected in the RCS1 sensitivity cases, where favourable and adverse conditions were compared to assess the spread of achievable performance.

Objective 5 was to synthesise the nominal, statistical and sensitivity results into an architecture-level evaluation and recommendation. This was completed in Chapter 6 through the evaluation framework, the separate RCS1 and RCS2 architecture assessments, and the combined RAMSES-level recommendation. The final assessment identified ISL as the preferred communication-architecture solution for each RAMSES CubeSat and as the common RAMSES CubeSat communication baseline, while the RCS2 data-volume target remains the main point requiring further consolidation.

7.4Limitations and future work

This research was intentionally broad in coverage because it compared two communication architectures across two different CubeSat mission cases, under complex representative mission, geometry and design assumptions. This breadth supported the architecture-level assessment but necessarily limited the level of detail applied to individual cases. Several mission, operations and landed-environment assumptions were therefore treated at architecture-assessment level, rather than modelled in full detail.

By identifying a common ISL architecture as the preferred baseline, the thesis reduces the set of viable options for the next design step. Future work can therefore focus less on broad trade-space comparison and more on detailed refinement of the selected ISL architecture.

The first area for future work is to update the ISL analysis as the RAMSES mission definition matures. This should include updated RAMSES, RCS1 and RCS2 trajectories, refined operations timelines including combined ISL communications, and an improved asteroid rotational model capturing possible spin-state evolution due to the close encounter.

A second area is to extend the ISL modelling over the full CubeSat operational sequences. For RCS1, this would include deployment, commissioning, orbit insertion, data generation and onboard storage. For RCS2, the sequence is shorter but more time-critical, because deployment, landing, transition to operations and transmission occur within the compressed close-encounter timeline.

The third area concerns the landed RCS2 ISL case. Future work should focus on a shortlisted set of landing cases linked to RAMSES project-led site selection, including scientific interest and technical feasibility. A loose-control model, in which touchdown location and conditions are dispersed around targeted values, should also be included. This should be complemented by a more detailed near-horizon propagation assessment, including local terrain effects and ISL terminal FHSS robustness. This would consolidate whether the RCS2 ISL baseline can achieve the targeted data-volume under more representative mission conditions.

Finally, the sensitivity analyses could be extended into a more complete Global Sensitivity Analysis (GSA). The local and scenario-based sensitivities used in this thesis identify the main robustness drivers, but a full GSA would quantify the relative contribution of parameter uncertainties across the complete model. This would provide a stronger basis for prioritising future design effort.

7.5Final conclusion

This thesis has developed and applied a mission-specific communication assessment for the RAMSES CubeSats, comparing DTE and ISL architectures across the distinct orbiting and landed mission cases. The work shows that, for small-body CubeSat missions, communication-architecture selection is strongly shaped by geometry and timing: changes in relative geometry, landing location or asteroid rotational phase can turn a favourable communication opportunity into a weak link condition. It also identifies the key strengths, limitations and sensitivity drivers of both communication solutions.

Table 7.1 summarises the achievement of the research objectives and the principal outcomes supporting the final conclusions.

Table 7.1 - Summary of research-objective achievement

Objective	Status	Key outcomes / quantities
#1 - Modelling Framework	**Met**	Common DTE/ISL modelling framework established.
#2 - Nominal performance	**Met**	RCS1: 39.6 / 62.1 Gbit DTE/ISL; RCS2 median: 0.59 / 1.5 Gbit
#3 - Sensitivity and robustness	**Met**	Main robustness drivers quantified; RCS1 ISL robust to ≈1.5× nominal ISD
#4 - Statistical characterisation	**Met**	RCS2 DTE: ≈10–11 Gbit favourable vs 0.59 Gbit median; ISL P10 ≈0.7–0.8 Gbit
#5 - Architecture evaluation	**Met**	Common ISL recommended

The thesis demonstrates that an ISL-based architecture is feasible and preferred for the RAMSES CubeSats, both individually and as a common baseline. This conclusion is strongest for RCS1, where ISL meets the data-return target with margin and reduces onboard resource demand. For RCS2, ISL is also preferred relative to DTE, but the achievable data-volume remains the main area requiring refinement.

References

[1] C. J. Benson *et al.*, "Spin state evolution of (99942) Apophis during its 2029 Earth encounter," *Icarus*, vol. 394, Art. no. 115452, 2023.

[2] P. Michel, M. Lazzarin, M. Küppers, and S. Green, "Science objectives of RAMSES: ESA's Rapid Apophis Mission for SpacE Safety," in *Apophis T-4 Years*, LPI Contribution No. 3083, 2025.

[3] P. Martino, M. Küppers, M. Lazzarin, and P. Michel, "Operational design of RAMSES: ESA's Rapid Apophis Mission for SpacE Safety," in *Apophis T-4 Years*, LPI Contribution No. 3083, 2025.

[4] OHB System and GMV, "Rapid Apophis Mission for Security and Safety (RAMSES): Executive Summary, Pre-Phase A/Phase A," Tech. Rep., Apr. 24, 2024.

[5] P. Martino *et al.*, "The ESA RAMSES mission concept: A rendezvous with the asteroid Apophis during its close encounter with Earth in 2029," in *Proc. 75th Int. Astronautical Congr. (IAC)*, Milan, Italy, Oct. 2024, Paper IAC-24,A3,4,1.

[6] P. Fortescue, G. Swinerd, and J. Stark, Eds., *Spacecraft Systems Engineering*, 4th ed. Chichester, U.K.: Wiley, 2011.

[7] P. Fortescue, J. Stark, and G. Swinerd, Eds., *Space Mission Engineering: The New SMAD*. Hawthorne, CA, USA: Microcosm Press, 2011.

[8] J. R. Lesh, *Practical Introduction to the Design of Telecommunications Systems for Deep Space Missions*. Pasadena, CA, USA: Jet Propulsion Laboratory, DESCANSO Design and Performance Summary Series, 2003.

[9] J. H. Yuen, *Deep Space Telecommunications Systems Engineering*. Pasadena, CA, USA: Jet Propulsion Laboratory, DESCANSO Design and Performance Summary Series, 2013.

[10] European Space Agency, "RAMSES Consolidated Report on Mission Analysis (CReMA)," ESA/ESOC, Tech. Rep. ESA-RAM-OPSOPH-RP-05001, Issue 1.0, Nov. 27, 2024.

[11] P. Pravec *et al.*, "The tumbling spin state of (99942) Apophis," *Icarus*, vol. 233, pp. 48–60, 2014.

[12] H.-J. Lee *et al.*, "Refinement of the convex shape model and tumbling spin state of (99942) Apophis using the 2020–2021 apparition data," *Astron. Astrophys.*, vol. 661, Art. no. L3, 2022.

[13] M. Brozović *et al.*, "Goldstone and Arecibo radar observations of (99942) Apophis in 2012–2013," *Icarus*, vol. 300, pp. 115–128, 2018.

[14] ESA/ESOC (OPS-ONX), "RAMSES—Space to Ground Interface Control Document (SG-ICD), Volume 1: Radio Frequency (RF)," ESA/ESOC, Tech. Rep. ESA-RAM-ESC-IF-00001, Issue 1.0, 2025.

[15] European Space Agency, "RAMSES System Requirements Document," ESA, Tech. Rep. ESA-TECSH-RS-2023-000382, Rev. 1, 2023.

[16] F. Carden, R. Henry, and R. Jedlicka, *Telemetry Systems Engineering*. Norwood, MA, USA: Artech House, 2002.

[17] European Cooperation for Space Standardization, "Space Engineering—Radio Frequency and Modulation," ESA Requirements and Standards Division, ESTEC, Noordwijk, The Netherlands, Standard ECSS-E-ST-50-05C Rev. 2, 2011.

