Assessing Long-Term Endogenous Risk in Satellite Constellations

PRISM: A Plane-Resolved Integrated Stochastic Model for Constellation Safety

Master's Thesis

Massimo Strani

Delft University of Technology

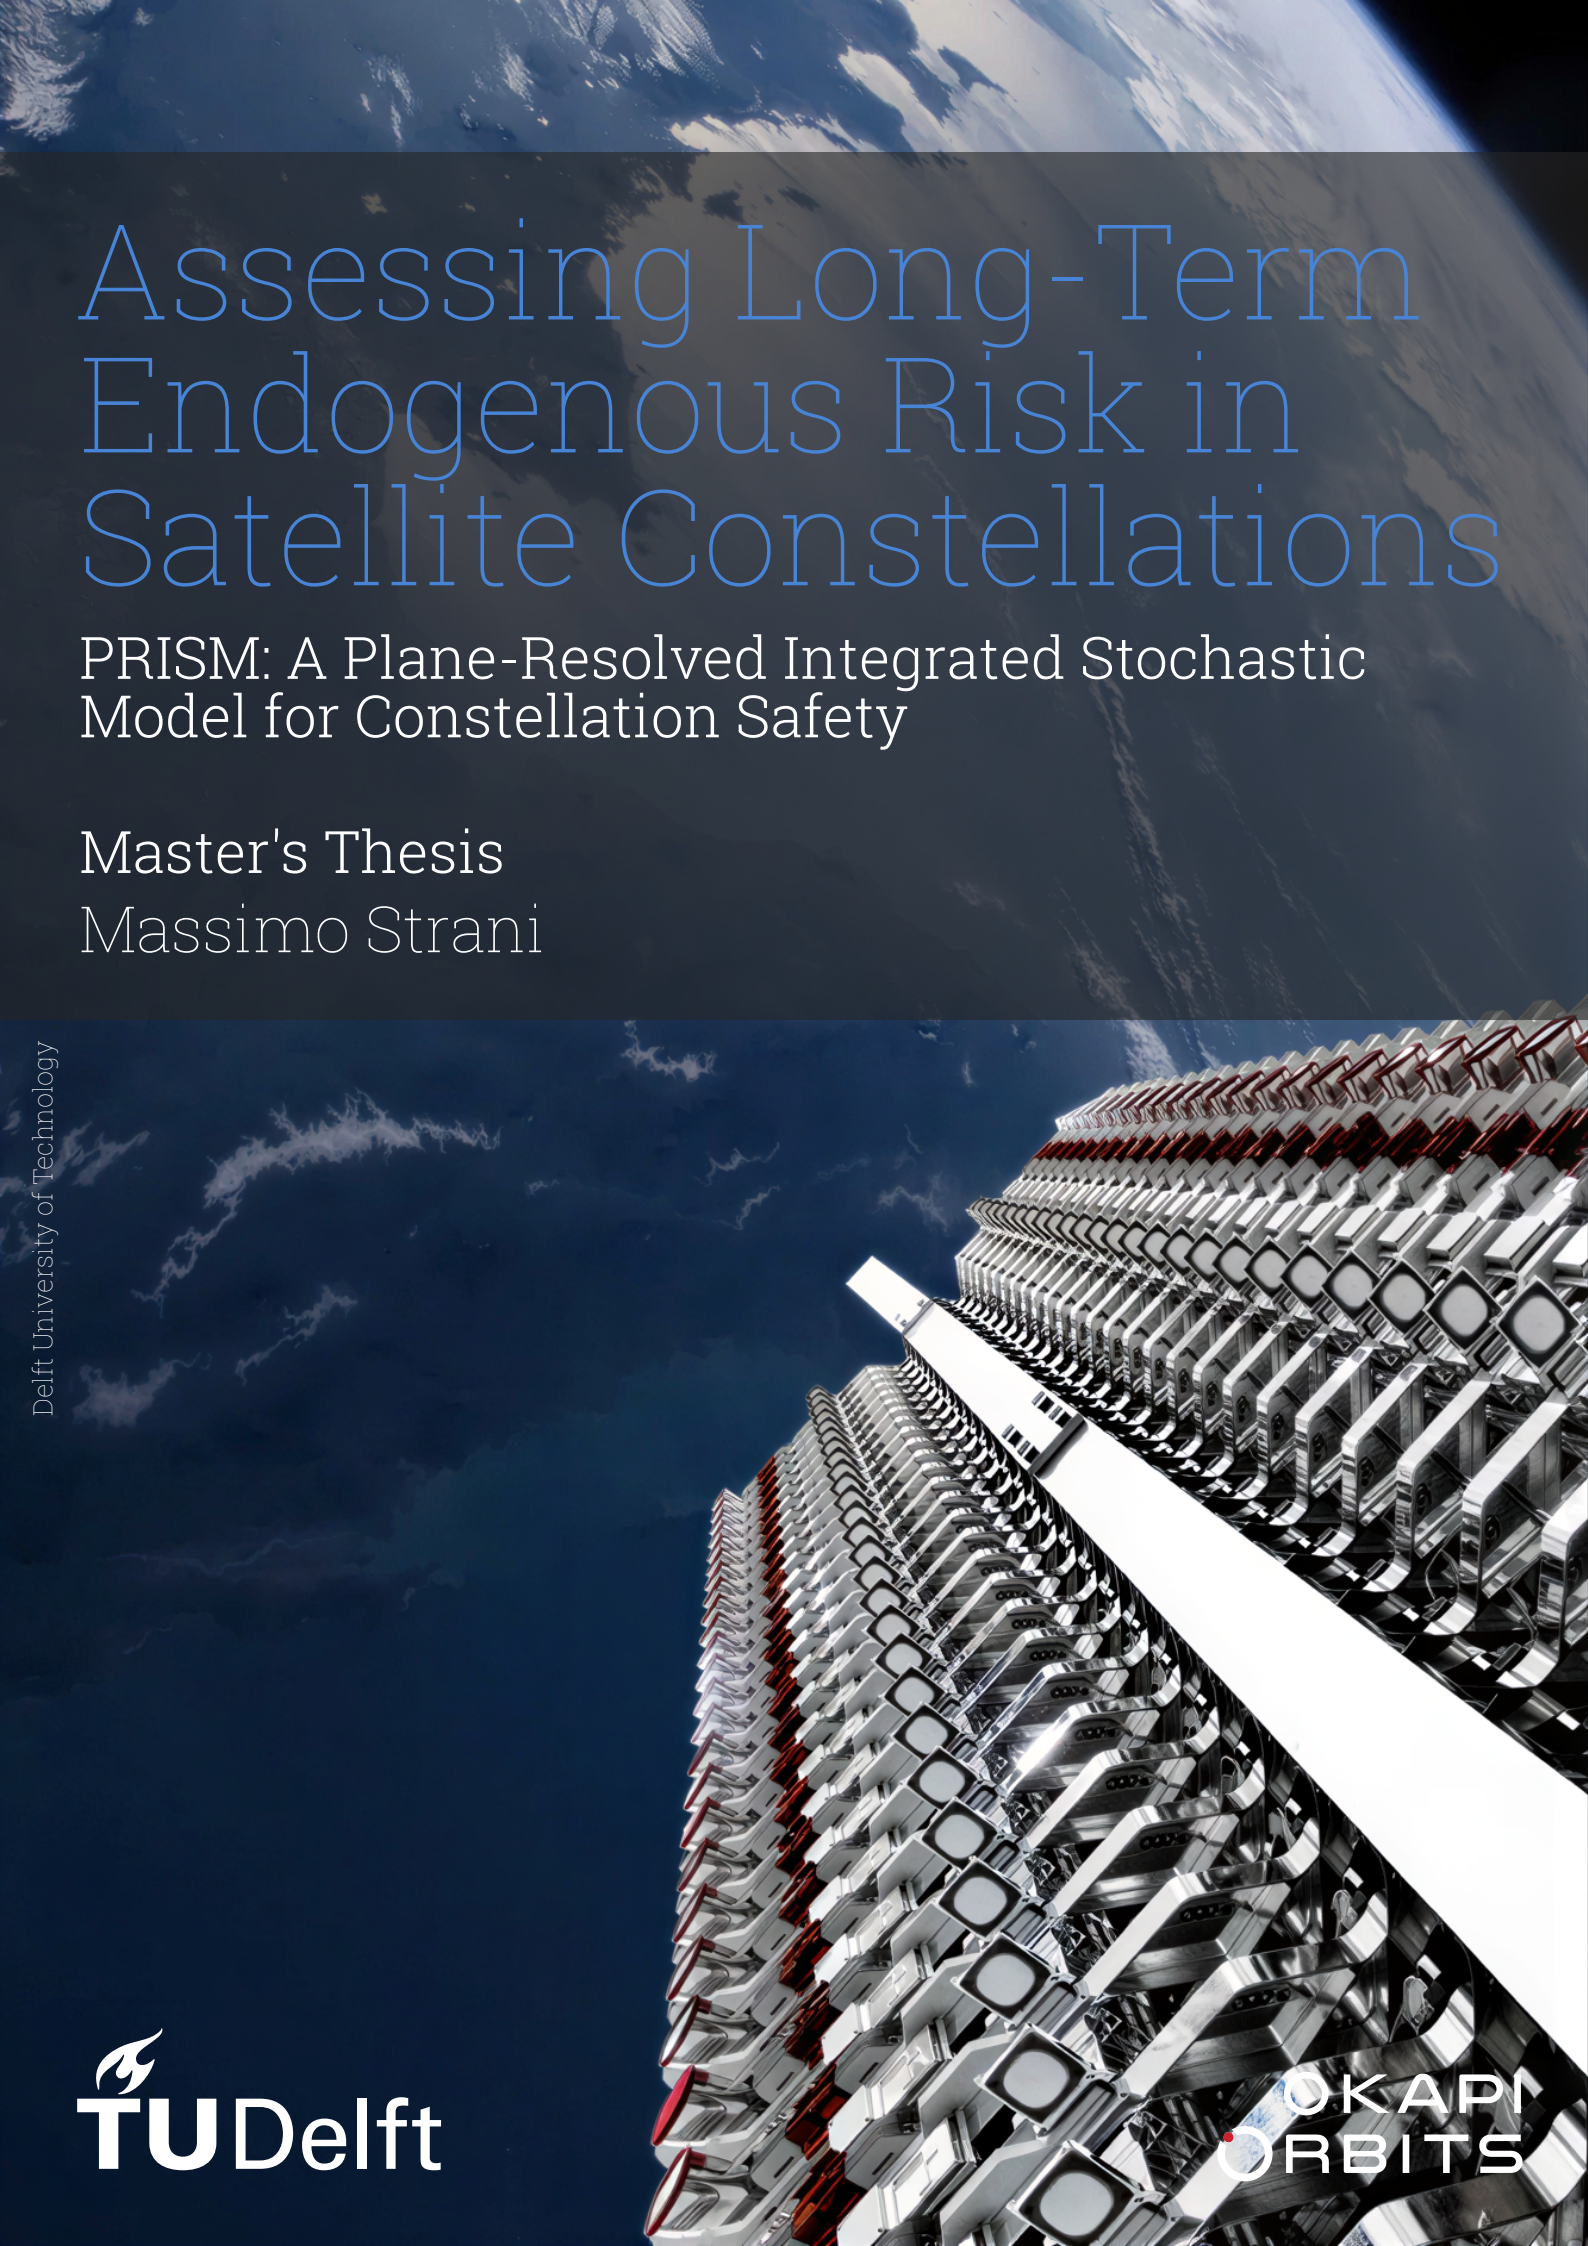

TU Delft

OKAPI ORBITS

Assessing Long-Term Endogenous Risk in Satellite Constellations

PRISM: A Plane-Resolved Integrated Stochastic Model for Constellation Safety

by

Massimo Strani

to obtain the degree of Master of Science in Aerospace Engineering at Delft University of Technology,
to be defended publicly on Wednesday September 9, 2026 at 10:45.

Student number:	5563208	
Project duration:	February 2026 – September 2026	
Thesis committee:	Dr. S. Gehly,	TU Delft, Supervisor
	Ir. K.J. Cowan,	TU Delft, Chair
	Dr. I. Akay,	TU Delft, Independent Examiner
	F. Läuferts,	External Supervisor

Cover: *Starlink Mission* by SpaceX (2019) under CC BY-NC 2.0. https://www.flickr.com/photos/spacex/47926144123 (upscaled)

TU Delft

OKAPI ORBITS

Preface

This thesis marks the end of my M.Sc. studies in Aerospace Engineering and of my time as a student at Delft University of Technology. The last two years have been extremely challenging, but fundamental to my personal and professional development: I wouldn't be honest if I didn't admit this journey came with many sleepless nights, but now I can look back with deep satisfaction, aware of the extensive engineering knowledge my studies have given me. This thesis in particular has pushed my knowledge to new boundaries and resurfaced my passion for mathematics, and I feel proud to contribute, even if just in small part, with my own tool PRISM to solving the new challenges faced by the space sector.

This project merges a wide variety of topics and skills, and I wouldn't have been able to make it through myself without the knowledge and expertise of my colleagues and supervisors, who have supported me all the way from the start. First, I would like to thank Dr. Steve Gehly, my supervisor at TU Delft, for the unwavering support at every step of this work. Your words of encouragement, your friendliness, and your thorough understanding of the subject have shown me that no task is too challenging to tackle. Thank you for spending countless hours helping me decipher mysterious math and for all our conversations, about the thesis and everything else. I am grateful every day to have had you as my supervisor. I would also like to thank the graduation committee, Ir. K.J. Cowan and Dr. I. Akay, for agreeing to be part of my defense.

This project came to be as a follow-up to my internship at OKAPI:Orbits. First, I would like to thank my whole team, who I have always been able to lean on, both professionally and personally. You have made my experience in Braunschweig a lot better, and I cherish every serious and silly moment spent together. To Frederik Läuferts, I want to thank you for taking the time to supervise me despite your own workload, for fostering my skills, showing great leadership, and guiding me through my first professional experience, well before the start of this thesis. A heartfelt thanks goes to Sven Flegel, who has been like a mentor to me through every difficulty life has thrown at me this year.

I am beyond grateful to everyone who has been by my side during the last couple of years. No words can express the gratitude and appreciation I feel, but I'll try my best. I would like to thank all my friends at Pegasus, who have made my time in Delft truly special, and showed me that it is never too late to chase my own dreams. I wish to see you all soon, and thank you for giving me an excuse to always come back to Delft. A special thanks to my friends in Trieste, who always make me feel like I never left, and across the world, especially Felix, Filippo and Alessio, for always being there even if we haven't talked in months. To Amanda, thank you for your unconditional support and for doing all of this side by side, I wouldn't have made it without you.

Finally, and most importantly, I want to thank my family, especially Adriana, Gianni and Daniela, who have made me the person I am today. To my mother Rossella, who has been a monumental example of resilience and strength, I dedicate this thesis. *Anche se lontani, vi porto sempre nel mio cuore.*

Massimo Strani
Trieste, August 2026

Statement on the Use of AI

Throughout the course of this project, several large language models (LLMs) were used in a supporting role, both in the preparation of this document and in the implementation of the model it describes. Their use was mainly limited to rephrasing sentences and correcting grammar, assisting in navigating literature and existing codebases, refactoring and optimizing code, and adjusting figures. The main ideas, analyses and conclusions presented in this thesis, as well as the implementation of PRISM, are entirely my own; all AI-assisted suggestions were critically reviewed and edited, and I take full responsibility for the final content of this document and the software accompanying it.

Abstract

Large satellite constellations have become the driving force of the New Space economy. Starlink, OneWeb, Amazon Leo, and TeraWave already operate or hold approval for fleets numbering in the thousands, and recent filings by Chinese operators propose constellations approaching 200 000 satellites. At this scale a constellation is not a collection of independent spacecraft sharing an orbital regime but a single interdependent system, whose satellites launch, station-keep, maneuver, fail, and de-orbit in a shared, congested volume of space, and whose internal traffic can rival the risk posed by the surrounding debris environment. Research on constellation safety has nonetheless concentrated almost entirely on the exogenous problem, how a constellation interacts with the pre-existing debris population and with other operators' assets, while the endogenous risk a constellation's own satellites pose to one another, through injection and disposal traffic, collision-avoidance cascades, and internal fragmentation, has received comparatively little attention. The few studies addressing endogenous risk directly treat it as a short-term, many-versus-many conjunction-screening problem over a fixed geometry at a single epoch. Essentially none address the long-term stability of a constellation over its full operational horizon, nor embed the problem in an evolutionary framework able to propagate stochastic failures, maneuvers, and fragmentation across decades.

This thesis develops PRISM: a Plane-Resolved Integrated Stochastic Model for the long-term, system-wide evolution of a satellite constellation's endogenous collision risk. The tool integrates the natural orbital dynamics of drag and J_2 precession, station-keeping, orbit-raising and de-orbiting maneuvers, stochastic disposal failures, catastrophic and non-catastrophic collisions, fragmentation, and debris-mitigation compliance, over horizons of decades. Its distinguishing feature is that it resolves risk not only in aggregate, but down to individual orbital planes, so that the specific shells and planes that drive a constellation's collision risk can be identified rather than averaged away. Methodologically, the model is built as a network whose nodes map to the orbit bins of a geometric lattice, defined jointly by altitude h, inclination i, and RAAN Ω, together with a specie (station-keeping, rising, de-orbiting, non-maneuverable, or fragment). The collision rate between any two nodes is derived from first principles by transforming the classical Cartesian collision-rate integral into the (h, i, Ω, u) phase space native to a Walker constellation, averaging the argument of latitude u so that the rate between two bins reduces to a closed-form overlap of great circles, evaluated through incomplete elliptic integrals and a symmetry-exploiting numerical quadrature.

The network's evolution is governed by an evolutionary model in which every stochastic event is treated as a step-wise homogeneous Poisson jump process over each sufficiently small discrete time step. Catastrophic and non-catastrophic collisions are resolved through the NASA Standard Breakup Model (SBM), which generates a fragment cloud that is redistributed across the lattice; collision avoidance enters as a success-rate factor on each pair rate, weighted differently depending on the objects' ability to maneuver. Each run is a single stochastic realization, so the evolution of the population is estimated as an ensemble mean over independent seeds. Because plain Monte Carlo estimates of collision statistics for realistically sparse constellations are prohibitively expensive, a conditioning scheme draws the rare first-collision event directly from its exact arrival-time distribution. Finally, the links of the network are defined as a closed-form rate evaluated on a recorded population state, so the entire edge set of any epoch can be rebuilt on demand from the ensemble-averaged node counts. Node strengths, source-sink balances, and plane-level concentration measures then identify which orbital planes, shells, and species drive collision risk and debris production.

PRISM is verified by comparison against MOCAT-MC, an independently developed, open-source Monte Carlo tool that detects collisions through a structurally unrelated, spatial-density method. Run over a century-scale evolution of the real LEO population as of January 2023,

the two models agree closely on both cumulative collision counts and population evolution by specie. The model is then applied to Starlink, OneWeb, and Iridium-NEXT. In the Walker-star patterns of OneWeb and Iridium-NEXT, the collision rate per link rises by four to five orders of magnitude from co-planar pairs to the 180° seam, where ascending and descending planes meet head-on, so that a small number of plane pairs carry almost all of the constellation's internal risk. A freshly generated fragment cloud is born as a single-plane object, with 88 % to 99.9 % of its fragments in the parent plane, and spreads only slowly under differential nodal drift, so a model that distributes the cloud uniformly on day one misstates the near-field risk by close to two orders of magnitude. Finally, a reduced, shell-level form of the mean-field system yields a carrying-capacity margin for each altitude shell, defined as the factor by which its present population could grow before collisional debris production exceeded removal by atmospheric drag.

This work contributes a computationally tractable, physically consistent tool for the long-term, plane-resolved assessment of endogenous collision risk, enabling operators and regulators to run the systematic parameter sweeps across failure rates, disposal compliance, and launch profiles needed to design safer constellations, and establishing a methodological foundation for extending evolutionary space-environment models beyond the aggregate, whole-population perspective they have so far been confined to.

Contents

Preface iii

Statement on the Use of AI v

Abstract vii

Nomenclature xiii

1 Introduction 1

1.1 Motivation . . . 3
1.2 Problem Statement . . . 4
1.3 Research Question . . . 6
1.4 Outline of the Report . . . 7

2 State of the Art 9

2.1 Constellation Dynamics . . . 9
2.1.1 Constellation Geometries . . . 9
2.1.2 Reference Missions . . . 11
2.2 Evolutionary Models for the Space Environment . . . 12
2.2.1 Semi-deterministic Models . . . 12
2.2.2 Source-sink Models . . . 14
2.2.3 Continuum-based Models . . . 16
2.3 Collision Rate Computation . . . 17
2.3.1 Pairwise Orbital Intersection Methods . . . 17
2.3.2 Spatial Density and Flux Methods . . . 18
2.3.3 Phase-Space Density Formulation . . . 20
2.4 Network Theory . . . 20
2.4.1 Network Theory for Space Applications . . . 20
2.4.2 A Network-based Evolutionary Model of the Space Environment . . . 21

3 Astrodynamics 23

3.1 Reference Frames . . . 23
3.2 Coordinate Systems . . . 24
3.2.1 Cartesian Coordinates . . . 24
3.2.2 Equatorial Coordinate System . . . 24
3.2.3 Keplerian Elements . . . 24

3.3 Equations of Motion . . . 25
3.4 Perturbations . . . 26
3.4.1 Earth Oblateness . . . 27
3.4.2 Atmospheric Drag . . . 28
3.4.3 Osculating and Mean Elements . . . 29

4 Model Formulation **31**
4.1 The Network's Evolutionary Model . . . 32
4.1.1 Constellation Dynamics as a Poisson Point Process . . . 32
4.1.2 Definition of the Stochastic Dynamic Equations . . . 33
4.2 Collision Rate Formulation . . . 34
4.2.1 Overlap Integral . . . 37
4.2.2 Relative Velocity . . . 44
4.2.3 Average Impact Cross-sectional Area . . . 45
4.3 Natural Dynamics . . . 45
4.4 Orbit Rising and De-orbiting . . . 48
4.5 Break-up Model . . . 50

5 Software Framework **53**
5.1 Software Architecture . . . 53
5.1.1 Initial Population and Data Sources . . . 55
5.1.2 Scenario Creation . . . 55
5.2 Stochastic Simulation . . . 56
5.2.1 Step Ordering . . . 56
5.2.2 Collision Sampling . . . 57
5.2.3 Conditioning . . . 58
5.3 Verification . . . 60

6 Sensitivity Analysis and Robustness **63**
6.1 Global Sensitivity Analysis . . . 64
6.2 Lattice Discretization . . . 67
6.2.1 Altitude Binning . . . 67
6.2.2 Selecting Altitude Binning from the Mission Profile . . . 72
6.2.3 Inclination and RAAN Binning . . . 79
6.3 Time Discretization . . . 84

7 Network Theory **87**
7.1 Definition of the Network Links . . . 87
7.2 Network Topology . . . 91
7.2.1 Adjacency Matrices . . . 91
7.2.2 Node Strength . . . 91

7.3 Network Dynamics 92

8 Results **95**
8.1 Baseline Constellation Evolution 95
8.1.1 Plain Monte Carlo 96
8.1.2 Forced Collision 99
8.2 Network Topology and Plane Resolution 100
8.2.1 Population per Plane 101
8.2.2 Risk per Plane 102
8.2.3 Node Strength 105
8.2.4 The Case for Plane Resolution 107
8.3 Carrying Capacity 108
8.3.1 The Capacity Indicator 108
8.3.2 Baseline Capacity per Constellation 111
8.4 Case Studies 112
8.4.1 Impact of Different CAM Policies 112
8.4.2 Impact of Different PMD Success Rates 113
8.4.3 Impact of Different Launch Strategies 115

9 Conclusion **117**
9.1 Conclusions 117
9.2 Recommendations for Future Work 121

References **123**

A Research Planning **127**
A.1 Work Package Breakdown 127
A.2 Milestones and Planning Overview 129

B Probability Foundations and Stochastic Processes **131**
B.1 Common Definitions 131
B.2 Generalized Dirac-Delta Function 133
B.3 Poisson Point Processes 135

C Numerical Methods **137**
C.1 Gauss-Legendre Quadrature 137
C.2 Fourier Methods for Circulant Sums 137
C.2.1 The Discrete Fourier Transform 138
C.2.2 Application to the Aggregate Collision Rate 139

Nomenclature

Abbreviations

Abbreviation	Definition
ADR	Active Debris Removal
CAM	Collision Avoidance Maneuver
CDM	Conjunction Data Message
CNES	Centre National d'Études Spatiales
CSI	Criticality of Spacecraft Index
DAMAGE	Debris Analysis and Monitoring Architecture to the Geosynchronous Environment
DELTA	Debris Environment Long-Term Analysis
DISCOS	Database and Information System Characterising Objects in Space
ECOB	Environmental Consequences of Orbital Breakups
EOL	end-of-life
ESA	European Space Agency
FCC	Federal Communications Commision
FDS	Flight Dynamics Service
GEO	Geostationary Orbit
GPS	Global Positioning System
IADC	Inter-Agency Space Debris Coordination Committee
ISO	International Organization for Standardization
ITU	International Telecommunication Union
JAXA	Japan Aerospace Exploration Agency
JB2008	Jacchia-Bowman 2008
LEGEND	LEO-to-GEO Environment Debris
LEO	Low Earth Orbit
LUCA	Long Term Utility for Collision Analysis
MASTER	Meteoroid and Space Debris Terrestrial Environment Reference
MEDEE	Modelling the Evolution of Debris in the Earth Environment
MEO	Medium Earth Orbit
MIT	Massachusetts Institute of Technology
MOCAT	MIT Orbital Capacity Tool
MOID	Minimum Orbital Intersection Distance
NASA	National Aeronautics and Space Administration
NEO	Near-Earth Object
NEODEEM	Near-Earth Orbital Debris Environment Evolutionary Model
NESSY	NEtwork model for Space SustainabilitY
ODE	Ordinary Differential Equation
PDE	Partial Differential Equation

Abbreviation	Definition
PDF	Probability Density Function
PIB	Particle-in-a-Box
PMD	Post Mission Disposal
PRISM	Plane-Resolved Integrated Stochastic Model
RAAN	Right Ascension of the Ascending Node
RSO	Resident Space Object
RSONet	Resident Space Object Network
SBM	Standard Breakup Model
SDM	Space Debris Mitigation
SGP4	Simplified General Perturbations
SIS	susceptible-infectious-susceptible
SRM	Solid Rocket Motor
SRP	Solar Radiation Pressure
SSA	Space Situational Awareness
SST	Space Surveillance and Tracking
STAT	STochastic Analog Tool
STELA	Semi-analytic Tool for End of Life Analysis
STM	Space Traffic Management
TCA	Time of Closest Approach
TLE	Two Line Element
USSPACECOM	United States Space Command

Symbols

Symbol	Definition	Unit
A	Active station-keeping satellite or cross-sectional area	- / m^2
A_c	Impact cross-sectional area	m^2
a	Semi-major axis	m
a_{inj}, a_{tgt}	Injection and target semi-major axis	m
B	Ballistic coefficient	m^2/kg
$\boldsymbol{B}(\cdot,\cdot)$	Bilinear mean-field collision operator	
C_D	Drag coefficient	
C	Courant number	
$\mathrm{Cov}(\cdot,\cdot)$	Covariance of the given random variables	
D	Active de-orbiting satellite	
$d^q_{S_i}$	Diameter of the q-th object in node S_i	m
$d^{(1)}_{S_i}$, $d^{(2)}_{S_i}$	First and second diameter moments of node S_i	m / m^2
E	Eccentric anomaly	°
$\mathbb{E}[\cdot]$	Expectation of the given random variable	
EMR	Energy-to-mass ratio	J/kg
e	Eccentricity	
F	Fragments	

Symbol	Definition	Unit
$F(\varphi, k^2)$	Incomplete elliptic integral of the first kind of amplitude φ and modulus k^2	
$\boldsymbol{f}$	Perturbing acceleration vector	m/s^2
$f_{(\cdot)}$	Probability density function of the given random variable	
$\boldsymbol{g}$	Exogenous forcing term on the system's evolution	
$\boldsymbol{G}(\cdot)$	Reduced mean-field dynamics vector field	
G_F	Fragment component of $\boldsymbol{G}(\cdot)$	
H	Atmospheric scale height	m
h	Altitude	m
Δh	Altitude bin width	m
$\mathcal{I}(s_\delta)$	Set of argument of latitude values for which two objects are found at the same position	
$\mathcal{I}_i(s_\delta)$	Inclination-controlled subset of $\mathcal{I}(s_\delta)$	
$\mathcal{I}_\Omega(s_\delta)$	RAAN-controlled subset of $\mathcal{I}(s_\delta)$	
$I_{(\cdot)}^{(\ell)}$	ℓ-th lattice bin in the given dimension	
$\mathbb{1}\{\cdot\}$	Indicator function of the given event	
i	Inclination	°
Δi	Inclination bin width	°
$\mathcal{J}(s_\delta)$	Analytical integral of the inner node	
$\boldsymbol{J}$	Jacobian matrix	
J_ℓ	ℓ-th geopotential coefficient	
J_2	Second zonal harmonic of the geopotential	
J	Total number of bins in the lattice	
K	Number of co-located object pairs in a bin	
L_c	Characteristic length of a fragment	m
$\boldsymbol{L}(\cdot)$	Linear mean-field collision operator	
M	Mean anomaly	°
M_{tot}, m_t, m_p	Total, target, and projectile mass in a collision	kg
m	Mass	kg
N	Non-maneuverable faulty satellite	
$N_{(\cdot)}$	Number of lattice bins in the given dimension	
$n_{(\cdot)}$	Phase-space number density in the given phase-space	
n_{eff}	Effective number of participating RAAN bins	
P_c	Probability of collision	
P_ℓ	Legendre polynomial of degree ℓ	
$Q_{(\cdot)}$	Number of Gauss-Legendre quadrature points in the given dimension	
R	Active rising satellite	
$R_\oplus$, R_{eq}	Earth's mean and equatorial radius	m
$\boldsymbol{r}$, r	Position vector and its magnitude in the ECI frame	m
r_p, r_a	Perigee and apogee radii	m
S	Object specie	
s_δ	Sine of declination, $s_\delta \equiv \sin\delta$	
s_{CAM}	Collision avoidance success rate	
$s_{S_i}^{\mathcal{C}}$, $s_{S_i}^{F}$	Collision and fragment-production node strength of node S_i	
$s_{S_i}^{\text{out}}$, $s_{S_i}^{\text{in}}$	Outgoing and incoming transport node strength of node S_i	