[18] L. J. Ippolito, Jr., *Satellite Communications Systems Engineering: Atmospheric Effects, Satellite Link Design, and System Performance*, 2nd ed. Chichester, U.K.: Wiley, 2017.

[19] J. S. Border, *Radiometric Tracking Techniques for Deep-Space Navigation*, vol. 1. Pasadena, CA, USA: Jet Propulsion Laboratory, DESCANSO Design and Performance Summary Series, Oct. 2000.

[20] Consultative Committee for Space Data Systems, "TM Synchronization and Channel Coding," CCSDS, Washington, DC, USA, Recommendation for Space Data System Standards 131.0-B-3, Issue 3, 2017.

[21] Jet Propulsion Laboratory, "DSN Telecommunications Link Design Handbook 810-005, Module 301: Coverage and Geometry, Rev. M," Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA, Rep., 2020.

[22] ESA/ESOC, "Use of ESTRACK Downlink Performance Models in SGICD," ESA/ESOC, Tech. Rep. ESA-OPS-GSY TN, OSP-GS/2839/ML, Issue 2.0, 2021.

[23] T. Delamotte *et al.*, "An RF-based solution for standardized LEO Inter-Satellite-Links," in *Proc. 12th Advanced Satellite Multimedia Systems Conf. and 18th Signal Processing for Space Communications Workshop (ASMS/SPSC)*, Barcelona, Spain, Sep. 2025.

[24] Consultative Committee for Space Data Systems (CCSDS), "Proximity-1 Space Link Protocol—Rationale, Architecture, and Scenarios," CCSDS, Washington, DC, USA, Informational Report, Green Book, Issue 1, CCSDS 210.0-G-1, Aug. 2007.

[25] NASA, "SpaceX Demonstration Mission: Press Kit," National Aeronautics and Space Administration, Washington, DC, USA, Tech. Rep., May 2012.

[26] A. Babuscia *et al.*, "Telecommunications systems testing and ground-compatibility verification for EM-1 CubeSat missions," in *Proc. 70th Int. Astronautical Congr. (IAC)*, Washington, DC, USA, Oct. 21–25, 2019, Paper IAC-19-B4.8.5.x50437.

[27] M. M. Kobayashi, M. Shihabi, and J. Taylor, *Mars Cube One Telecommunications Subsystem Design*. Pasadena, CA, USA: Jet Propulsion Laboratory, DESCANSO Design and Performance Summary Series, Sep. 2021.

[28] S. Simonetti *et al.*, "LICIACube on DART mission: An asteroid impact captured by Italian small satellite technology," in *Proc. 34th Annu. Small Satellite Conf.*, Logan, UT, USA, 2020.

[29] C. B. Haskins and C. C. DeBoy, "Deep-space transceivers—An innovative approach to spacecraft communications," *Proc. IEEE*, vol. 95, no. 10, pp. 2009–2018, Oct. 2007.

[30] G. Cucinella *et al.*, "CubeSAT Deep Space X-Band TT&C Transponder (C-DST)," presented at L'impegno Italiano nel settore dei CubeSat: tecnologie e missioni future, Jul. 2–4, 2024.

[31] IMT, "C-DST CubeSat Deep Space X-Band TT&C Transponder Flyer," IMT, Tech. Data Sheet, version 2.0, 2024.

[32] J.-L. Issler and C. Dudal, "Inter-Satellite-Links for Rosetta/Philae, Hayabusa2/MASCOT, and next-generation miniaturized ISL in S- and Ka-band," presented at an ESA/CNES technical briefing, Jun. 2017.

[33] P. Michel *et al.*, "The ESA Hera mission: Detailed characterization of the DART impact outcome and of the binary asteroid (65803) Didymos," *Planet. Sci. J.*, vol. 3, no. 7, Art. no. 160, Jul. 2022, doi: 10.3847/PSJ/ac6f52.

[34] M. Santos *et al.*, "Adaptive signal processing and SDR techniques in GAMALINK for power-efficient deep space Inter-Satellite-Links," in *Proc. TT&C 2025*, ESA/ESOC, Darmstadt, Germany, Nov. 25–28, 2025.

[35] M. Santos *et al.*, "GAMALINK: An enabler for precision formation flying for ESA Proba-3 mission," in *Proc. TT&C 2025*, ESA/ESOC, Darmstadt, Germany, Nov. 25–28, 2025.

[36] NASA/JPL Solar System Dynamics Group, "HORIZONS Web-Interface and On-Line Ephemeris System," Jet Propulsion Laboratory, California Institute of Technology, 2026. [Online]. Available: https://ssd.jpl.nasa.gov/horizons/. Accessed: Jul. 6, 2026.

[37] J. S. Seybold, *Introduction to RF Propagation*. Hoboken, NJ, USA: Wiley-Interscience, 2005.

[38] M. Sanchez Net, "Simulation of multipath reflections from planetary bodies: Theory and application to the lunar south pole," *IPN Prog. Rep.*, vol. 42-226, 2021.

[39] ESA/ESTEC, "Link Budget Tool Technical Description," ESA/ESTEC, Tech. Rep. ESA-TEC-ETC TN, ESA-TECETC-TN-2005-000937, Issue 1.0, 2005.

[40] Jet Propulsion Laboratory, "Atmospheric and Environmental Effects, DSN Telecommunications Link Design Handbook, 810-005, Module 105, Rev. E," Jet Propulsion Laboratory, California Institute of Technology, Pasadena, CA, USA, Rep., 2015.

[41] W. J. Conover, *Practical Nonparametric Statistics*, 3rd ed. New York, NY, USA: John Wiley & Sons, 1999.

[42] OHB Italia S.p.A., "Downlink Performance Analysis," OHB Italia S.p.A., Milan, Italy, Tech. Rep. RAMSES-OHBI-COM-RP-0086, Issue 04, Feb. 11, 2026.

[43] ISISPACE B.V., "Deep Space Deployer Executive Summary," ISISPACE B.V., Delft, The Netherlands, Tech. Rep. ISIS-6U_DSD-RPT-0005, Rev. 01_00, 2020.

[44] GomSpace A/S, "NanoPower MSP Datasheet," GomSpace A/S, Tech. Rep. GS-DS-NANOPOWER-MSP_1016340, Rev. 3.2, 2026.

[45] TEKEVER, "ISL User Manual," TEKEVER, Tech. Rep. TEK-ISL-MA-0001, Issue 1.0, 2026.

[46] C. Cappellin, M. M. Bilgic, J. R. de Lasson, and O. Borries, Antennas on CubeSat platforms: Accurate RF predictions, TICRA, Copenhagen, Denmark, 2020.

[47] O. Borries, M. F. Palvig, and T. Rubæk, High-frequency CubeSat platform scattering using higher-order method of moments, TICRA, Copenhagen, Denmark, 2019.

[48] Anywaves, "S-Band TT&C Antenna Datasheet," Anywaves, Toulouse, France, Tech. Data Sheet DTS-V1-ASA-052026, May 2026.

[49] D. K. Misra, *Radio-Frequency and Microwave Communication Circuits: Analysis and Design*, 2nd ed. Hoboken, NJ, USA: Wiley, 2004.

[50] European Space Agency, "ESA awards contracts for Ramses mission to Apophis," ESA Space Safety News, Feb. 10, 2026. [Online]. Available: https://www.esa.int/Space_Safety/Planetary_Defence/ESA_awards_contracts_for_Ramses_mission_to_Apophis. Accessed: Jul. 6, 2026.

[51] European Space Agency, "First time controlling two spacecraft with one dish," *ESA Enabling & Support*, Jan. 21, 2025. [Online]. Available: https://www.esa.int/Enabling_Support/Operations/First_time_controlling_two_spacecraft_with_one_dish. Accessed: Jul. 6, 2026.

[52] D. S. Abraham *et al.*, "Antenna beam-sharing: Progress toward multiple uplinks per antenna," in *Proc. 17th Int. Conf. Space Operations (SpaceOps 2023)*, Dubai, United Arab Emirates, Mar. 6–10, 2023, Paper SpaceOps-2023, ID 139.

[53] R. J. Serfling, *Approximation Theorems of Mathematical Statistics*. New York, NY, USA: John Wiley & Sons, 1980.

[54] P. J. Huber and E. M. Ronchetti, *Robust Statistics*, 2nd ed. Hoboken, NJ, USA: John Wiley & Sons, 2009.

[55] T. Rainforth, R. Cornish, H. Yang, A. Warrington, and F. Wood, "On nesting Monte Carlo estimators," 2017, arXiv:1709.06181.

[56] M. Giles and O. Sheridan-Methven, "Analysis of nested multilevel Monte Carlo using approximate normal random variables," *SIAM/ASA J. Uncertain. Quantif.*, vol. 10, no. 1, pp. 200–226, Feb. 2022, doi: 10.1137/21M1399385.

Appendix

Appendix I Baseline assumptions and analysis requirements

This appendix lists the complete numbered set of baseline assumptions and analysis requirements used to define the RAMSES CubeSat communication-architecture assessment. The assumptions are grouped according to their role in the analysis: mission and geometry, communication system, and analysis requirements.

Table Ap.1 defines the mission and geometric baseline, including the RAMSES timeline, Earth–Apophis trajectory, CubeSat operational cases and landed-geometry assumptions. Table Ap.2 defines the communication-system baseline, including the DTE and ISL link assumptions, ground-segment representation, link-performance criteria, data-mode definitions and antenna baselines. Table Ap.3 defines the analysis requirements and implementation constraints, including resource, accommodation, technology-maturity and data-return assumptions relevant to the architecture assessment.

Prefixes distinguish the three groups of assumptions, M, C and R. The final column of each table identifies the source of the corresponding assumption or the type of assumption introduced for the analysis.

Table Ap.1 - Mission and geometric baseline assumptions

ID	Assumption / input	Definition adopted in the thesis	Value or profile	Source / assumption type
M1	**Apophis closest approach**	Reference event for the RAMSES mission and CubeSat analyses.	CA = 2029-04-13, 21:46:00 UTC	Ephemeris / mission input [10], [36]
M2	**RAMSES mission timeline**	Launch and arrival dates used to frame the time-critical mission context.	Launch in April 2028; arrival at Apophis in February 2029, approx. two months before CA	Public RAMSES mission description [3], [5], [10]
M3	**Earth–Apophis trajectory**	Earth position relative to Apophis used for DTE range and direction.	JPL HORIZONS ephemerides; minimum range approximately 32,000 km	Ephemeris-derived input [10], [36]
M4	**Ground Station viewing geometry**	Ground Stations definition for DTE viewing-angle evaluation.	Three equally spaced equatorial Ground Station positions	Thesis modelling assumption
M5	**CubeSat cruise phase**	CubeSats are carried by RAMSES during interplanetary cruise.	No CubeSat communication during cruise	RAMSES mission concept / thesis scope boundary [5], [4], [15]
M6	**RCS1 operational case**	RCS1 is the orbiting CubeSat operating in proximity to Apophis; the analysis starts after orbit insertion.	Circular Apophis-centred orbit of 1.2 km radius; deployment, transfer and commissioning not analysed	RAMSES RCS1 concept [3], [50] / thesis analysis assumption
M7	**RCS1 mission window**	RCS1 DTE/ISL performance is assessed over the extended	100 days from 15 March 2029; approximately CA -1 month to	Thesis analysis assumption based on

		operational interval.	CA +2 months	[3]
M8	**RCS1 DTE geometry**	RCS1-to-Ground Station range is assumed equal to the Apophis-to-Ground Station range; a simplified asteroid-occultation model is necessary	RCS1 orbital radius neglected for range calculation; asteroid occultation to be considered and treated with a simplified conservative model	Thesis modelling assumption
M9	**RCS1 ISL geometry**	RCS1–RAMSES local geometry used for ISL analysis.	Relative trajectory / link-geometry profile used in simulations	Mission-analysis input [10]
M10	**RCS2 operational case**	RCS2 is the landed CubeSat operating from the Apophis surface; the analysis starts after landing.	Lander/surface case; descent, touchdown and commissioning not analysed	RAMSES RCS2 concept [3], [50] / thesis scope assumption
M11	**RCS2 surface-operation windows**	RCS2 is analysed from landing through early and extended post-CA operation.	Landing at CA−20 h; 20 h covers pre-CA return, 40 h includes early post-CA return, 60 h represents extended operation	RAMSES RCS2 concept / thesis analysis assumption [3]
M12	**RCS2 DTE geometry**	Landed CubeSat-to-Earth geometry used for DTE analysis.	Earth direction varies with Apophis rotation, Earth–Apophis motion and active Ground Station geometry	Thesis modelling assumption
M13	**RCS2 ISL geometry**	Landed CubeSat-to-RAMSES geometry used for ISL analysis.	RAMSES trajectory relative to the rotating Apophis surface	Mission-analysis input / thesis modelling assumption
M14	**RCS2 landed geometry cases**	Statistical surface assessment; RCS2 assumed to rest on the surface without burial or penetration.	Monte Carlo landing-site population; initial rotational phase phi0 in [0 deg, 360 deg], sampled at 5 deg increments	Thesis analysis assumption
M15	**Apophis shape and rotation**	Shape and rotation model used for landed-geometry analysis.	Triaxial ellipsoid, approximately 375 m × 170 m × 170 m; Tilt–Spin–Precession model: Tspin = 30.56 h, α = 30 deg, Tprec = 264.0 h	Literature-derived input [1], [11], [12], [13] / thesis modelling assumption
M16	**Minimum-elevation condition**	Near-horizon visibility constraint for landed links.	Baseline 10° minimum elevation	Thesis modelling assumption [13], [37], [38]

Table Ap.2 - Communication-system baseline assumptions and reference data

ID	Assumption / input	Definition adopted in the thesis	Value or profile	Source / assumption type
C1	**Candidate communication architectures**	Both DTE and ISL are assessed for RCS1 and RCS2 as telemetry-return architectures.	DTE: CubeSat → Ground Station; ISL: CubeSat → RAMSES relay link.	Thesis architecture assumption
C2	**DTE link-performance framework**	RCS DTE uses the RAMSES-compatible telemetry framework and is assessed using an Eb/No-based feasibility criterion.	X-band downlink, residual-carrier modulation assumptions, Turbo coding and Eb/No thresholds.	Thesis analysis assumption; values RAMSES-derived [14], [17], [20], [39]
C3	**DTE telemetry-rate set**	RCS DTE mode selection uses the RAMSES-defined TM rate set.	25 TM rates from 8 bps to 308 kbps; Turbo 1/2 and Turbo 1/4 coding options.	RAMSES-derived reference data [14]