Symbol	Definition	Unit
$\mathrm{sgn}(\cdot)$	Sign of the given quantity	
T	Orbital period	s
T_{EOL}	Mean time to end-of-life	s
Δt_k	Time step of the discrete-time system	s
t	Time	s
t_f	Final time of the simulation	s
$\mathcal{U}_{\ell mn}$	Region of the (h, i, Ω, u) phase-space identified by the ℓ-th altitude bin, m-th inclination bin, and n-th RAAN bin	
U	Gravitational potential	J/kg
$\boldsymbol{u}$	Element of the (h, i, Ω, u) phase-space	
u	Argument of latitude	°
$\mathrm{Var}(\cdot)$	Variance of the given random variable	
$\boldsymbol{v}$, v	Velocity vector and its magnitude in the ECI frame	m/s
v_{air}	Relative velocity with respect to the atmosphere	m/s
v_{rel}, v_{12}	Relative velocity (between nodes S_i and S_j)	m/s
v_{rel}^{osc}	Relative velocity from the osculating orbital elements	m/s
Δv	Velocity kick	m/s
W	Effective width of an altitude occupancy distribution	
w_q, ξ_q	q-th Gauss-Legendre quadrature weight and abscissa	
w_κ	Half-spectrum Parseval weight at frequency κ	
$w_{q,\ell}^{(k)}$	Fraction of the k-th step that object q spends in the ℓ-th altitude bin	
$\boldsymbol{X}$, $\boldsymbol{X}^*$, $\boldsymbol{X}_k$	State vector, its equilibrium point, and its value at step k	
x_{S_i}	Number of objects in node S_i	
x_S	Total number of objects of species S	
$\boldsymbol{Y}(\cdot)$	Poisson-jump propensity vector	
$Z(\geq L_c)$	Cumulative number of fragments with characteristic length above L_c	
α	Right ascension	°
β	Latitude	°
α_{CAM}	Collision avoidance scaling factor	
Γ_ℓ	Altitude factor of the collision rate in the ℓ-th shell	
γ	Failure rate of post mission disposal maneuvers	
$\Delta_{S_iS_j}$	Observed number of collisions between nodes S_i and S_j	
$\Delta_{S_iS_j}^{cat}$, $\Delta_{S_iS_j}^{nc}$	Observed number of catastrophic and non-catastrophic collisions between nodes S_i and S_j	
$\Delta_{S_i}^{col}$, $\Delta_{S_i}^{dam}$	Observed number of collision losses and non-catastrophic damage events of node S_i	
$\Delta_{S_i}^{\gamma}$	Observed number of failed post mission disposal maneuvers	
$\Delta_{S_i}^{EOL}$	Number of satellites that have reached end-of-life	
$\Delta_{R_i}^{dep}$	Number of satellites deposited from node R_i into the active specie	
Δ^{inj}	Number of satellites injected by new launches	
$\delta(\cdot)$	Dirac delta function	
δ_{ij}	Kronecker delta function	
δ	Declination	°
$\varepsilon_{S_i}^{+}$, $\varepsilon_{S_i}^{-}$	Inflow and outflow of objects into node S_i	

Symbol	Definition	Unit
ε	Discretization error	
ϵ	Low-thrust acceleration	m/s^2
$\zeta_{S_j S_\ell}$	Total number of fragments generated by the collision between nodes S_j and S_ℓ	
$\zeta_{S_j S_\ell}^{S_i}$	Number of fragments generated by the collision between nodes S_j and S_ℓ flowing to node S_i	
$\dot{\eta}$, $\dot{\eta}_{S_i S_j}$	Collision rate (between nodes S_i and S_j) without accounting for collision avoidance	1/s
θ	True anomaly	°
θ_{12}	Crossing angle between two orbital planes	°
ϑ	Object flow rate	1/s
Λ	Total expected number of collisions	
$\Lambda_{S_i S_j}^{(k)}$	Mean number of collisions between nodes S_i and S_j at step k	
Λ_k	System-wide mean number of collisions at step k	
λ_k	System-wide collision rate accounting for collision avoidance	1/s
μ	Mean of the given random variable	
$\mu_\oplus$	Earth's gravitational parameter	m^3/s^2
μ_2, μ_3	Second and third occupancy moments of a bin	
μ', μ_N, μ_F	Mixed moments of the N-F cross-species bin occupancy distribution	
$\Xi^{\mathcal{C}}$, $\Xi^{\mathcal{T}}$	Collision and transport adjacency matrices	
$\xi_{S_i S_k}^{F_j}$	Fraction of fragments from the collision between nodes S_i and S_k allocated to node F_j	
ρ	Atmospheric density	kg/m^3
ϱ	Population scaling factor of the carrying-capacity indicator	
ϱ_{crit}	Critical value of ϱ at the carrying-capacity threshold	
Σ	Total number of collisions occurred	
$\Sigma_{S_i S_j}^{(k)}$	Number of collisions between nodes S_i and S_j at step k	
Σ_k	System-wide number of collisions at step k	
σ	Standard deviation	
τ	Elapsed time since the start of a maneuver	s
τ_∞	Time taken by a altitude band of derelicts to fully develop	s
$\Phi_{m_1 m_2 \Delta n}$	Overlap tensor over two inclination bins and a RAAN bin offset	
$\Phi_{m_1 m_2 \Delta n}^{\text{sym}}$	Swap-averaged overlap tensor	
$\Phi_{m_1 m_2 \Delta n}^{v}$	Velocity-weighted overlap tensor	
ϕ	Geocentric latitude	°
$\chi_{S_i S_j}^{(\tau)}$	Transition rate of type τ between nodes S_i and S_j	1/s
Ψ_{12}	Angular overlap hyper-volume of two nodes	
Ψ_{sym}	Swap-averaged angular overlap hyper-volume of two nodes	
ψ	Angle between the object's position and the ascending node	°
Ω	Right ascension of the ascending node	°
$\Delta\Omega$	RAAN bin width	°
ω	Argument of periapsis	°

1

Introduction

With the rapid commercialization and democratization of space brought about by the New Space economy, the market is witnessing a 21st-century race to orbit, shifting from government-driven missions to an arena increasingly open to private actors. One consequence is the opportunity for small companies and institutions to carry out space missions whose costs would have been prohibitive just a decade ago, especially thanks to ride-share launch providers. Yet the most striking change is the rise of planned large-scale satellite constellations, which are reshaping both the economic and physical landscape of near-Earth space.

Satellite constellations are groups of satellites working together as a system. Their main advantage over individual satellites is their ability to provide quasi-permanent global coverage, such that at any time at least one satellite is visible from any point on Earth, making them particularly suited for navigation, telecommunication, and broadband internet services. However, the strict low-latency requirements posed by such applications make Low Earth Orbit (LEO) and Medium Earth Orbit (MEO) their preferred orbital regimes, significantly adding to the congestion that already characterizes these regions. Starlink alone, with a population of more than 9600 active satellites, comprises about 55% of all in-orbit payloads[1] and, in January 2026 received approval from the Federal Communications Commision (FCC) to increase its fleet to a total of 15 000 satellites (FCC, 2026a). The latest FCC ruling on OneWeb, the second-largest broadband internet provider after Starlink and already comprising roughly 650 active satellites[2], reports approval to expand the constellation to 6372 satellites (FCC, 2023). Two of Jeff Bezos's companies, Amazon and Blue Origin, are likewise developing new mega-constellations: Amazon Leo (formerly Project Kuiper), with 7736 satellites (FCC, 2026b), and TeraWave, with 5408 satellites (Blue Origin LLC, 2026), respectively. In late December 2025, China submitted paperwork to the International Telecommunication Union (ITU) for several mega-constellations totaling nearly 200 000 satellites distributed over 3660 orbital planes (Jones, 2026). The shift is already visible in orbit: fig. 1.1 shows constellation satellites overtaking every other object in LEO within the last few years.

Contrary to what these numbers might suggest, the capacity of viable operational orbits is not unlimited, and significant efforts are required to ensure a safe and sustainable use of space. In this regard, inter-governmental bodies such as the Inter-Agency Space Debris Coordination Committee (IADC), the International Organization for Standardization (ISO), and the European Space Agency (ESA) have been publishing increasingly stringent guidelines and standards on space traffic management and debris mitigation, thereby driving advances in Space Situational Awareness (SSA), Space Surveillance and Tracking (SST), and Space Traffic Management (STM) technologies. Figure 1.2 puts that effort in context, with the number of tracked objects rising steadily since the beginning of spaceflight and sharply over the last decade.

[1] Counts of active (operational) Starlink satellites versus total active payloads, derived from the United States Space Command (USSPACECOM) Space-Track.org catalog, https://www.space-track.org/, accessed February 21st, 2026.

[2] See note 1.

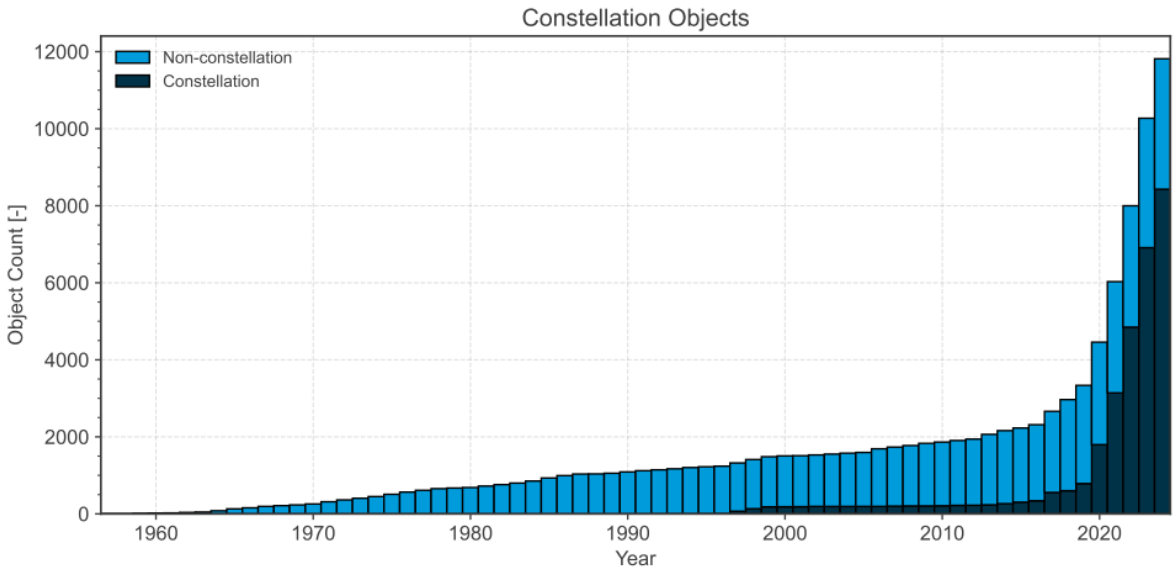

Figure 1.1: Evolution of the number of constellation and non-constellation objects in LEO (ESA, 2025).

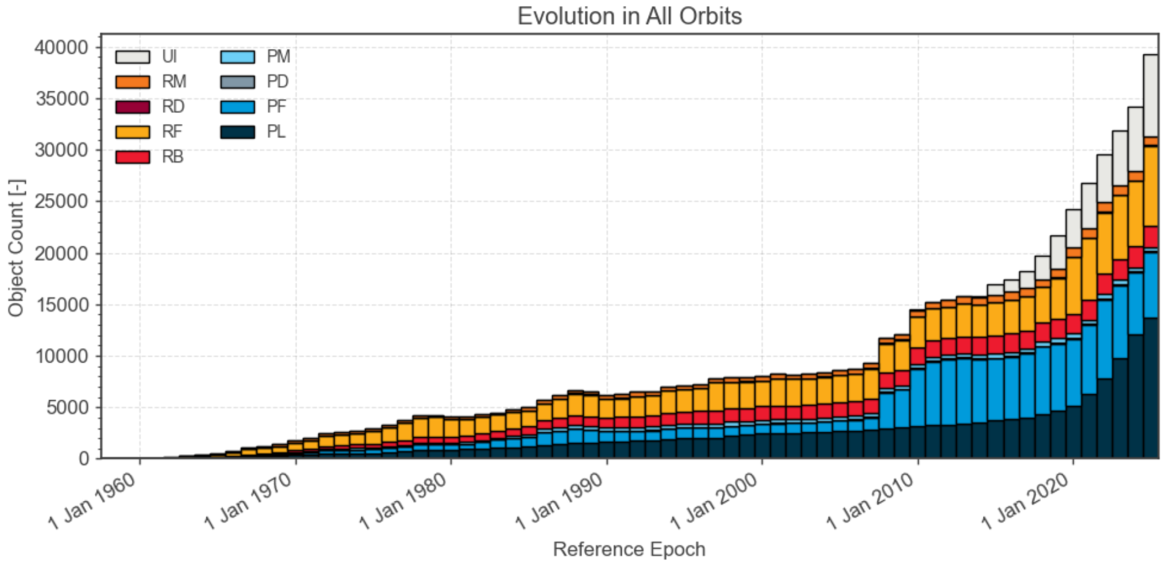

Figure 1.2: Number of tracked objects in Earth orbit over time (ESA, 2025).

The space environment is notoriously difficult to model. Public services track approximately 40 000 objects with diameters as small as 10 cm, but smaller particles such as Solid Rocket Motor (SRM) dust, paint flakes, ejecta, and sodium–potassium droplets constitute the vast majority, an estimated 1.2 million objects between 1 and 10 cm and 140 million between 1 mm and 1 cm, and can still pose significant threats to active satellites despite their small size (ESA, 2025). Because these particles cannot be tracked individually, their spatial distribution is typically described only in a statistical sense. Moreover, the motion of every object in space is influenced by a diverse range of perturbations, from atmospheric drag to solar radiation pressure to the complex gravitational field of the Earth and other celestial bodies. If left at the mercy of such perturbations, objects can drift away from their nominal trajectories and collide, giving rise to the notorious Kessler syndrome (Kessler & Cour-Palais, 1978), a scenario in which fragments from a collision further increase the local debris density, triggering a chain reaction of collisions

in a positive-feedback loop. Even when objects are maneuverable, close approaches are not uncommon and require careful planning and execution of Collision Avoidance Maneuvers (CAMs). Within this dynamic environment, satellite constellations play a major role: due to their large populations, their operations can be easily disrupted by the external environment, and in turn they pose a significant risk to other active satellites. While several studies (Arroyo-Parejo et al., 2021; Reiland et al., 2021; Ruan et al., 2024) focus on the interaction between constellations and the surrounding debris environment, interactions between satellites within the same constellation are often only marginally addressed, if not entirely neglected. This project aims to contribute to filling this gap by developing an accurate and fast model for the long-term assessment of constellation safety.

1.1. Motivation

Satellite constellations face challenges unseen in single-satellite missions. First, after being deployed from a launcher, satellites have to be raised from their injection orbit to their target orbit, spiraling up through a wide range of altitudes. Not all satellites are launched at the same time, but are generally deployed in phases, or as replacement of failed or de-orbited peers. Considering an average operational life of 5 years, more than 10 replacements per day would be necessary for the thousands of constellation satellites approved by the FCC as of 2022, generating substantial two-way traffic (Polli et al., 2023). Once in their target orbit, satellites must perform continuous station-keeping throughout their mission, and at end-of-life (EOL) they must be properly disposed of either by controlled re-entry and burn-up in the atmosphere for LEO missions, or by transfer to a designated graveyard orbit at higher altitudes. At every phase of their mission profile, satellites might cross paths with other objects, including members of their own constellation, and can be forced to execute a CAM. Polli et al. (2023) show that injection and disposal maneuvers through congested shells can carry collision probabilities up to and above $P_c = 10^{-3}$, depending on constellation density and crossing geometry, for instance, when a satellite from an upper shell performs Post Mission Disposal (PMD) and spirals through the lower shells of a multi-shell constellation (such as Starlink). It is therefore fundamental that a constellation satellite be equipped with a propulsion system that ensures maneuverability: a failed satellite that can no longer maneuver cannot rise to its target orbit, cannot perform station-keeping, cannot perform CAMs in case of conjunction events, and cannot execute a controlled re-entry, essentially turning into a loose cannon in orbit. Lewis et al. (2017) identify conjunctions involving failed satellites (92.5% of all intra-constellation conjunctions) as the dominant long-term self-induced collision driver in a satellite constellation.

To scale traffic management to the huge number of objects in these systems, the orbital maneuvers required to safeguard constellation infrastructure are increasingly assessed and executed autonomously. OneWeb's Flight Dynamics Service (FDS) continuously processes more than 50 000 Conjunction Data Messages (CDMs) per day generated by the 19th Space Defence Squadron, with high risk conjunctions ($P_c \geq 2 \times 10^{-5}$ or miss distance $\leq 2\,\mathrm{km}$ within a 48 h window to Time of Closest Approach (TCA)) triggering evasive maneuvers that further distort nominal trajectories (Yoon et al., 2023). Most operators still adopt a local-pairwise maneuver paradigm, where each satellite will only react to its most threatening upcoming conjunction; a strategy that fails to capture the intricate cascading phenomena emerging in autonomously operated constellations. When one satellite maneuvers, it reduces inter-satellite spacing with its neighbors, forcing them in turn to maneuver and propagating a domino-like response throughout the constellation. Chen et al. (2024) quantify this endogenous instability for Starlink, showing that cascaded effects can deplete up to 81.4% of the total budget allocated for collision avoidance (Starlink satellites carry fuel for approximately 350 CAMs SpaceX, 2022), and effectively waste 79.5% of the overall network's lifetime.