C4	**DTE Ground Station support**	RCS DTE uses the RAMSES Ground Station performance and contact-availability representation, assuming compatibility with simultaneous RAMSES/RCS support.	Time-dependent G/T profile and contact-time / duty-cycle profile; frequency separation and multi-spacecraft support techniques (e.g. MUPA) assumed where required.	Thesis analysis assumption; profiles RAMSES-derived [14], [22], [51], [52]
C5	**DTE feasibility criterion**	A telemetry mode is feasible when the available link performance satisfies the required threshold and margin.	Eb/No-based margin criterion applied to each available TM mode.	Thesis modelling assumption [17], [20], [39]
C6	**DTE Telemetry rate selection**	The selected DTE rate is the highest feasible telemetry mode at each analysis epoch.	Selection from the RAMSES-compatible TM rate set.	Thesis modelling assumption
C7	**DTE carrier recovery**	DTE feasibility includes carrier-recovery assessment where relevant.	Residual-carrier C/No or equivalent margin evaluated for selected cases.	Thesis modelling assumption [14], [17], [19], [39]
C8	**RCS1 DTE antenna baseline**	RCS1 DTE uses an MGA for far-range communication and LGA coverage for close-encounter or degraded-link cases.	MGA boresight gain ≈20 dBi; LGA gain and axial-ratio behaviour taken from antenna-pattern inputs.	Thesis analysis assumption / antenna-pattern input
C9	**RCS2 DTE antenna baseline**	RCS2 DTE uses a surface-fixed LGA configuration after landing.	LGA gain and axial-ratio behaviour taken from antenna-pattern inputs and applied according to landed viewing geometry.	Thesis analysis assumption / antenna-pattern input
C10	**ISL link-performance framework**	RCS ISL feasibility is assessed by comparing received power with mode-specific receiver sensitivity thresholds.	Received-power threshold criterion applied to each ISL data mode.	Adopted reference / supplier-type data [34], [35], [45]
C11	**ISL bandwidth allocation**	ISL data rates depend on the bandwidth allocation assigned to the CubeSat link.	Low-priority and high-priority rate columns represent bandwidth-sharing cases between CubeSat links.	Thesis architecture assumption [34], [35], [45]
C12	**ISL data-mode set**	ISL mode selection uses the available ISL data modes.	Six data modes; high-priority rates from 0.52 kbps to 20.48 kbps; carrier frequency 2.43 GHz.	Adopted reference / supplier-type data [34], [35], [45]
C13	**ISL data-mode selection**	ISL is treated as an adaptive modem link with autonomous mode selection based on link feasibility.	Highest feasible ISL mode selected at each simulation step; data transfer is assumed whenever the ISL link can be established.	Supplier-published capability; applied as thesis modelling assumption [34], [35], [45]
C14	**ISL relay boundary**	RAMSES-to-Earth communication is assumed not to limit the RCS ISL data-return capability.	Only the RCS-to-RAMSES link is modelled; the RAMSES-to-Earth leg is not modelled separately.	Thesis architecture assumption

Table Ap.3 - Analysis requirements and implementation constraints

ID	Assumption / input	Definition adopted in the thesis	Value or profile	Source / assumption type
R1	**CubeSat accommodation**	The communication system must be compatible with the RAMSES CubeSat platform while preserving payload volume.	6U-class CubeSat; at least 3U retained for scientific payload; communication-system goal <1U excluding antenna. 1U maximum mass 1.33 kg.	RAMSES CubeSat concept / thesis analysis assumption [15], [43]
R2	**Electrical interface and power**	Communication-system power demand is assessed in relation to CubeSat-class resources and mission-specific energy constraints.	28 V input; power demand assessed against representative 6U CubeSat resources and mission-specific energy constraints.	CubeSat reference design / thesis analysis assumption [6], [43]
R3	**Technology maturity and environment**	Candidate solutions must be compatible with the RAMSES schedule and deep-space environment.	Existing or near-existing designs only; minor modifications acceptable.	RAMSES schedule constraint / thesis architecture assumption [3], [5]
R4	**ISL data allocation**	ISL data return is assessed in relation to the RAMSES relay allocation for each CubeSat.	At least 9.2 Gb per CubeSat over the mission duration, including 6 MB/day platform housekeeping and remaining capacity for science data.	RAMSES-derived input / thesis analysis assumption [15]
R5	**DTE data-return reference**	DTE is assessed as a direct-return architecture not capped by the RAMSES ISL relay allocation.	Direct-return reference greater than the ISL allocation; Goal ≥12 Gb.	RAMSES-derived input / thesis analysis assumption [15]
R6	**Architecture evaluation criteria**	Architecture comparison considers communication performance together with resource, implementation and operational implications.	Link feasibility, coverage, selected rate/mode, cumulative data-volume, robustness, onboard resources, architectural dependence and operational flexibility.	Thesis analysis assumption

Appendix II RCS2 landing-site population

This appendix lists the 200 RCS2 landing locations introduced in Chapter 3 and used for the landed-population analyses in Chapter 5. The same surface-location population is used for both the DTE and ISL analyses. Table Ap.4 reports the surface coordinates for each location in the Apophis body-fixed frame, in metres.

Table Ap.4 – RCS2 landing-site body-frame coordinates

Landing site ID	X [m]	Y [m]	Z [m]
1	11.2040779	78.14668557	-31.09581381
2	23.84784596	-77.75244872	30.63523302
3	-80.5720304	65.71426002	-36.88867442
4	-63.11887594	-44.84158075	-62.24016358
5	-18.00822312	-34.14433171	72.84672281
6	0.392786321	35.63351431	72.63065646
7	-24.66820586	-43.08267531	-68.13248312
8	-56.31239426	21.54610173	73.45603372
9	-26.61897477	-77.67360264	-30.38455615
10	-6.980155091	-68.8874235	-46.76839702
11	63.21452666	35.58183668	-67.31451246
12	-51.66914722	81.30437737	-6.722424401
13	-6.635456441	64.08341858	52.47867624
14	19.63842827	52.76661863	-62.1409462
15	101.761273	66.08280247	24.39838051
16	59.890047	-77.46886414	20.32564506
17	97.26267988	-44.29274867	53.7881487
18	27.74459758	-33.30190348	72.620069
19	-40.84607036	-12.19767573	77.17654038
20	-119.8520945	3.775594558	-60.8378009
21	-125.277861	39.1628698	45.89723766
22	-24.20775659	43.45573917	67.95326127
23	-3.556853034	79.96944814	-27.06952885
24	-4.50489782	80.89967358	-24.46984015
25	157.2579479	4.964119517	41.87724838
26	-8.295564382	40.2069633	-70.39265615
27	-53.73288226	-48.48062047	-61.46630416
28	-29.57903907	83.16026787	-10.50884579
29	44.51285794	61.4307794	-51.83294774
30	-173.4466998	28.21675847	-8.314547039
31	-2.529023116	76.89875741	-34.06786994
32	-13.48888694	73.03443042	-40.50954917
33	-95.62130856	10.39763723	-67.78223761

34	-109.3953196	3.004068798	-64.45262171
35	35.8884643	-57.76109496	-56.60187937
36	45.4715454	-53.47375015	-58.99508929
37	0.275796358	-82.61161712	18.83069652
38	18.45336462	70.31910395	44.22865954
39	0.211757903	45.57842288	-67.52636452
40	-82.47351504	-75.57416515	8.29347802
41	-46.29714632	-52.2000524	-59.87887886
42	0.902963683	78.29145622	31.14818888
43	54.58854342	30.40080775	70.88079394
44	-117.3505249	56.40139349	31.73248338
45	-25.37021058	60.59601633	-55.01856243
46	95.69338193	5.846765921	-68.24469261
47	18.23787866	-19.42120219	77.48346278
48	-128.34085	49.14919432	-34.35230119
49	-4.740814591	28.82125318	75.23286189
50	134.0042898	51.31317864	-26.63961088
51	37.51667257	66.8713819	-46.64372848
52	12.66293658	-66.36364031	-49.68678098
53	-130.8407772	-44.38477559	-38.12699142
54	-142.7797381	-6.53707834	-50.49751325
55	-15.6300072	74.94860023	-37.12741925
56	112.9105897	67.14076755	-4.781251018
57	28.02381103	-25.22970456	75.42738589
58	49.65625232	8.629024388	76.63519875
59	62.7966635	60.11227128	49.6156856
60	57.36449286	-43.62497638	-64.02205826
61	-44.32917476	78.9072874	-22.74079115
62	-56.35635256	80.8883144	3.206824518
63	172.7689846	9.773778818	-27.08617241
64	35.26585435	-34.10728413	-71.67260277
65	-180.0866735	18.06858101	-6.800013051
66	6.071628285	-29.2526066	75.06732175
67	-20.37176767	-73.73158728	-38.81760709
68	-42.33645326	-77.13975117	-28.17378401
69	9.01252013	-78.78374583	29.77711071
70	67.79533093	-9.634785483	73.88024196
71	113.358143	33.79112871	-54.64071843
72	65.2020102	-76.44760666	20.6909213
73	-21.21997956	31.5395666	-73.72000863
74	0.503581577	-6.212592381	-79.78573472
75	10.81201572	7.135622344	79.58038058
76	9.084981473	-55.0884906	60.79762821
77	-79.87153728	11.5424364	-71.33750435
78	-159.305352	-34.53570097	-24.45016977