Finally, in the event of a catastrophic collision or explosion, the resulting debris cloud would disperse along the constellation orbits, imposing disproportionate inter- and intra-plane risks on the remaining satellites, in a process closely resembling a localized Kessler syndrome. Overall, the interplay between CAM cascades, failure-induced non-maneuverability, and high frequency

shell-crossing traffic creates a complex feedback system where operational decisions (thruster sizing, PMD policy, CAM thresholds) propagate through the constellation life-cycle, affecting not only individual satellite safety but also the systemic sustainability of the constellation as a whole.

1.2. Problem Statement

Assessing the endogenous collision risk of a constellation rigorously requires propagating every satellite and comparing all possible pairs, of which there are $N(N-1)/2$ for N objects, so that the cost grows with the square of the constellation size. For systematic parameter sweeps across failure rates, PMD compliance, launch profiles, or fragmentation events, this cost becomes prohibitive. In order to reduce the computational cost, classical geometric sieves are applied to pre-filter pairs that cannot intersect, progressively eliminating candidates based on apogee-perigee separation, minimum orbital intersection distance, and time-phase coincidence (Kerr & Sánchez Ortiz, 2021). These filters are effective for heterogeneous populations that span a wide range of altitudes and inclinations. For the endogenous risk problem, however, all intra-constellation pairs share nearly identical semi-major axes, rendering the geometric filters unfit. The time-phase filter is equally ineffective in the opposite direction: because inter-plane phasing is deliberately designed to avoid node crossings, it suppresses virtually every pair regardless of actual risk. Consequently, the sieve pipeline offers little practical reduction for intra-constellation conjunction screening. The suppression is also misdirected. Designed phasing is a property of satellites that are actively holding their slots, and those are not where the endogenous risk sits: it is carried by objects that have left the designed configuration, satellites that failed and now drift, satellites crossing shells while raising or de-orbiting, and fragments released by a breakup. A filter resting on nominal phasing discards pairs correctly for the population that poses little risk and cannot speak for the population that poses most of it.

For the pairs that survive screening, detailed pairwise collision probabilities must be computed. This is generally done by means of covariance-based methods (Alfano, 2005; Patera, 2001) and deterministic methods pioneered by Öpik (1951) and Wetherill (1967) for impacts of celestial bodies, and adapted by JeongAhn and Malhotra (2017) for circumterrestrial space. In the theory of Öpik and Wetherill, the orbits are assumed to be fixed in time, and the mean anomalies of target and field objects are assumed independent and uniformly distributed. Then, a collision rate for an ensemble of bodies is obtained by averaging the probability of collision between the orbits over the Euler angles. This is equivalent to Kessler (1981) gas-kinetic approach, where all objects within an altitude shell are assumed to be uniformly distributed over the longitude of the node and argument of periapsis. This assumption is well-motivated for Solar-System dynamics, but breaks down for Earth-orbiting satellites whose elements evolve rapidly due to Earth's oblateness, drag, and lunisolar perturbations. JeongAhn and Malhotra (2017) address this with a Monte Carlo treatment of the field object: its trajectory is forward-propagated under the full equations of motion and sampled at random epochs, and the collision probability is the average of the pairwise probabilities taken over that sample. This approach is substantially more accurate, but at the cost of many additional propagations, and it remains a deterministic, short-term assessment of a fixed constellation geometry, not a tool for long-term, stochastic, multi-scenario risk evaluation. Indeed, Reiland et al. (2021) applied the JeongAhn-Malhotra method for detecting self-induced collisions within mega-constellations, reporting a required computational time of 13 h and 40 h for OneWeb and Starlink, respectively, on the 200 CPUs of the University of Arizona's High Performance Computing cluster.

A comprehensive endogenous risk assessment must account for a range of stochastic drivers: random satellite failures, PMD failures, shell-crossing traffic during raising and de-orbiting phases, fragmentation events from both collisions and explosions, and cascading of collision avoidance maneuvers. Brute-force pairwise methods offer no natural mechanism for embedding these processes: one must manually prescribe which satellites fail, which fragment, and which change orbit, then rerun the full propagation for each scenario realization. Evolutionary models for the space environment represent a more natural solution, since they incorporate stochastic events alongside orbital propagation. High-fidelity tools propagate each object individually

and resolve every collision or fragmentation event explicitly. This category includes NASA's LEGEND (J. C. Liou, 2006; J. C. Liou et al., 2004), ESA's DELTA (Virgili, 2016), the British DAMAGE (Lewis et al., 2001), the German LUCA (Radtke et al., 2017), the Italian SDM (Rossi et al., 2009), the Japanese NEODEEM (Kawamoto et al., 2018), the French MEDEE (Di Costanzo et al., 2013), and the American MOCAT-MC (Jang et al., 2024). Each run produces a single stochastic realization of the debris environment, and a statistically meaningful picture requires a large number of Monte Carlo runs. Two fundamental limitations persist even with this approach. First, these models assess collision risk using spatial-density methods that assume all objects within a given volume element are uniformly distributed and equally likely to collide, such as the widely used CUBE method (J. C. Liou, 2006; J. C. Liou et al., 2003). This assumption is well-motivated for whole-population simulations where secular perturbations scatter objects over long timescales, but it is demonstrably invalid for constellation-only populations whose active satellites are locked to precise, designed orbits. Lewis et al. (2019) showed explicitly that the CUBE method can report collisions between satellites whose orbits never intersect: any two objects falling in the same cube at the same sampling instant are treated as candidates for collision, so when the altitude separation between two planes is smaller than the cube, pairs that stay kilometres apart are counted as if they could meet, and that achieving reliable estimates for constellation geometries would require cube sizes and time steps so small that a single Monte Carlo run would take days to weeks. Second, high-fidelity methods produce time series of object counts and collision events with no closed-form relationship between species, the classes into which the population is divided, such as active satellites, derelicts and fragments. The dynamics are effectively a black box: risk attribution requires post-hoc statistical analysis across many runs, and extracting systematic trends or identifying dominant risk drivers is neither direct nor efficient.

Source-sink models, such as MOCAT-SSEM (D'Ambrosio et al., 2023), address the efficiency limitation by grouping the population into species bins defined by orbital parameters and operational status, and evolving the bin populations via coupled ordinary differential equations governing transfer rates between species. For this reason, they are also popularly known as Particle-in-a-Box (PIB) models. Because no individual trajectories are integrated, these models run in seconds to minutes, enabling parameter sweeps, and their formulation naturally allows for equilibrium analysis and the identification of carrying-capacity thresholds. The NEtwork model for Space SustainabilitY (NESSY) of Wang et al. (2025) occupies an intermediate position: it advances the system using finite-difference equations governed by Poisson-jump stochastic events, so each run is a single realization and multiple runs are needed to recover mean behavior (analogous to Monte Carlo) but without orbit propagation or numerical integration, making it substantially faster than high-fidelity tools. Crucially, NESSY couples population dynamics to a network structure, enabling the direct application of network-theoretic metrics to identify risk drivers and systemic trends.

Despite these advantages, all existing source-sink models, including NESSY, discretize orbital space into altitude shells (and, in NESSY, inclination bands), assuming within each bin that objects are uniformly distributed over the Euler angles. This assumption is valid for century-scale simulations of the full debris population, because the secular precession of RAAN and argument of periapsis driven by J_2 randomizes the in-plane and cross-plane phases on timescales of weeks to months. For a decade-scale simulation of a constellation-only population the assumption fails on two counts: active satellites are strictly bound to their designed planes and phasing slots, and drifting failed satellites or freshly generated debris disperse in RAAN far more slowly than the precession rate on its own would suggest. What sets the pace is not the drift, which at these altitudes carries a single object around the full circle in a matter of months, but the *differential* drift across the population: objects released at nearly the same semi-major axis and inclination precess at nearly the same rate, so they stay clumped and spread only as fast as their spread in those two elements allows. The spread in those elements within a single constellation, or within a cloud released by one breakup, is narrow, and the resulting dispersion is correspondingly slow. Applying a shell-level uniformity assumption to such a population systematically misattributes collision risk, precisely the error identified empirically by Lewis et al. (2019) and Ruan et al. (2024) in the context of spatial volume-element methods. This constitutes an unresolved gap in the literature: no existing source-sink model incorporates plane-level resolution, simultaneously

binning by altitude shell, inclination, and RAAN, sufficient to respect the geometric structure of a large constellation. This thesis addresses that gap directly by adapting the NESSY framework to a constellation-only environment with explicit orbital-plane resolution, and demonstrating that the resulting model can accurately reproduce endogenous collision dynamics, at a computational cost that renders systematic parameter sweeps feasible.

1.3. Research Question

The problem statement identifies a gap at the intersection of three main shortcomings. First, evolutionary models have exclusively been applied to the full, heterogeneous space debris population; no existing source-sink model has been validated, or indeed formulated, for a constellation-only scenario. Because the phenomena governing a constellation's long-term evolution (e.g., structured orbital-plane geometry, repetitive shell-crossing traffic, and correlated failure and replenishment cycles) are fundamentally different from those of the broader debris environment, the adaptation is non-trivial and requires both architectural changes to the species taxonomy and sensitivity studies to identify which stochastic processes drive endogenous risk. Second, the spatial-uniformity assumptions embedded in all existing collision rate formulations have been shown to be invalid precisely for the geometries that characterize large constellations, systematically misattributing risk. Third, current high-fidelity methods produce per-object time series with no closed-form relationship between species populations: risk attribution is reduced to post-hoc statistical analysis across large Monte Carlo ensembles, making the identification of dominant risk drivers neither direct nor systematic. A modeling framework that naturally encodes inter-species relationships and exposes global-scale risk structure without requiring such auxiliary analysis would therefore constitute a distinct methodological advance. Together, these three shortcomings motivate the central research question:

> **To what extent can evolutionary space environment models assess long-term endogenous collision risk in satellite constellations?**

To aid in answering this question, four sub-questions have been formulated.

Q1 How can traditional evolutionary models be adapted for constellation-only populations?

Existing source-sink models bin the population by altitude shell and assume angular uniformity within each bin, an assumption only valid for century-scale simulations of the space debris population, but demonstrably invalid for decade-scale constellation dynamics. This question investigates the structural modifications to the bin taxonomy, the collision computation, and the distributional assumptions required to make an evolutionary model physically consistent with the constrained geometry of a large constellation.

Q2 Which stochastic processes need to be modeled to accurately capture the constellation's evolution?

A constellation evolves through a combination of operational events (satellite failures, launch replenishment, PMD outcomes, shell-crossing raising and de-orbiting maneuvers) and catastrophic events (fragmentation by collision or explosion). This question establishes, through sensitivity analysis, which of these processes contribute most significantly to endogenous risk, and therefore determines the minimum level of modeling fidelity required for reliable long-term projections.

Q3 How does the network formalism compare to traditional models in accuracy and computational efficiency, and what unique constellation-level insights does the network approach provide?

Beyond reproducing aggregate collision statistics, the network representation naturally encodes species relationships as directed edges, exposing systemic risk drivers through

graph-theoretic metrics such as node centrality and carrying capacity. This question assesses whether such a formalism can match or surpass classical source-sink and semi-deterministic methods in predictive accuracy at a comparable cost, while delivering global risk insights that are structurally inaccessible to single-node models.

Q4 How well does the stochastic dynamic model generalize to different constellation geometries and mission profiles?

The operational value of the model depends on its capacity to produce reliable risk assessments across a range of configurations (e.g., varying shell altitudes, inclination patterns, failure rates, and PMD compliance levels) without requiring recalibration for each scenario. This question evaluates the model's robustness and the stability of its network-level diagnostics under different mission profiles.

1.4. Outline of the Report

These questions are answered by building the Plane-Resolved Integrated Stochastic Model, or PRISM: a model that resolves a constellation's population by orbital plane, evolves it stochastically over a mission horizon, and reads the resulting risk as a network. The report is arranged in the order in which that model is built, applied and tested. The first two chapters set the ground it stands on. Chapter 2 reviews the state of the art in two halves, the conjunction screening and pairwise probability methods used to assess a constellation over days or weeks, and the evolutionary and source-sink models used to project a whole population over decades, and shows that neither answers the question posed here: the first cannot span the horizon, the second cannot resolve the orbital plane. Chapter 3 then fixes the astrodynamics the rest of the work relies on, the reference frame and the three parametrizations of a state that recur throughout, and identifies the two perturbations retained, Earth's oblateness and atmospheric drag, together with the secular rates they produce.

The model itself occupies the next three chapters. Chapter 4 defines the orbit-bin lattice over altitude, inclination and RAAN, the species an object can belong to, and the stochastic balance equations that move objects between them; its central result is a closed-form collision rate between two bins, derived from the phase-space density and reduced to a form cheap enough to evaluate over an entire lattice. Chapter 5 turns that formulation into a simulator, describing the step ordering, the aggregation that makes collision sampling affordable, and the conditioning scheme that extracts converged statistics from events too rare for a plain Monte Carlo to see, and it verifies the implementation against MOCAT-MC. Chapter 6 asks how finely that lattice and its time step must be resolved before the answers can be trusted, and quantifies what a coarse choice costs.

The last three chapters use the model. Chapter 7 recasts the population as a network whose nodes are orbit bins and whose links are the interactions between them, and defines the topological and dynamical measures that identify which regions carry the risk and which channels transport it. Chapter 8 applies all of it to Starlink, OneWeb and Iridium-NEXT, following each constellation's population over its own horizon, resolving the risk over the orbital planes, measuring how much traffic a shell can absorb before it stops clearing its own debris, and sweeping the collision-avoidance, disposal and launch policies an operator controls. Chapter 9 answers the research questions and sets out what should be built next.

2

State of the Art

This chapter provides an overview of the state of the art in evolutionary models for the space environment and collision-rate computation. First, an overview of constellation dynamics is provided in section 2.1, with examples from the Starlink and OneWeb reference missions. Then, the current models for the evolution of the space debris environment are described in detail in section 2.2, with emphasis on their different approaches to collision detection. In this regard, the different formulations for collision rate are reviewed in section 2.3. Section 2.4 provides a detailed overview of NESSY as well as the current advances in network theory for space safety assessment.

2.1. Constellation Dynamics

2.1.1. Constellation Geometries

A satellite constellation is a coordinated group of satellites acting together as a system, whose orbits are deliberately designed to achieve a desired spatial coverage. Satellite constellations can be broadly partitioned into two families: symmetric, or Walker, constellations, which exploit a regular geometry to provide near-uniform global or zonal coverage from circular orbits, and non-Walker configurations, which sacrifice spatial uniformity to extend coverage to regions inaccessible to circular low-inclination constellations (Wertz & Larson, 1999).

Walker Constellations

The dominant design paradigm for large LEO constellations is the one formalized by (Walker, 1984), who introduces the notation i: T/P/F, where i is the common orbital inclination, T is the total number of satellites, P is the number of equally spaced orbital planes, and F is an integer phasing factor specifying the in-track phase offset between satellites in adjacent planes. The phasing factor F controls the angular offset between adjacent planes and influences coverage continuity during initial deployment; it also determines which inter-plane encounter geometries are possible.

The Walker Star configuration employs a near-polar inclination, typically between 85° and 90°, with ascending node spacings distributed uniformly over 180° (Walker, 1984). As a consequence, planes in one half of the configuration are co-rotating while those in the opposing half counter-rotate, meeting at a longitude band known as the counter-rotating seam (see fig. 2.1). At this seam, satellites on adjacent planes travel in nearly opposite directions, producing relative encounter velocities approaching twice the orbital speed, far exceeding those between co-rotating planes Wertz and Larson (1999). From a coverage standpoint, the near-polar inclination delivers continuous polar coverage inaccessible to low-inclination configurations, such as Walker Delta, at the cost of reduced uniformity at equatorial latitudes.