79	33.89930555	55.73408243	-58.59623152
80	-1.92140419	37.48801356	71.79436681
81	127.5458607	-10.5971251	57.08286658
82	-12.56613338	-73.79926258	39.31971329
83	-83.44158744	-9.156595933	-70.87847757
84	41.7280055	-40.36298495	68.05331566
85	-10.23530546	-39.38534378	70.75543638
86	178.3572457	-22.53385444	1.249748263
87	3.919249993	-82.34322505	19.77256171
88	74.54706358	65.62978855	-39.31116937
89	54.34164265	24.63239926	72.87195064
90	174.59342	24.36768739	13.18293242
91	26.60655308	54.40770238	-60.37748132
92	28.36159073	48.54285004	64.51546549
93	-145.8501102	5.069635235	-48.98334959
94	180.9108981	-0.89819569	16.70569564
95	-28.01348185	80.44735525	-22.81375397
96	-18.11213594	-51.18137416	-63.38957832
97	2.713209329	-52.47861118	62.92129926
98	-168.5226457	26.54903615	-21.56195059
99	62.23608136	78.3012605	-15.64282517
100	7.581028749	-83.49182408	-14.64074953
101	-24.90840224	-23.3702241	-76.15890616
102	12.77769563	66.36070941	-49.6847546
103	35.28070446	-4.192942579	-78.43255253
104	6.086974407	27.48984885	-75.65493885
105	-61.01470615	-73.31714535	30.69596461
106	-17.32818535	25.31879603	-76.00005253
107	-71.37961603	30.88315694	-67.8408095
108	-114.1594272	56.61777242	-33.51770432
109	10.96137769	59.67153376	56.77520705
110	-31.51014343	-83.57182922	-5.252699841
111	-45.66390118	33.3149131	-70.90080641
112	-96.06193649	-72.63493991	1.000084827
113	23.87509666	-16.14206856	77.86266333
114	20.92084173	-53.23511855	61.70720911
115	-158.1984225	4.185946442	-41.28598973
116	-176.0031688	23.36982273	11.11583003
117	25.19394705	-55.61495168	-59.5102146
118	-41.9689889	-30.71800882	72.35173115
119	13.22395543	75.6688389	-35.99074828
120	49.66104911	-75.89797056	28.91528911
121	140.9452917	-31.07296751	42.77735854
122	75.91820967	-25.01961411	69.04869229
123	-31.21231922	-60.94299835	-54.10981051

124	12.96337083	55.58347312	60.26471247
125	-43.15131644	74.10384915	34.45963308
126	-133.5892459	57.43518845	-11.86169124
127	84.08814414	5.595468325	71.06362464
128	-31.66073552	77.88593486	-28.96578275
129	116.043362	56.03930713	33.1672253
130	15.38977036	-84.68129285	1.901070735
131	-100.490972	14.16680338	65.83185773
132	101.1491118	70.35357964	10.11705346
133	-0.571874475	-81.52975742	-22.62408166
134	77.59980178	-76.5051769	-9.447125614
135	18.36542152	16.37605165	78.0984939
136	-100.8065582	28.66935506	61.41381963
137	71.99082564	38.55448999	64.14151999
138	33.18147734	46.89878711	-65.15957932
139	-14.02162463	83.63685482	-12.91753678
140	-172.2151482	15.96729391	-25.06324516
141	153.6476265	33.87786357	31.12515908
142	-1.170824209	-23.15511123	76.97276668
143	-42.14039771	-14.92069553	76.62062892
144	88.6027675	-27.5834588	-65.25346369
145	70.50281582	31.23779374	-67.86842889
146	-2.573988215	-73.56330535	-40.06407558
147	29.86166503	25.41947128	75.23883806
148	-13.22502324	-29.83296948	74.69213328
149	-33.16550595	60.77047436	-54.06439562
150	25.27981479	29.33241223	74.28559992
151	36.34220802	60.72212745	53.72961799
152	13.6263449	55.50126181	60.30446696
153	130.2243393	48.43736569	-33.91968015
154	-20.44287714	14.47695757	78.33390555
155	91.08387202	-68.18448006	-27.02518818
156	-95.20866184	48.7504708	50.9871748
157	-34.0128886	-41.17473153	68.42437977
158	-34.74826179	-83.43529762	2.769131279
159	87.18048075	-62.76944419	-38.58274884
160	-65.93225425	72.50756292	-30.4972472
161	-30.96108584	2.043253066	-78.8482582
162	-120.7165609	36.79981131	-49.75324477
163	3.091404959	72.15180807	-42.27023692
164	25.09082028	-81.47857179	-20.03938223
165	166.1045975	-20.19562242	-29.65302735
166	-25.9235304	-76.03958091	-33.94914605
167	33.38526345	-64.02734841	-50.59826648
168	-12.15155976	76.14259887	-35.16693111

169	58.26703497	61.54721201	49.08788895
170	-57.84993042	26.55214878	-71.76125315
171	-165.0046271	-38.20633505	-3.956813716
172	-125.1461087	-4.130149975	58.78957752
173	132.9458237	59.05619549	-2.344731339
174	24.01512859	8.071536948	78.95848883
175	-35.4709238	-55.81509629	-58.35345807
176	67.27854737	-37.22738279	-65.77190577
177	-21.53011823	27.9050174	-74.99029323
178	-60.78673148	63.7093268	45.97423532
179	58.69365139	56.42943102	54.17679566
180	-75.62994254	37.29723323	64.01682036
181	22.93177484	44.26608616	67.57163126
182	-4.938930396	-34.33505911	-73.15160099
183	-43.78880772	5.742017316	-77.53859237
184	-14.16244536	-81.02617751	-23.38622984
185	8.007308412	-52.35629324	62.92722542
186	30.46739926	61.08508773	54.04724284
187	-81.40664072	71.07924516	-26.18015072
188	-14.61765673	-21.3139264	77.18570345
189	-29.81930992	-66.02903178	-48.70036709
190	61.13024989	0.167882275	-75.50616882
191	-11.48253984	39.81353695	-70.50692488
192	35.59211664	15.75045937	77.09320051
193	52.62812337	74.05910791	-31.99380537
194	165.4134504	-36.07273134	11.43598459
195	-27.92430594	-74.15846625	-37.18434184
196	-32.43789526	-55.69283727	-58.78535576
197	-28.09547989	40.73866488	69.15390281
198	-10.92599932	-84.57344579	-6.461389124
199	-104.9160984	28.42588316	60.2152662
200	-94.89525186	70.92164693	-16.14095843

Appendix III Monte Carlo convergence assessment for RCS2 populations

This appendix presents the Monte Carlo convergence check used to verify that the RCS2 landed-case population adopted in the thesis is sufficiently stable for the comparative DTE and ISL analysis in Chapter 5. The check is reported through six figures, covering respectively the DTE and ISL cases for 20 h, 40 h and 60 h of surface operation durations.

The convergence assessment monitors three cumulative data-volume indicators: $M = \mathrm{P50}$, $B^* = \mathrm{P95}$ and $W^* = \mathrm{P5}$. These represent the median, favourable and adverse population cases, respectively. The percentile-based indicators are used instead of the absolute best and worst realisations, since sample quantiles provide a mathematically robust measure of population behaviour that is less sensitive to extreme tail-end outlier noise [53], [54].
Each convergence figure is arranged in three vertical panels, one for each monitored indicator. The top panel shows B^*, the middle panel shows M and the bottom panel shows W^*. In each panel, the indicator value is plotted as the sampled population is increased, allowing the convergence of the favourable, median and adverse cases to be assessed separately.

The convergence path is evaluated in two stages, corresponding to the two parts of the x-axis in each panel. In the first stage, the number of realisations per surface location is increased while the number of surface locations is kept fixed. This checks the convergence of the asteroid initial phase and antenna-heading sampling for each landing location. In the second stage, the number of realisations per surface location is kept at the converged value from the first stage, while the number of sampled surface locations is increased. This checks the convergence of the landing-site population itself. This two-stage sampling path follows established Nested Monte Carlo (NMC) simulation frameworks, which are specifically formulated to isolate inner-loop conditional uncertainties from outer-loop macro-environmental parameters [55].

The convergence criterion is based on a 2% tolerance over three consecutive stable steps. The vertical markers in each figure indicate the points at which each stage satisfies the adopted convergence criterion. This criterion serves as a sequential stopping rule [56]. Requiring stability across consecutive steps ensures that the simulation has reached a true statistical asymptote rather than an accidental, temporary fluctuation. Stabilisation of M, B^* *and* W^* is therefore used as evidence that the selected RCS2 population is adequate for the comparative analysis.

Figure 103 to Figure 105 show the convergence assessment for the RCS2 DTE case at 20 h, 40 h and 60 h surface-operation durations. Figure 106 to Figure 108 show the corresponding assessment for the RCS2 ISL case over the same three time marks.