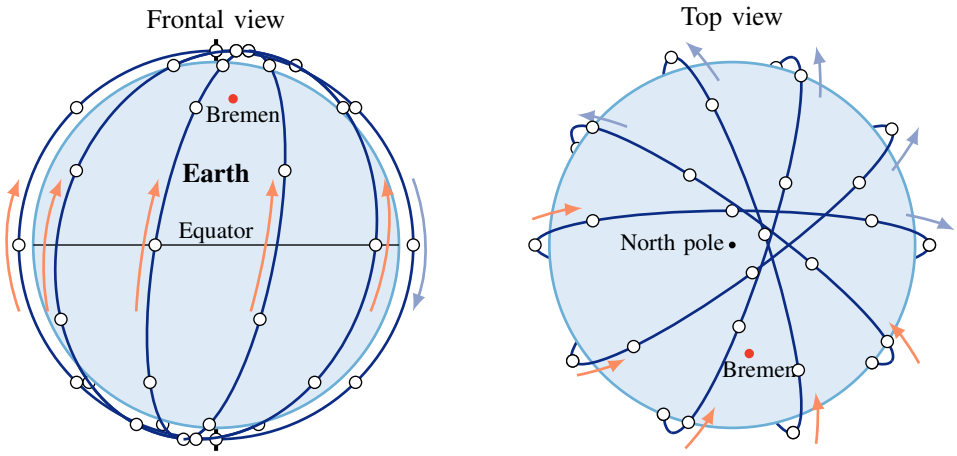

Figure 2.1: Illustration of a 80°: 40/5/1 Walker star constellation (Matthiesen et al., 2023).

In the Walker Delta configuration, the RAAN of each plane is uniformly distributed over the full 360°, enforcing rotational symmetry about Earth's polar axis (Wertz & Larson, 1999). This regularity ensures that satellite footprints tile the sphere with near-uniform density at all longitudes, so that coverage holes and revisit times are primarily driven by inclination rather than local longitude: a Walker Delta shell provides continuous coverage equatorward of i when the inter-plane spacing and minimum elevation angle permit, while coverage degrades progressively toward the poles (see fig. 2.2). The Global Positioning System (GPS) constellation (55°: 24/6/2) and the Galileo constellation (56°: 27/3/1) are canonical Walker Delta implementations at medium Earth orbit, providing the navigational redundancy and coverage continuity that global positioning requires (Wertz & Larson, 1999). For large broadband constellations such as Starlink and OneWeb, the limitations of a single Walker shell at a fixed inclination have motivated multi-shell hybrid architectures combining layers at different altitudes and inclinations to simultaneously address equatorial latency, polar coverage, and capacity.

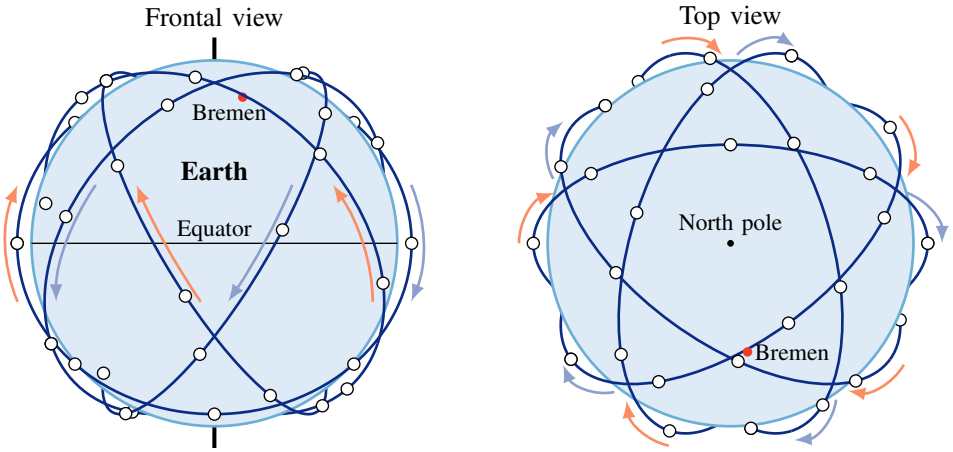

Figure 2.2: Illustration of a 60°: 40/5/1 Walker delta constellation (Matthiesen et al., 2023).

Non-Walker Constellations

Beyond the Walker family, a broad class of non-Walker designs can be organized under the name of flower constellations, introduced by Mortari et al. (2004), which includes all satellite constellations whose members share a common repeating ground track in an Earth-fixed frame. This framework encompasses a wide family of existing designs, including the Molniya orbit (12-hour period, eccentricity ≈ 0.65, inclination $\approx 63.4°$) and the Tundra orbit (24-hour period, eccentricity ≈ 0.27, same critical inclination), both of which exploit the prolonged apogee dwell

time at high inclination to provide near-continuous coverage of polar and high-latitude regions with a small number of satellites. The trade-off is a strongly non-uniform spatial density: because satellites spend the large majority of their period near apogee, the uniform mean anomaly distribution that underlies the Walker coverage analysis does not hold, complicating both design optimization and risk assessment.

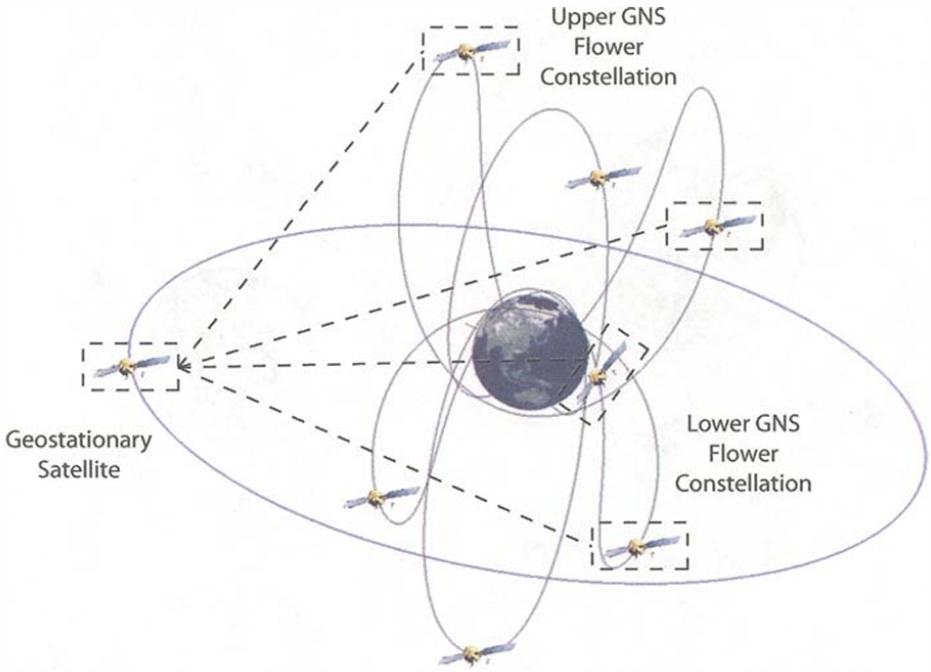

Figure 2.3: Example flower constellation in the Earth-centered fixed frame (Mortari et al., 2004).

2.1.2. Reference Missions

The three missions examined in this section are selected to span the principal large constellation designs: Starlink represents a Walker Delta multi-shell architecture, OneWeb a near-polar Walker Star architecture at medium LEO altitude, and Iridium the historically first large-scale commercial LEO Walker Star, whose 2009 collision event provides the unique observational benchmark for model validation.

Starlink

SpaceX's Starlink is the largest commercial satellite constellation ever authorized. Its second-generation system, authorized by the FCC in January 2026, covers up to 15 000 satellites organized across nine orbital shells between 340 and 485 km at inclinations ranging from 28° to 96.6°, complementing a first-generation system of approximately 4400 satellites near 550 km (FCC, 2026a). The multi-shell design serves three concurrent objectives: inclination determines the geographic coverage band, population density sets the capacity for that band, and the deliberately low operating altitudes ensure passive atmospheric disposal within years of a satellite failure. The Gen2 deployment explicitly exploits this last property by testing satellites at the lowest shells, where drag rapidly clears non-responsive objects, before raising healthy units to their assigned altitude. From a safety standpoint, SpaceX triggers collision avoidance maneuvers at a probability threshold of one in one million (two orders of magnitude below common industry practice) and the authorization conditions continued deployment on ongoing disposal-failure-rate reporting (FCC, 2026). These documented failure, disposal, and conjunction avoidance rates constitute the primary empirical inputs for the Starlink component of the network model developed in this thesis.

OneWeb

OneWeb operates a near-polar broadband constellation at 1200 km, targeting high-latitude connectivity unreachable by lower-inclination Walker Delta configurations. Its Phase 1 configuration, authorized by the FCC in 2023, consists of 716 satellites split between 12 near-polar planes at 87.9° inclination and 8 moderately inclined planes at 55° (FCC, 2023). Each satellite, with a mass of approximately 150 kg and equipped with a Hall Effect electric propulsion system, is injected at 450 or 600 km and orbit-raised to its operational altitude over days to weeks via one of two ascent profiles: a direct raise with an optional recovery pause near 900 km, or an indirect raise that includes an extended RAAN-phasing drift at approximately 600 km (Yoon et al., 2023). The resulting transit sub-population occupies lower shells, including those used by Starlink, for operationally significant durations. At 1200 km, atmospheric drag is negligible, so end-of-life disposal relies entirely on active propulsion; every OneWeb satellite additionally carries a magnetic grappling fixture to enable Active Debris Removal (ADR) should it become unresponsive before disposal (Yoon et al., 2023).

Iridium

The Iridium constellation, introduced in 1998 as the first commercial LEO system offering global coverage, operates 66 active satellites in 6 planes of 11 satellites each at 780 km altitude and 86.4° inclination (Iridium Satellite LLC, 2013). This configuration was preserved in Iridium's second generation replacement programme, Iridium NEXT, completed in 2019. On 10 February 2009, the operational Iridium 33 collided with the derelict Cosmos 2251 at approximately 790 km and a relative velocity of roughly 11.6 km/s, generating over 1800 trackable debris fragments and producing the first confirmed hypervelocity satellite–satellite collision on record (Anz-Meador & Liou, 2010).

2.2. Evolutionary Models for the Space Environment

This section surveys the principal classes of models used to simulate the long-term evolution of the orbital debris environment. Semi-deterministic models propagate individual or representative objects forward in time, computing each object's full trajectory at high computational cost. Source-sink models replace per-object propagation with coupled rate equations governing the flow of objects between orbital classes, trading fidelity for tractability and enabling rapid Monte Carlo campaigns over century-scale horizons. Continuum-based models take this abstraction further, treating the population as a smooth spatial density field governed by partial differential equations that admit analytical results, such as carrying capacity and equilibrium stability conditions. Together, these three approaches span a spectrum from high-resolution individual-object tracking to tractable population-level analysis, and their respective strengths and limitations directly motivate the network-based framework developed in this thesis.

2.2.1. Semi-deterministic Models

Semi-deterministic evolutionary models project future orbital populations by combining deterministic orbital propagation of discrete objects with Monte Carlo sampling for stochastic source-sink events, such as future launch traffic, fragmentation, and mitigation compliance. At each timestep, all discrete objects are propagated forward, collision probabilities are evaluated pairwise or through spatial density estimates, fragmentation events are triggered stochastically via break-up models, and mitigation rules are applied. This paradigm has historically served as the foundation for long-term space environment forecasting.

The first generation of full-scale three-dimensional tools, namely NASA's LEGEND (J. C. Liou, 2006; J. C. Liou et al., 2004), ESA's DELTA (Virgili, 2016), and the Italian SDM (Rossi et al., 2009), established the standard for century-scale environment projection across LEO, MEO and Geostationary Orbit (GEO). LEGEND (LEO-to-GEO Environment Debris model), developed by the NASA Orbital Debris Program Office, extends the one-dimensional EVOLVE model (Yates & Jonas, 1995)

by adding latitudinal and longitudinal resolution, critical for modeling Sun-synchronous orbits and GEO populations where spatial density varies heavily by longitude and latitude. DELTA (Debris Environment Long-Term Analysis) employs a target-centered, flux-based collision prediction scheme based on Klinkrad (1993) theory: rather than evaluating every object pair, each target accumulates an orbit-integrated mean flux from three-dimensional inertial control volumes, and a Poisson distribution determines whether a collision occurs within each time interval. Fragments from collisions or explosions are placed at the point of peak flux along the target orbit. The SDM (Space Debris Mitigation) addresses the propagation cost explicitly by employing three regime-specific propagators: a faster lower-fidelity propagator for fragment clouds, a high-fidelity propagator which propagates the full dynamics of the objects under geopotential, solar radiation pressure and lunisolar perturbations, and an analytically averaged propagator for GEO objects. Collision evaluation is handled either via Öpik's analytical method for near-circular LEO orbits with eccentricities below 0.1, or via the CUBE method for eccentric or high-altitude orbits.

To improve spatial resolution and computational efficiency, subsequent models introduced more sophisticated software architectures. DAMAGE (Debris Analysis and Monitoring Architecture for the GEO Environment), developed at the University of Southampton, specifically targets GEO, where standard fixed-grid bins cannot resolve the localized density peaks caused by satellite bunching in longitude deadband slots. To address this, DAMAGE introduces a dynamic control volume consisting of constant-radius spheres ("bubbles") centered on large objects, limiting the number of computational cells to the number of tracked objects and reducing both memory requirements and evaluation time (Lewis et al., 2001). LUCA (Radtke et al., 2017) and MEDEE (Di Costanzo et al., 2013) instead focus on modularity and propagation fidelity. LUCA (Long Term Utility for Collision Analysis), developed by the Institute of Space Systems at TU Braunschweig, is structured around an object-oriented Fortran architecture in which physical process plugins for propagation, collision rate computation, breakup, PMD, ADR, and constellation traffic can be independently activated or replaced. This makes it a flexible testbed for scenario construction, though execution time scales directly with the number and complexity of active plugins. MEDEE (Modelling the Evolution of Debris in the Earth's Environment), developed by Centre National d'Études Spatiales (CNES), uses STELA as its orbital propagator: a semi-analytical integrator that averages short-period perturbation analytically to permit large integration steps. Despite this efficiency, propagating the full six-element state of every fragment over long simulation horizons across many Monte Carlo runs remains costly, and MEDEE requires a massively parallel compute environment (360 cores, 24 GB RAM, 4 Tflops/s) to operate within practical time constraints.

More recent frameworks have explicitly focused on mega-constellation traffic. NEODEEM (Near-Earth Orbital Debris Environment Evolutionary Model), developed jointly by Kyushu University and Japan Aerospace Exploration Agency (JAXA), propagates all objects larger than 10 cm with a 5-day timestep semi-analytical propagator (Kawamoto et al., 2020). The model was extended to handle modern constellation traffic, including Walker delta deployment patterns and nano/micro-satellite populations with limited or no propulsion (Kawamoto et al., 2018). Collisions are evaluated by computing the overlap probability between error spheres centered on each object, weighted by the relative velocity and effective collision cross-section (Kawamoto et al., 2020). Due to the low frequency but high impact of catastrophic collision events, 100 independent Monte Carlo runs are required to extract statistically significant trends, each taking approximately 60 hours of elapsed time on JAXA's supercomputer, consuming over 560 000 core-hours per study (Sugita et al., 2021). The model has been applied to investigate environment capacity, i.e., the number of additional satellites that can be inserted at a given altitude while maintaining long-term sustainability, revealing a critical sensitivity to both PMD compliance and operational altitude: constellations operating below 60% PMD compliance drive persistent long-term population growth regardless of other mitigations. At the same time, deployments above 900 km trigger collision cascades that do not decay on human timescales (Kawamoto et al., 2021). Disposal strategy is equally relevant: gradual de-orbit via electric propulsion or drag augmentation complies with the 5-year rule yet increases the short-term collision probability due to the extended dwell time through congested shells (Kawamoto et al., 2018).

Finally, MOCAT-MC (MIT Orbit Capacity Analysis Tool), developed at Massachusetts Institute

of Technology (MIT), is an open-source model that pushes Monte Carlo scalability to its limit by replacing semi-analytical propagation with a purely analytical Simplified General Perturbations (SGP4) propagator. The reduced per-object propagation cost allows tens of millions of trackable objects to be simulated over a 200-year horizon on a single processing thread, extending the previous state-of-the-art of approximately 400 000 objects by roughly two orders of magnitude (Jang et al., 2024). This comes at the cost of lower accuracy than semi-analytical alternatives. More fundamentally, the object-by-object Monte Carlo structure remains prohibitively expensive for systematic parameter sweeps across failure rates, PMD compliance levels, or launch profiles, a limitation that motivates the source-sink and density-based formulations discussed in the following subsections.

Table 2.1: Comparison of Semi-deterministic Models

Model	Target Regime	Propagator	Collision Detection
LEGEND	LEO to super-GEO	Semi-analytical, numerical	CUBE
DELTA	LEO to GEO	Semi-analytical	Klinkrad
SDM	LEO to GEO	Semi-analytical, numerical	Orbit-Trace, CUBE
DAMAGE	2000 km to super-GEO	Semi-analytic	Object-centered
MEDEE	LEO to GEO	Semi-analytical	CUBE
LUCA	LEO to GEO	Plugin-dependent	Orbit-Trace, CUBE
NEODEEM	LEO	Semi-analytical	Pairwise
MOCAT-MC	LEO	Analytical	CUBE

Table 2.1 provides a comparative overview of the models discussed in this section. Despite their different architectural choices, these models are broadly consistent in their physical assumptions and event modeling. The most systematic comparison was conducted as IADC Action Item 27.1 (J. C. Liou et al., 2013), in which SDM, DELTA, LEODEEM (the low-Earth precursor to NEODEEM), LEGEND, and DAMAGE are run under a strictly controlled scenario sharing identical initial population, launch traffic cycle, PMD success rate, and a common breakup model. Across 725 Monte Carlo runs, all six models predicted a net population increase of approximately 30% over 200 years, with catastrophic collisions occurring every 5 to 9 years, confirming the instability of the LEO environment even under optimistic mitigation assumptions. This mutual consistency, however, is not equivalent to independent validation: since the future state of the environment is unobservable, true validation is structurally impossible for long-term projections, and cross-model agreement under identical conditions is the closest available substitute. Partial historical comparison is possible (e.g., models such as LEGEND include a backward simulation component reproducing the known catalog from 1957 until 2001), but this is limited to trackable objects above approximately 10 cm, with smaller objects remaining less accessible to direct observational verification.

2.2.2. Source-sink Models

An alternative approach consists of discretizing the physical space or a suitable phase-space into bins and defining a set of coupled Ordinary Differential Equations (ODEs) which describe the evolution of the object count in each bin, under the effect of sources (new launches, fragmentation events, etc.) and sinks (natural decay, PMD, ADR)

$$\frac{\mathrm{d}x_k}{\mathrm{d}t} = \sum_{l \neq k} \vartheta_{lk} x_l - \sum_{l \neq k} \vartheta_{kl} x_k + S_k - L_k \tag{2.1}$$

where x_k is the object count in bin k, ϑ_{lk} is the transfer rate from bin l to bin k due to slow dynamics, and S_k and L_k are the fast source and sink terms. Source-sink models (also known as

Particle-in-a-Box models) do not resolve each trajectory individually and thus do not require orbit propagation, thereby considerably improving computational efficiency at the cost of accuracy.

The source-sink paradigm was first applied to the full debris environment by Rossi et al. (1994), who integrated a system of 150 coupled nonlinear first-order ODEs with 15 altitude shells and 10 logarithmic mass bins. Collision rates in each shell are obtained by multiplying the average intrinsic collision probability, computed from a catalog of 2700 actual orbits using a method adapted from Öpik's; atmospheric drag is accounted for through shell-specific residence times. Rossi et al. (1997) subsequently implements the STochastic Analog Tool (STAT), a fully discretized finite-difference model whose per-run cost is independent of the total population count and thus more efficient for century-scale projections. The code was already run with multiple random-number-generator seeds to characterize the spread of outcomes, recognizing that the stochastic nature of rare, high-energy collision events cannot be adequately represented by a single deterministic trajectory.

This observation points to a structural limitation of ODE-based source-sink models: because they integrate continuous rate equations, they return non-integer expected population counts. For large, slowly evolving debris populations, this is entirely appropriate: the continuous value is the expectation of the Poisson-distributed population, and the law of large numbers ensures it closely tracks any single realization. The approximation becomes less defensible in the two regimes most relevant to the present work: constellations, where the population of each node is measured in tens to hundreds of satellites and individual failure or collision events are consequential, and short time-horizons, where rare-event statistics cannot be averaged away. Fanto et al. (2023) addresses this limitation by recasting the source-sink framework as a stochastic jump process: borrowing the Gillespie algorithm from physical chemistry, the number of occurrences of each process (launches, PMD, drag decay, and collisions) is drawn from a Poisson distribution at each time step, and a large ensemble of such realizations is aggregated to recover the full probability distribution of population outcomes.

The most significant recent advance in the source-sink family is the MOCAT-SSEM developed at MIT. D'Ambrosio et al. (2023) introduce a multi-shell, multi-species ODE model in which the LEO region between 200 and 2000 km is divided into 36 spherical shells of 50 km thickness, and the population is partitioned into active satellites, derelict satellites, and debris. Collision rates follow the kinetic theory of gases, with species-dependent impact parameters and the NASA SBM used to compute fragment yields; atmospheric drag provides the sole inter-shell flux for non-manoeuvring objects. Crucially, D'Ambrosio et al. define *orbital capacity* as the stable equilibrium point of the dynamical system: the launch rate that maximizes the number of active satellites at equilibrium while keeping the fractional collision loss below a specified threshold. This equilibrium-based capacity metric, derived from the eigenvalue structure of the Jacobian of Eq. (2.1), provides a rigorous and computationally tractable criterion for assessing the sustainability of any given launch cadence.

Gusmini et al. (2024) extends MOCAT-SSEM to account for shell-crossing objects, dropping the assumption that newly launched objects are instantaneously added to their target bin. The active satellite population is subdivided into three sub-species (operating, raising, and deorbiting satellites), and the inter-shell flux associated with orbit transfer maneuvers is modeled explicitly. Assuming a low-thrust continuously applied control in the velocity direction, the secular variation of the semi-major axis as a function of time follows an analytically derived expression, which is inverted to give the epoch at which the transfer crosses each shell boundary and hence the residence time in each shell during transit. This residence time then populates the inter-shell transition matrices that couple the shells during raising and de-orbiting phases. The central finding of Gusmini et al. is that ignoring orbit-transfer fluxes systematically underestimates the number of collisions, particularly in the heavily trafficked injection-orbit shells, providing both physical justification and direct numerical evidence for explicitly modeling transiting populations.

2.2.3. Continuum-based Models

The continuum approach treats the fragment population as a fluid whose density evolves under the influence of orbital perturbations, replacing the propagation of individual trajectories with the integration of a single partial differential equation over the relevant phase space

$$\frac{\partial n}{\partial t} + \nabla \cdot (n\boldsymbol{f}) = \dot{n}^{+} - \dot{n}^{-} \tag{2.2}$$

where $\nabla \cdot (n\boldsymbol{f})$ models the slow, continuous forces acting on the system (such as orbital perturbations) and the right-hand side represents the fast, discontinuous, source ($\dot{n}^{+}$) and sink ($\dot{n}^{-}$) events of the system. The continuum approach is equivalent to a PIB model with bins of infinitesimal volume: integrating over the volume of a discrete bin and applying the divergence theorem converts the divergence term into a sum of fluxes across the bin boundaries, yielding Eq. (2.1). While the source-sink approach approximates the density field as piecewise constant within discrete bins, the continuity equation maintains a continuous density field over the entire domain.

The earliest fully analytical model of this kind applicable to LEO was introduced by McInnes (1993), who formulates Eq. (2.2) in 1D and applies the method of characteristics to obtain a closed-form solution for the long-term evolution of a fragmentation cloud under atmospheric drag only. The method of characteristics converts the Partial Differential Equation (PDE) into a set of ODEs along characteristic curves, each of which tracks the density as if following a fluid parcel through phase space; the analytical tractability depends on the simplicity of the assumed dynamics, which in McInnes's formulation required quasi-circular orbits and an exponential atmosphere model. In a later work, McInnes (2000) extends the same formalism to large nanosatellite constellations, deriving the continuity equation for the mean spatial number density of a constellation under drag decay, on-orbit failures, and the continuous deposition of new satellites. In particular, he solved for the asymptotic steady-state distribution and posed the inverse problem of determining the deposition rate required to maintain a desired constellation density, thereby establishing the concept of a carrying capacity set by the balance between replenishment and drag-induced loss.

Letizia et al. (2015) operationalized this framework by coupling it to the NASA SBM to initialize the density of the debris cloud at band formation, and demonstrated that the area-to-mass ratio distribution among fragments, which governs the differential drag decay of different subpopulations, could be handled by partitioning the cloud into independent bins, each propagated numerically under its own characteristic equation. Letizia, Colombo, and Lewis (2016) then extended the method of characteristics to multiple dimensions in Keplerian element space, making it possible to model the early dispersal phase driven by Earth's oblateness analytically, and simultaneously introduced eccentricity as an additional phase-space variable, extending validity below 800 km where the circular-orbit assumption breaks down. The Starling suite (Frey et al., 2019) estimates the non-linear evolution of densities by propagating the characteristics numerically under any chosen force model and recovering the density field through a Gaussian Mixture Model fit.

All these formulations were originally designed to propagate the density of a single fragmentation cloud in isolation. Colombo et al. (2016) made the first step toward applying the continuum framework to the global debris population by superimposing the evolving cloud density onto a background spatial density. The culmination of this line of work is the probabilistic debris evolutionary model of Giudici et al. (2024), which embeds the density-based fragment propagation within a full end-to-end environment propagator. The intact object population is classified into payloads, rocket bodies, mission-related objects, and individual constellation species, each characterized by its own mission profile, post-mission disposal rate, and explosion rate, and propagated individually. Fragmentation debris are instead described through a continuous density distribution and evolved through a hybrid scheme: the background fragment population is propagated via the method of characteristics in the (a, e, i) space, while freshly generated fragmentation clouds are handled by a Finite Volume Method that integrates the continuity equation analytically over each control volume, exploiting the sparsity of the resulting ODE system for efficiency. Explosion and collision probabilities are estimated at each time step for every intact object, fragmentation events are triggered stochastically, and the resulting fragment

density is immediately fed back into the propagating population. Applied to the LEO environment over a 200-year horizon, the model provides quantitative projections of debris population growth as a function of launch traffic and post-mission disposal compliance, demonstrating the viability of density-based methods as computationally efficient alternatives to fully deterministic environment propagators.

2.3. Collision Rate Computation

This section provides an overview of the most widely used methods for computing collision rates in long-term evolutionary models. Importantly, the goal of these methods is to estimate a statistical, time-averaged collision rate, a fundamentally different problem from computing the collision probability for a specific conjunction event. Methods such as those developed by Patera (2001) and Alfano (2005) condition on a predicted close approach and propagate the state covariance of both objects to the TCA, yielding an event-specific probability that requires full knowledge of the positional uncertainty at the encounter. While an all-vs-all pairwise assessment is theoretically possible, its cost scales quadratically with the number of objects and becomes infeasible for large populations such as constellations, as already discussed in section 1.2. The methods presented here instead produce statistical estimates of the collision rate and can be divided into two families:

- *Orbital intersection methods* determine the Minimum Orbital Intersection Distance (MOID) between two orbits under simplified assumptions about the distribution of orbital elements. If the MOID falls below the combined hard-body radii of the objects, a collision rate is computed as the time-averaged probability of geometric intersection.
- *Spatial density methods* discretize the physical or phase space into bins. The local density of objects within each bin is used to compute a collision rate drawing from the kinetic theory of gases, where the rate is proportional to the product of the local object densities, the relative velocity, and the effective collision cross-section.

These two families did not develop in isolation: they historically co-evolved as successive attempts to overcome their respective limitations, and remain closely related in their underlying assumptions.

2.3.1. Pairwise Orbital Intersection Methods

The classical way to evaluate collision probabilities between two circular Keplerian orbits was pioneered by Öpik (1951) in the context of collisions between celestial bodies. The method is essentially divided into two steps: first, the semi-major axis a, eccentricity e, and inclination i of the orbits are assumed fixed in time; if the two orbits are sufficiently close in space i.e. they intersect, or their MOID is within the combined hard body radii of the objects, then a collision can occur. Assuming that the mean anomalies M are independent and uniformly distributed, the probability of collision P_c is proportional to the time Δt spent by each object in the region around the orbits' intersection points

$$P_c = \frac{2\Delta t}{T_1 T_2} = \frac{2 r_c v_{\text{rel}}}{\|\boldsymbol{v}_1 \times \boldsymbol{v_2}\| T_1 T_2} \tag{2.3}$$

where r_c is the collision radius, $\boldsymbol{v}_i$ and T_i are the velocity vectors and orbital periods of the objects, respectively, and $v_{\text{rel}} = \|\boldsymbol{v}_1 - \boldsymbol{v}_2\|$ is the relative velocity of collision. LUCA and SDM implement this method under the name of Orbit-Trace. Since Orbit-Trace does not take into consideration the actual position of the objects along their path, it systematically overestimates collisions of objects on synchronized orbits (Radtke et al., 2013).

The second step consists of assuming a uniform distribution over RAAN Ω and argument of periapsis ω, allowing for the calculation of the long-term collision rate of an ensemble of bodies

on a target orbit as the average of all possible encounter geometries. The method is used by Wetherill (1967) to estimate collisions between asteroids and comets, and Kessler and Cour-Palais (1978) for collisions between artificial satellites, who generalizes it to non-circular orbits and solves singularity problems associated with tangential approaches.

These assumptions are valid for deep-space phenomena, for which the orbits can be considered constant in time, and the time horizons considered are sufficiently large to justify a uniform spread over Euler angles due to secular precession. However, they fail to model objects subject to strong dissipative forces, such as Near-Earth Objects (NEOs), objects with non-uniform precessing rates of Ω and ω, objects in mean-motion resonance with each other, such as in satellite constellations, or when the considered time-scales are too short. Vokrouhlický et al. (2012) formulates an Öpik-like method accounting for the large oscillations in eccentricity and inclination and the non-linear evolution of RAAN and argument of periapsis found in highly-inclined and highly-eccentric orbits. JeongAhn and Malhotra (2017) relaxed the fixed-orbit assumption by propagating a set of clones drawn from the actual perturbed dynamics and summing their individual Öpik-like collision probabilities, averaged by the clone multiplicity.

The method of averaging the pairwise collision probability over a population of field objects has been generalized by Kessler (1981) in the form of his well-known gas-kinetic formula

$$\dot{\eta} = A_c n_{\boldsymbol{r}} v_{\text{rel}} \tag{2.4}$$

where $\dot{\eta}$ is the rate of impacts on a target object of collisional cross-sectional area A_c traversing a control volume characterized by a spatial density of impactors $n_{\boldsymbol{r}}$. Given a, e, and i of the target orbit, and assuming a uniform distribution of Ω, ω, and M, Kessler obtains the spatial density as a function of radial distance r and latitude β only

$$n_{\boldsymbol{r}}(r,\beta) = \frac{1}{2\pi^3 r a \sqrt{(\sin^2 i - \sin^2 \beta)(r - r_p)(r_a - r)}} \tag{2.5}$$

where r_p and r_a are the perigee and apogee radii of the target orbit, respectively. The expression is purely geometric because it describes one object: it is the fraction of an orbital period that the object spends per unit volume around (r, β), and it integrates to unity over the volume its orbit can reach. The number of impactors enters by scaling, a population of N objects on that orbit having spatial density $N n_{\boldsymbol{r}}$. The collision rate in all volume elements accessible to two objects with spatial densities $n_{\boldsymbol{r}}^{(i)}$ and $n_{\boldsymbol{r}}^{(j)}$ is

$$\dot{\eta}_{ij} = \int_V A_c n_{\boldsymbol{r}}^{(i)} n_{\boldsymbol{r}}^{(j)} v_{ij} \, \mathrm{d}V \tag{2.6}$$

where v_{ij} is the relative velocity between the two objects (Frey & Colombo, 2021; Kessler, 1981). This formulation lays the groundwork for the density-based methods, which will be addressed in the following subsections.