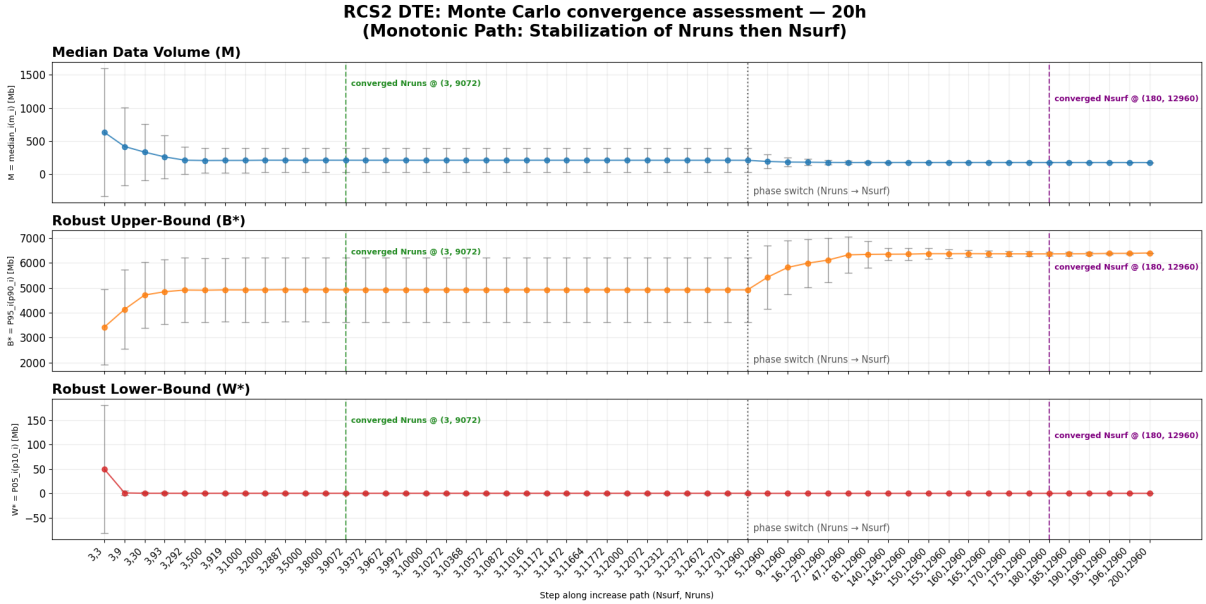


Figure 103: RCS2 DTE Monte Carlo convergence assessment at 20 h mission time

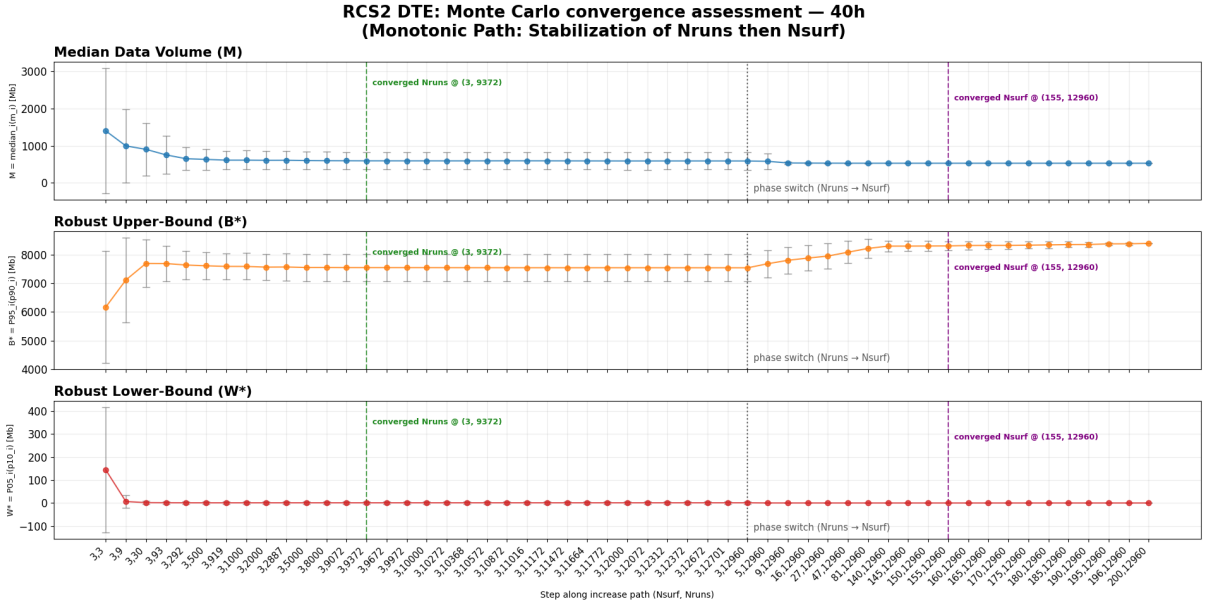


Figure 104: RCS2 DTE Monte Carlo convergence assessment at 40 h mission time

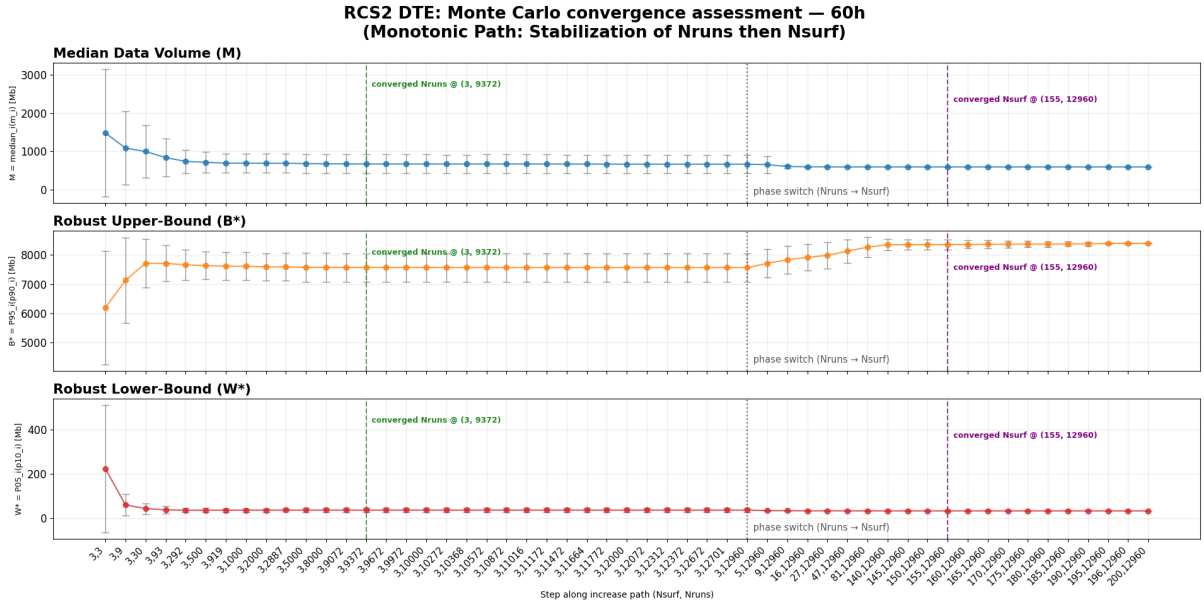


Figure 105: RCS2 DTE Monte Carlo convergence assessment at 60 h mission time

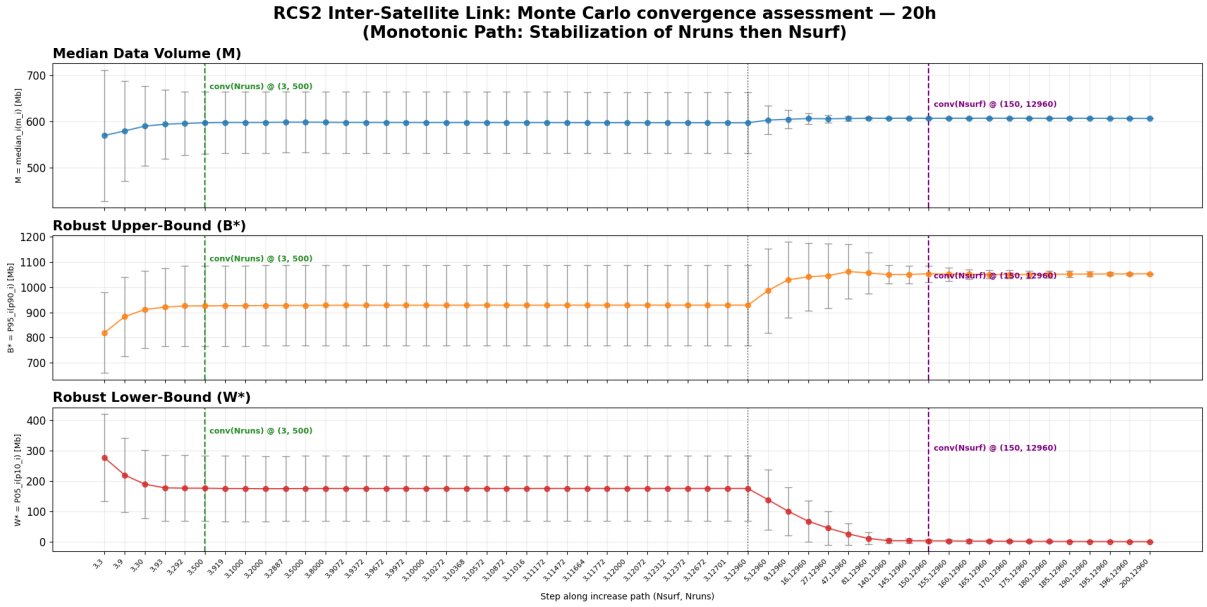


Figure 106: RCS2 ISL Monte Carlo convergence assessment at 20 h mission time

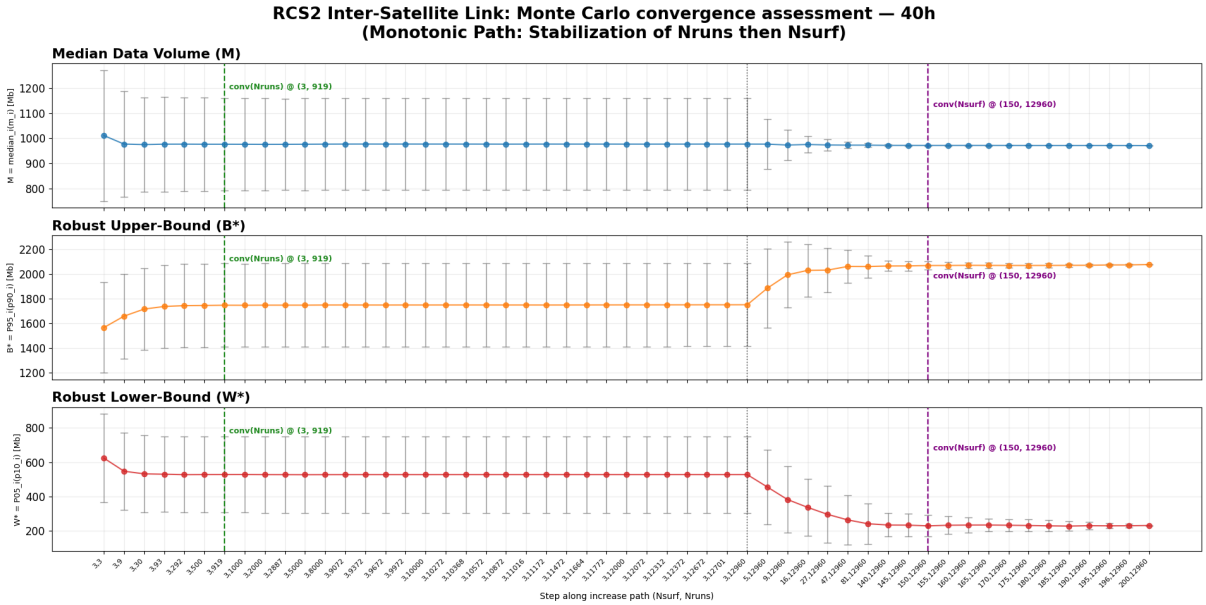


Figure 107: RCS2 ISL Monte Carlo convergence assessment at 40 h mission time

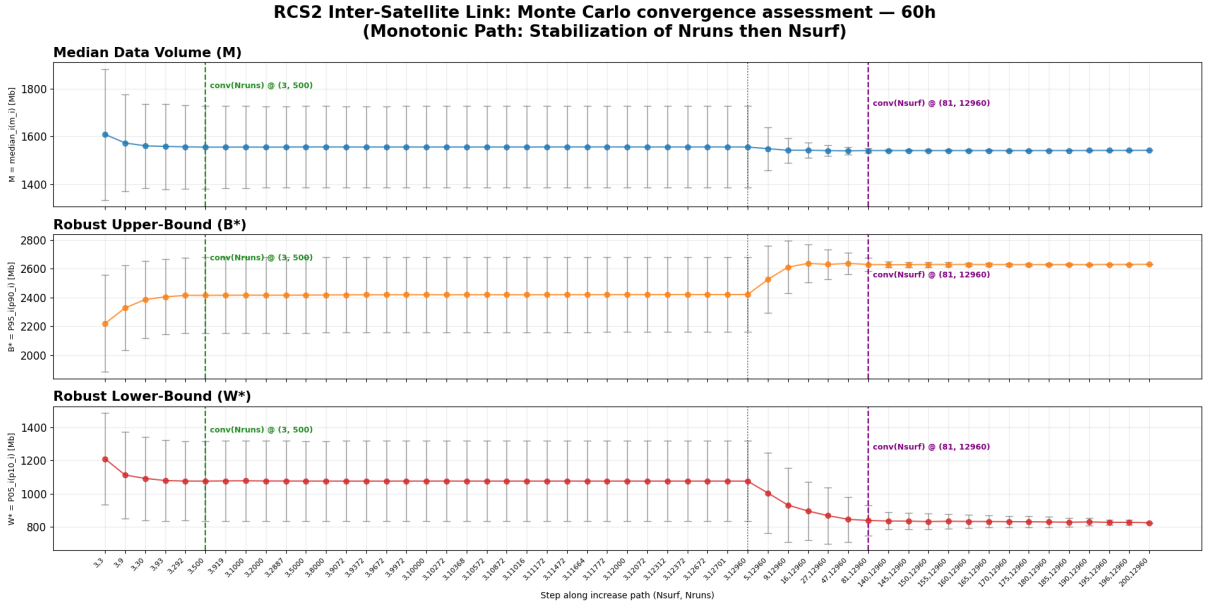


Figure 108: RCS2 ISL Monte Carlo convergence assessment at 60 h mission time

Across the six convergence checks shown in Figure 103 to Figure 108, the monitored indicators satisfy the adopted convergence criterion within the selected population size. This supports the use of the RCS2 landed-case population adopted in the thesis for the DTE and ISL nominal and sensitivity assessments.

Appendix IV DTE TM rates and ISL data-mode definitions

This appendix reports the full set of discrete DTE telemetry rates and ISL data modes used by the communication models in Chapter 3 and subsequent analysis. Table Ap.5 provides the DTE telemetry information rates for the Turbo 1/2 and Turbo 1/4 configurations.

Table Ap.5 - DTE TM-rate definitions

#TMrate	Turbo Code 1/2 TM rate [bps]	Turbo Code 1/4 TM rate [bps]
TM01	15.86	7.94
TM02	41.11	20.59
TM03	74.60	37.30
TM04	115.98	58.10
TM05	231.97	116.19
TM06	356.87	178.76
TM07	715.95	358.62
TM08	1195.71	598.93
TM09	4829.64	2419.14
TM10	8916.68	4466.32
TM11	12879.23	6451.14
TM12	23182.73	11612.11
TM13	29676.67	14864.89