2.3.2. Spatial Density and Flux Methods

The CUBE method (J. C. Liou et al., 2003) overcomes the limitations of Öpik's method and Orbit-Trace, and builds directly on top of Kessler's gas-kinetic approach. Contrary to uniform sampling of space, this approach uses uniform sampling of the system in time, with updated objects and their orbital elements at each time step. Because the spatial density at each instant is obtained from the actual state of the system, no assumptions regarding secular variations in objects' orbital elements are required. The total number of collisions between two objects over L time intervals $[t_k, t_{k+1})$ with $k = 0, \ldots, L$ is

$$\eta_{ij} = \int_{k=0}^{k=L} \int_{t_k}^{t_{k+1}} \dot{\eta}_{ij} \, \mathrm{d}t \, \mathrm{d}k = \int_{k=0}^{k=L} \int_{t_k}^{t_{k+1}} A_c n_{\boldsymbol{r}}^{(i)} n_{\boldsymbol{r}}^{(j)} v_{ij} \Delta V \, \mathrm{d}t \, \mathrm{d}k \tag{2.7}$$

where $\dot{\eta}_{ij}$ is precisely the collision rate obtained from Eq. (2.6) over a cube of volume ΔV, assuming that the objects' spatial densities are constant over the volume (i.e., the objects are

uniformly distributed in the cube), valid for sufficiently small volumes. The cube size needs to be small enough to capture the true collisional nature of the system while balancing computational cost. The standard cube size adopted in LEGEND is 10 km (J. C. Liou, 2006). Just as Kessler's gas-kinetic formula does, the method assumes that all objects within the control volume can collide. This assumption proved especially wrong in the case of endogenous collisions in a constellation. The method is unable to distinguish between two satellites on orbits that intersect and two satellites that approach to within a few kilometers on orbits that do not intersect: if the cube size is larger than the separation between the constellation planes, objects on different orbital planes or even different shells are allowed to collide, even when they are constrained to orbits that do not intersect by design, severely overestimating collisions (Lewis et al., 2019). Lewis et al. (2019) finds that to provide good estimates of the true collision rate between two satellites, cube sizes of between 4 km to 16 km are needed. However, as the cube size decreases, a higher trajectory resolution is required, in turn requiring a smaller time step (less than 1 day for cube sizes smaller than 10 km), considerably increasing the computational load.

An alternative flux-based method was introduced by Klinkrad (1993) and later embedded at the core of ESA's Meteoroid and Space Debris Terrestrial Environment Reference (MASTER) (Bendisch et al., 2004; Klinkrad et al., 2001). Klinkrad's method decomposes orbital space into a fixed spherical grid of sectors bounded by altitude shells, declination bands, and right ascension meridians.

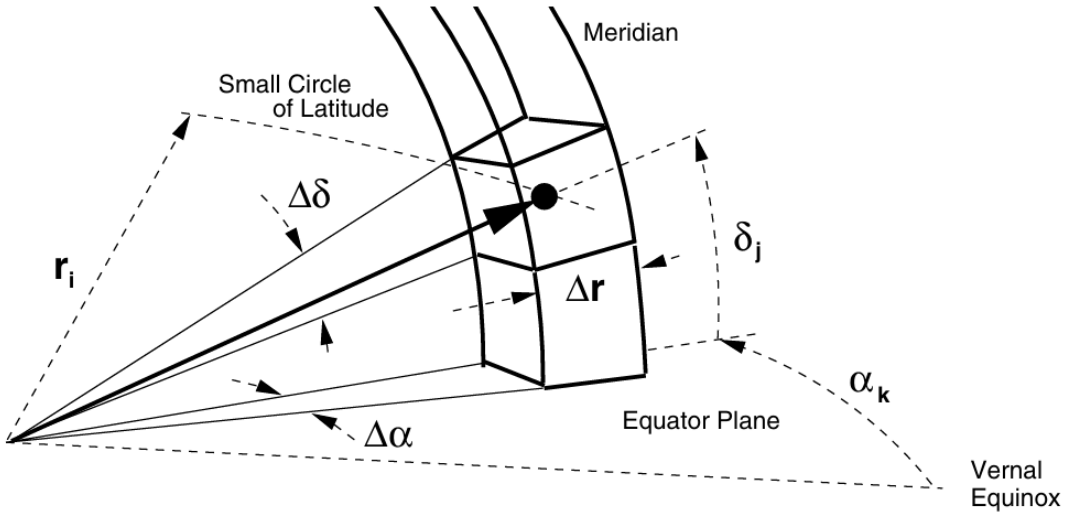

Figure 2.4: Definition of control volume cells in an inertial, geocentric reference frame (Klinkrad, 2006, p. 62)

Each object is propagated to a common reference epoch and its orbit is analysed for intersections with each sector; the resident probability within a given sector is computed as the fraction of the anomalistic period spent inside it, and the local spatial density is obtained as the sum of all contributions divided by the sector volume. The transient velocity of the object at the central crossing time is also stored in a local horizontal frame, enabling the full directional structure of the flux to be retained. To evaluate the collision flux on a specific target, its orbit is in turn crossed against the same grid, and at each intersected sector, the stored characteristics of the field population are retrieved to compute relative velocities and differential flux contributions, which are then summed along the target trajectory. This approach naturally discriminates flux by source type, altitude, and inclination, and can resolve directionality in a way that the isotropic cube formulation cannot. A further variant, the object-centered control volume method (e.g. DAMAGE's bubbles), extends this idea by attaching the control volume to each individual target rather than fixing it to a global grid, computing the flux of impactors sweeping through a sphere or box co-moving with the object; while conceptually equivalent, this formulation simplifies the bookkeeping for single-target risk assessments.

2.3.3. Phase-Space Density Formulation

In its most general form, the impact rate on a target at position $\boldsymbol{r}^*$ with velocity $\boldsymbol{v}^*$ is the kinetic integral

$$\dot{\eta}(\boldsymbol{r}^*, \boldsymbol{v}^*) = \int_{\mathbb{R}^3} A_c n_{\boldsymbol{r},\boldsymbol{v}}(\boldsymbol{r}^*, \boldsymbol{v}) \|\boldsymbol{v} - \boldsymbol{v}^*\| \, \mathrm{d}^3\boldsymbol{v} \tag{2.8}$$

where $n_{\boldsymbol{r},\boldsymbol{v}}$ is the phase-space density of the impactors in the cartesian state space (Frey & Colombo, 2021). Assuming a spherical target (A_c equal in every direction) and substituting $n_{\boldsymbol{r},\boldsymbol{v}} = f_{\boldsymbol{v}|\boldsymbol{r}} n_{\boldsymbol{r}}$ where $f_{\boldsymbol{v}|\boldsymbol{r}}$ is the conditional distribution of velocity given position, one obtains an expression identical to Eq. (2.4)

$$\dot{\eta}(\boldsymbol{r}^*, \boldsymbol{v}^*) = A_c n_{\boldsymbol{r}} \int_{\mathbb{R}^3} \|\boldsymbol{v} - \boldsymbol{v}^*\| f_{\boldsymbol{v}|\boldsymbol{r}} \, \mathrm{d}^3\boldsymbol{v} = A_c n_{\boldsymbol{r}} v_{\mathrm{rel}} \tag{2.9}$$

The central challenge is the evaluation of the phase-space density. Kessler's approach assumes full randomization of the Euler angles, collapsing the six-dimensional distribution in Cartesian elements onto a two-dimensional spatial distribution in altitude and latitude (see Eq. (2.5)). The phase-space density approach abandons this assumption and instead maintains the full distribution in Keplerian elements, obtaining the spatial density through a distribution transformation.

Giudici et al. (2023) shows that the kinetic impact rate integral in Eq. (2.8) admits an identical 3D representation in the (a, e, i) space. Given the target's position $\boldsymbol{r}^*$, for each triplet (a, e, i), the three remaining orbital elements (Ω, ω, M) are algebraically determined, yielding at most four intersecting orbit solutions. The velocity integral therefore reduces to a sum over these four branches, each weighted by the residence time of the field orbit at $\boldsymbol{r}^*$, in direct analogy with Kessler's (and Öpik's) original formulation. The method imposes no assumptions on eccentricity, inclination, or the randomisation of the angular elements, thereby extending the validity of density-based collision rate estimates to any orbital regime and impact geometry.

2.4. Network Theory

All the traditional evolutionary models previously introduced output results in the form of the number of objects over time and the cumulative number of catastrophic collisions. Semi-deterministic models provide the full trajectory of each object, but identifying risk drivers and attributing population-level metrics is not straightforward and requires considerable post-hoc analysis. Besides basic information, some metrics have been proposed to understand the health of the space environment. The Criticality of Spacecraft Index (CSI) (Rossi et al., 2015) measures the danger represented by typical classes of objects and their possible effects on the environment. The Environmental Consequences of Orbital Breakups (ECOB) (Letizia, Colombo, Lewis, & Krag, 2016) is a severity index that measures how dangerous a break-up event can be for other space objects in terms of the resulting collision probability for a set of target spacecraft with the breakup fragments over time. McInnes (2000) for the continuum formulation first, followed by D'Ambrosio et al. (2023) for MOCAT-SSEM , developed the concepts of carrying capacity from the equilibrium point of the dynamical system, as previously discussed. Though these metrics provide valuable insights into the characteristics of the space environment from different scopes, they remain inherently object-centric. This section therefore turns to the applications of network theory to the study of the space environment. By embedding the population into a graph whose edges capture the interactions between objects or orbital classes, it becomes possible to quantify systemic properties that no single-node metric can reveal.

2.4.1. Network Theory for Space Applications

The first application of graph theory to the space debris environment was introduced by Lewis et al. (2010), who uses outputs from DAMAGE to construct an undirected network in which vertices represent orbiting objects and edges encode conjunction events between them. By computing standard network metrics (degree centrality, betweenness, clustering coefficient, and

assortativity) on projected population snapshots, they show that a small number of high-degree objects act as hubs responsible for a disproportionate share of conjunction activity, and that removing these objects would efficiently disrupt the collisional network. Their analysis revealed disassortative mixing (high-degree nodes preferentially link to low-degree ones), a property shared with other technological networks such as the internet, with direct implications for the design of active debris removal campaigns. This work established the foundational vocabulary that all subsequent space-network studies have adopted.

Romano et al. (2024) introduces the Resident Space Object Network (RSONet), a fully data-driven model in which Resident Space Objects (RSOs) are connected by links weighted with conjunction probability, derived either from Two Line Element (TLE) propagation with a three-filter geometric screener or from CDMs. Rather than analysing network metrics in isolation, they combined degree, clustering, and betweenness into a single relevance score that ranks each RSO by its contribution to the population-wide collision risk, providing a practical tool for prioritising ADR targets. Lin et al. (2024) extends a similar conjunction-probability graph with Louvain community detection, identifying clusters of high-risk objects and showing that, after constellation deployment, average community risk increases by a factor of roughly 3.5 compared to the pre-deployment baseline, a result directly relevant to mega-constellation safety assessment. More operationally oriented, Owens-Fahrner et al. (2025) cast constellation design itself as a graph optimization problem: candidate orbital slots become nodes, collision probability and coverage metrics define edge weights, and simulated annealing selects the subgraph that minimizes risk while maximizing LEO observability, demonstrating how graph theory can inform not just post-hoc analysis but the design phase of large constellations.

All of these approaches share a common limitation: they are static snapshots of the population at a given moment. The network is reconstructed at each timestep from propagated orbits, but the population itself does not evolve within the network framework. Conjunction graphs capture the risk as it currently stands; they do not propagate launches, failures, fragmentations, and decay forward in time in a self-consistent way. Coupling topological analysis with stochastic population dynamics is precisely the gap addressed by NESSY, which is described in the following subsection.

2.4.2. A Network-based Evolutionary Model of the Space Environment

A qualitative shift occurred with Acciarini and Vasile (2020), who departs entirely from the static paradigm. In their work, they model the space environment as a multilayer temporal network, with a physical layer capturing collision propagation and an information layer modeling inter-satellite communication. The network dynamics are described by a susceptible-infectious-susceptible (SIS) epidemic model: disruptions spread probabilistically through the graph, and spectral properties of the supra-adjacency matrix (epidemic threshold, spectral gap) determine whether a cascade self-extinguishes or persists. This introduced the language of network dynamics into the space context and identified which nodes and links represent systemic vulnerabilities in the orbital environment as a whole.

In their subsequent work, Acciarini and Vasile (2021) made the decisive conceptual shift from object-level nodes to orbital-class population nodes: each node aggregates all objects of a given type (active, inactive) within a defined orbital class, and each directed edge encodes a time-varying collision rate between classes, in direct analogy with source-sink models. Collision events are sampled at each timestep from Poisson distributions parametrized by the gas-kinetic collision rate formula, and the NASA SBM generates the resulting fragment populations. Lyapunov analysis shows that the equilibrium solution of this network, under constant source and sink terms, is unconditionally stable, providing a theoretical grounding absent from classical evolutionary models.

NESSY refines and extends this into a fully operational framework. Nodes are populations binned by altitude shell and inclination, and operational status (payloads, upper stages, non-maneuverable satellites, and fragments), and edges encode a richer set of stochastic transition rates: satellite failures, post-mission disposal (with success and failure branches), atmospheric

decay, shell-crossing flows weighted by orbital residence time, collision rates, and fragmentation by collision. Population counts are advanced via finite-difference equations rather than integrated ODEs, with Poisson-drawn increments at each timestep ensuring that event counts remain integer-valued and statistically well-posed. The spatial density of each bin enters the collision rate integral, assuming uniform object distribution within each bin.

This architecture yields outputs that neither the static network models of Lewis nor RSONet, nor the traditional source-sink models can deliver on their own. From an evolutionary perspective, NESSY generates time series of collision counts, fragment accumulation, and carrying-capacity thresholds across the entire population. From the network side, the same simulation yields a time-varying, weighted, directed graph whose topological properties (degree distribution, node centrality, etc.) can be monitored alongside the population dynamics. A bin that acts as a network hub (high in-degree from neighboring shells) can be identified not only as a region of elevated collision rate but as a potential nucleation point for a Kessler-type runaway. In this regard, NESSY is the first framework to close the loop between structural network analysis and stochastic population evolution, making it the ideal backbone for the model developed within this project.

3

Astrodynamics

This chapter establishes the astrodynamics background on which the model of chapter 4 rests, fixing both the notation and the physical results used throughout. Section 3.1 defines the frame to which every orbit in this work is referred. Section 3.2 introduces the three equivalent parametrizations of a state that recur in the following chapters: Cartesian coordinates, in which the collision rate is formulated, the equatorial coordinates of right ascension and declination, in which the geometry of the intersection between two orbital planes is resolved, and the Keplerian elements, in which the population is binned and propagated. Section 3.3 presents the two-body problem and its analytical solution, followed by the perturbed form of the equations of motion and by the variational equations that recast the effect of a perturbation as a drift of the orbital elements. Finally, section 3.4 reviews the perturbations acting on a satellite in LEO, justifies the retention of Earth's oblateness and atmospheric drag alone, gives the secular nodal precession and orbital decay that the two produce, and closes on the distinction between osculating and mean elements that underlies the propagation machinery.

3.1. Reference Frames

A reference frame is defined by an origin, a fundamental plane, and a reference direction lying in that plane. It fixes the geometry against which a state is expressed, independently of the coordinates chosen to express it. All orbits in this work are referred to an Earth-centered inertial frame. Its origin is Earth's center of mass, its fundamental plane is the equatorial plane, and its reference direction is the vernal equinox, i.e., the intersection between the equatorial and the ecliptic planes in the direction of the ascending node of the Sun's apparent orbit. The third axis completes the right-handed triad along Earth's rotation axis, pointing North.