TM14	34692.47	17377.28
TM15	39097.52	19583.75
TM16	44381.11	22230.27
TM17	49263.02	24675.59
TM18	54135.23	27116.06
TM19	58646.57	29375.77
TM20	69384.65	34754.41
TM21	80759.02	40451.78
TM22	120153.70	60184.37
TM23	158913.25	79598.83
TM24	197052.36	98702.51
TM25	307894.15	154222.60

Table Ap.6 provides the ISL data modes with the corresponding receiver-sensitivity thresholds. Values are provided for the two bandwidth (BW) allocations considered in the thesis.

Table Ap.6 - ISL data-mode rates definitions

Mode	Symmetric-BW allocation [kbps]	High-priority BW allocation [kbps]	Receiver sensitivity [dBm]
Q1	8.19	20.48	-102.0
Q2	2.29	5.72	-103.0
Q3	0.82	2.08	-105.5
B1	4.95	9.88	-106.0
B2	1.04	2.60	-107.0
B3	0.21	0.52	-108.0

RCS1 is assumed to use the rates as defined in the Symmetric-BW allocation, while RCS2 uses the rates defined in the High-priority BW allocation.

Appendix V Detailed link budget plots

This appendix provides detailed DTE telemetry link budget plots supporting the reduced link budget discussions in Chapter 4. The main chapter presents the link contributors most relevant to the interpretation of the RCS1 DTE case, while this appendix retains two representative full term-by-term plots for traceability, including the selected turbo-code evolution over time. These figures are not used as separate results, but are provided as supporting material for the DTE performance discussion in the main text. Figure 109 shows the full RCS1 mission timeline, while Figure 110 shows a zoomed view around closest approach.

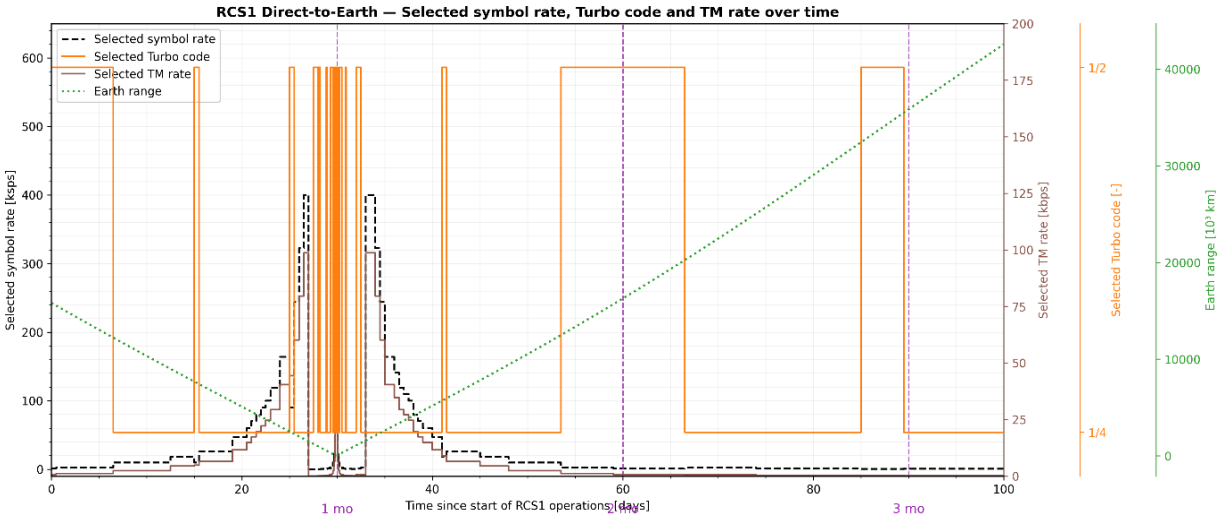


Figure 109: RCS1 DTE detailed link budget

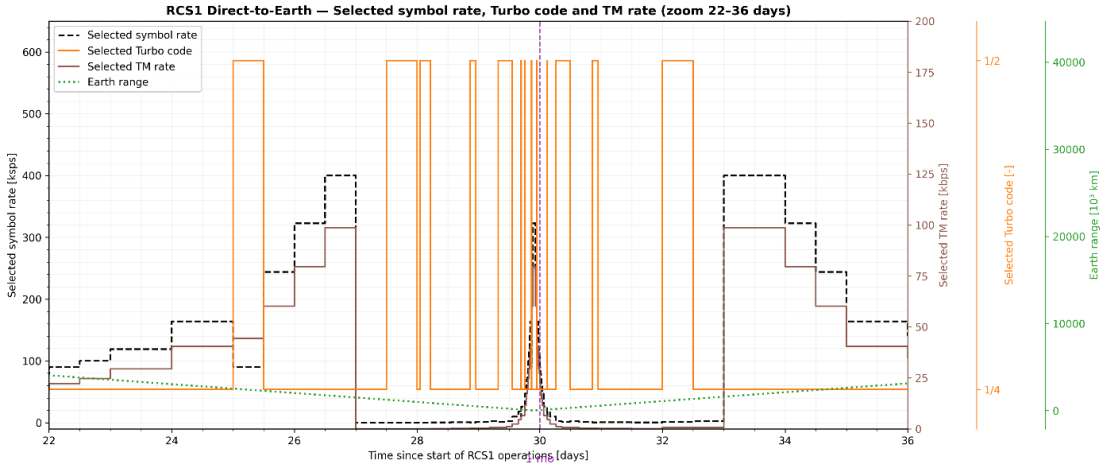


Figure 110: RCS1 DTE detailed link budget around close-approach