The frame is non-rotating, in the sense that it does not follow Earth's diurnal rotation. It is not strictly inertial, since Earth itself accelerates along its orbit around the Sun, but the residual acceleration is irrelevant at the scale of a geocentric orbit. The equatorial plane and the equinox direction are themselves not fixed: the lunisolar and planetary torques acting on Earth's equatorial bulge make them drift at a rate of roughly 50″ per year, so that a reference epoch must be attached to the frame for its reference direction to be unambiguous. This choice, however, has no bearing on the results of this work: the collision rate derived in chapter 4 depends on RAAN only through differences between objects, and the nodal precession rate is independent of RAAN altogether, so that a rigid rotation of the reference direction leaves both the dynamics and the collision rate of the population unchanged.

3.2. Coordinate Systems

The state of an object can be expressed via a variety of equivalent formulations, each one providing different advantages depending on the problem at hand. Traditionally, the position and velocity of an object are expressed in Cartesian coordinates. In astrodynamics, however, it is often common to describe the trajectory of an object by means of its orbital elements, which give a better insight in the shape and orientation of the orbit. Other representations can be defined to overcome numerical limitations associated with singularities of the classical elements set.

The collision rate in eq. (2.8) is defined in terms of Cartesian elements, but the density distribution of the objects in space is more intuitively defined in terms of orbital elements (Frey & Colombo, 2021). Likewise, the effect of perturbations on an object's orbit finds its natural description in terms of orbital elements, following Lagrange and Gauss perturbation theories, allowing to derive analytical and semi-analytical expressions for the propagation of the state over time.

3.2.1. Cartesian Coordinates

The state in Cartesian coordinates is the combination of the position $\boldsymbol{r} \in \mathbb{R}^3$ and velocity $\boldsymbol{v} \in \mathbb{R}^3$ vectors in the three-dimensional space

$$\begin{pmatrix} \boldsymbol{r} & \boldsymbol{v} \end{pmatrix}^T = \begin{pmatrix} x & y & z & v_x & v_y & v_z \end{pmatrix}^T \in \mathbb{R}^6 \tag{3.1}$$

The equations of motion are easily defined in this system, which provides the simplest and most straightforward representation of the object in space. On the other hand, understanding the shape and orientation of the object's orbit is less intuitive, and so is the relationship between position and velocity components at any point in space.

3.2.2. Equatorial Coordinate System

An object's position can also be expressed in spherical coordinates defined with respect to Earth's equatorial plane and a fixed reference direction, conventionally the vernal equinox

$$\begin{pmatrix} r & \delta & \alpha \end{pmatrix}^T \in \mathbb{R}^3 \tag{3.2}$$

where r is the radial distance from Earth's center, the declination $\delta \in [-\pi/2, \pi/2]$ is the angle between the position vector and the equatorial plane, positive Northward, and the right ascension $\alpha \in [0, 2\pi)$ is the angle measured eastward in the equatorial plane between the reference direction and the projection of the position vector onto that plane. The spherical coordinates are shown in fig. 3.1.

3.2.3. Keplerian Elements

A holistic view of the orbit is easily obtained in terms of the orbital, or Keplerian, elements

$$\begin{pmatrix} a & e & i & \Omega & \omega & M \end{pmatrix}^T \in \mathbb{R}^6 \tag{3.3}$$

where the semi-major axis a describes the size of the orbit, and the eccentricity e its shape. The orientation of the orbit in space is described by the combination of three angles: the orientation of the orbital plane with respect to a reference plane (usually the central body's orbital plane) is described by the inclination i, i.e., the angle between the two planes, and the longitude of the node, or RAAN, Ω, i.e., the angle between a reference direction and the ascending node (the point where the object moves North through the reference plane) as measured on the reference plane; the orientation of the orbit itself within the orbtial plane is instead described by the the argument of periapsis ω, i.e., the angle between the ascending node and the orbit's periapsis (the

point of closest distance to the central body) as measured on the orbital plane. All elements are shown in fig. 3.1.

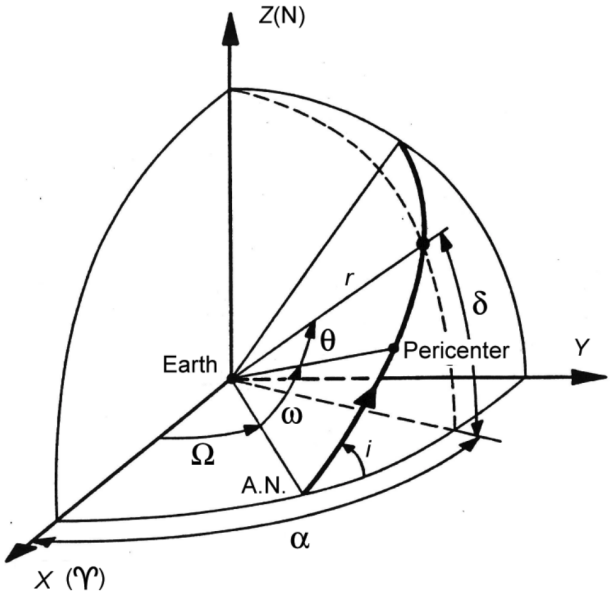

Figure 3.1: Definition of the orbital elements i, Ω, ω and θ and the position of a body in spherical coordinates (r, δ, α) (Wakker, 2015).

In the absence of perturbing accelerations, these five elements are fixed in time. The only time-varying element, which describes the angular position of the object on its orbit with respect to the periapsis, is the true anomaly θ. Because for non-circular orbits the true anomaly varies non-linearly with time, it is often substituted by the mean anomaly M, representing the time-proportional angle traveled by the object if it were subjected to uniform circular motion. The two are related by the eccentric anomaly E through Kepler's equation

$$M = E - e \sin E \tag{3.4}$$

and

$$\tan \frac{\theta}{2} = \sqrt{\frac{1+e}{1-e}} \tan \frac{E}{2} \tag{3.5}$$

The angle between the ascending node and the object's position vector, measured on the orbital plane, is called the argument of latitude $u = \omega + \theta$, and is the sum of the argument of periapsis and the true anomaly. It is particularly useful for circular orbits, where the argument of periapsis is undefined.

3.3. Equations of Motion

The motion of a body orbiting a larger central one is largely governed by the central body's point mass gravity. Assuming that the mass of the orbiting body is negligible, i.e., the gravitational force exerted by the smaller body onto the larger one is negligible, and that the mass of the central body is constant and concentrated in a single point, the following equation of motion can be directly derived from Newton's law of universal gravitation

$$\ddot{\boldsymbol{r}} = -\frac{\mu_\oplus}{r^3}\boldsymbol{r} \tag{3.6}$$

where $\boldsymbol{r}$ is the distance between the two bodies, and $\mu_\oplus$ is the gravitational parameter of the central body. For Eearth, its value is $3.986\,004\,418 \times 10^{14}$ m^3/s^2. In the absence of any external

perturbation, also known as a Keplerian orbit, eq. (3.6) can be integrated analytically

$$r = \frac{a(1-e^2)}{1+e\cos\theta} \tag{3.7}$$

which is the equation of a conic section in polar coordinates expressed in terms of the Keplerian elements. For $e < 1$ the orbit is closed: specifically, if $e = 0$ the orbit is circular, elliptical otherwise. When the only force acting on the system is the point mass gravity of the central body, the principle of conservation of mechanical energy is expressed in the form of the *vis-viva equation*

$$\frac{v^2}{2} - \frac{\mu_\oplus}{r} = -\frac{\mu_\oplus}{2a} \tag{3.8}$$

which allows to compute the instantaneous velocity of the orbiting body at a distance r from the central one only from the semi-major axis a of its orbit.

A more general form of the equations of motion can be obtained by also accounting for external perturbing accelerations $\boldsymbol{f}$

$$\ddot{\boldsymbol{r}} = -\frac{\mu_\oplus}{r^3}\boldsymbol{r} + \boldsymbol{f} \tag{3.9}$$

whose solution is generally obtained by means of numerical integration methods, or by deriving approximate analytical solutions under appropriate assumptions.

The equations above are written in Cartesian coordinates, but the effect of a perturbation is more naturally read in terms of orbital elements. The idea, known as the variation of parameters, is to retain the Keplerian solution as the description of the motion and let the six constants that define it become functions of time: at every instant the object is taken to move along the Keplerian orbit tangent to its true trajectory, and the perturbation manifests itself as a slow drift of that orbit. The three second-order equations in $\boldsymbol{r}$ are thereby replaced by six first-order equations, one per element, whose right-hand sides vanish when $\boldsymbol{f}$ does, recovering the constant elements of a Keplerian orbit. The reformulation is exact, and its advantage is that the fast Keplerian motion is already contained in the solution, leaving only the slow departure from it to be integrated.

Two equivalent formulations of these variational equations exist. When $\boldsymbol{f}$ is conservative, it derives from a disturbing potential, and the rate of change of each element can be written in terms of the partial derivatives of that potential with respect to the elements themselves. These are the *Lagrange planetary equations*. However, a non-conservative perturbation admits no potential, and the Lagrange form does not apply to it. The *Gauss planetary equations* lift this restriction by taking as their input the perturbing acceleration itself, resolved along the radial, transverse and out-of-plane directions, which makes them applicable to drag. Neither set is directly integrable, since the rates depend on the position of the object along its orbit and therefore vary over one revolution. Approximate analytical solutions are obtained by averaging them over one orbital period, which discards the periodic terms and retains the secular drift, and it is in this form that both sets are used in the remainder of this work. The complete systems of differential equations, together with their derivation, can be found in Vallado and McClain (2013).

3.4. Perturbations

The perturbing acceleration $\boldsymbol{f}$ collects every effect that the two-body problem leaves out. For a satellite in LEO the relevant contributions are the non-spherical part of Earth's gravitational field, atmospheric drag, the gravitational attraction of the Moon and the Sun, and Solar Radiation Pressure (SRP), followed by a tail of smaller effects such as solid and ocean tides, Earth's infrared and albedo radiation, and relativistic corrections. Their magnitudes span many orders of magnitude and, crucially, do not scale in the same way with altitude, so that which of them matters depends on where in LEO the object is.

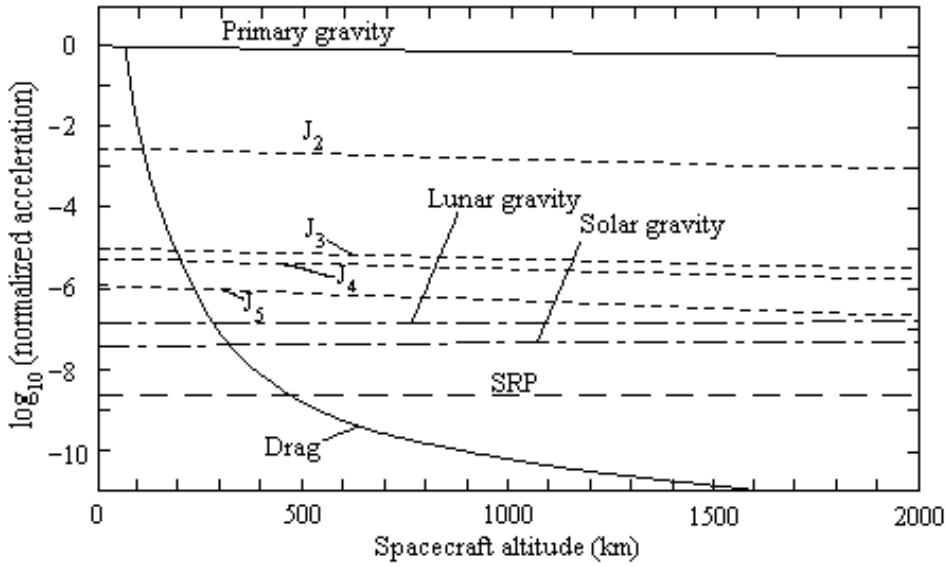

Figure 3.2: Perturbing accelerations acting on a satellite in LEO as a function of altitude (Virgili-Llop, 2014).

Figure 3.2 shows the hierarchy. The J_2 term stands out from all the others, at roughly three orders of magnitude below the primary gravity and two to three above the next contribution, and its dominance is nearly constant across the whole regime. Atmospheric drag is the only term whose magnitude varies strongly with altitude, falling off exponentially: below 200 km it rivals J_2 itself, it crosses the third-body and higher zonal terms between roughly 400 km to 500 km, and it becomes the smallest of the perturbations only above 800 km or so. Everything else, i.e., the higher zonal harmonics, the lunar and solar attraction, and SRP, sits five orders of magnitude or more below the primary gravity and is essentially flat with altitude.

Only J_2 and drag are modeled in this work, and magnitude alone is not the whole argument. A perturbation that only makes the elements oscillate about their mean values leaves the orbit where it was, while one that produces a secular drift displaces it steadily, and it is the latter that governs the evolution of an orbit on timescales of months and years. In this regard, the two retained terms are the ones that count. Earth's oblateness makes the line of nodes precess and the line of apsides rotate at rates of the order of several degrees per day, continuously reorienting the orbital plane in inertial space while leaving the size and shape of the orbit untouched. Drag, being dissipative, removes energy from the orbit and so shrinks the semi-major axis monotonically, circularizing the orbit and eventually driving the object to re-entry. Between them, the two perturbations therefore account for the whole secular evolution of the orbit: one moves the plane, the other moves the altitude. The lunar and solar attraction and SRP do contribute secular terms of their own on Ω and ω, but in LEO the oblateness torque dominates so strongly that these are four to five orders of magnitude slower than the J_2 nodal drift, and what remains of their effect is periodic and averages out over a revolution. The same two perturbations are retained by MOCAT-MC (Jang et al., 2024) and by the source-sink literature more broadly, and they are the two that admit a closed-form averaged solution (Martinusi et al., 2015).

3.4.1. Earth Oblateness

Equation (3.6) assumes a spherically symmetric mass distribution of the central body. However, in reality, Earth's shape, the geoid, is slightly pear-like, with a bulge at the equator and flattening at the poles. To account for this, the gravitational potential U is formulated as the superimposition of several terms that depend on geocentric latitude and longitude, which alternately add and subtract from the two-body potential (Wertz & Larson, 1999). Then, the acceleration due to

Earth's full gravitational field on an orbiting object is

$$\ddot{\boldsymbol{r}} = \nabla U(\boldsymbol{r}) \tag{3.10}$$

Zonal harmonics are the terms of the gravitational potential function that depend on geocentric latitude ϕ only, and are hence symmetrical about the polar axis, dividing Earth into spherical segments

$$U = \frac{\mu_\oplus}{r}\left[1 - \sum_{\ell=2}^{\infty} J_\ell \left(\frac{R_{eq}}{r}\right)^\ell P_\ell(\sin\phi)\right] \tag{3.11}$$

where $R_{eq} = 6378\,\mathrm{km}$ is Earth's equatorial radius, P_ℓ is the Legendre polynomial of degree ℓ, and J_ℓ are dimensionless geopotential coefficients representing Earth's oblateness. The second-degree term, $J_2 = 0.001\,082\,626\,9$, accounts for most of Earth's gravitational departure from a perfect sphere (Vallado & McClain, 2013). For a point-mass gravity, the summation term in eq. (3.11) disappears and eq. (3.6) is recovered.

Under J_2 only, the geopotential is

$$U = \frac{\mu_\oplus}{r}\left[1 - \frac{J_2}{2}\left(\frac{R_{eq}}{r}\right)^2 (3\sin^2\phi - 1)\right] \tag{3.12}$$

Combining the above expression with the Lagrange planetary equations and averaging the long and short-term periodic terms to isolate the secular variations produces

$$\dot{\Omega} = -\frac{3}{2}\frac{nR_{eq}^2 J_2}{(a(1-e^2))^2}\cos i \tag{3.13}$$

which is the nodal precession rate[1]. Effectively, under the effect of J_2, RAAN drifts westward (for a prograde orbit) or eastward (for a retrograde orbit) at a rate which depends on the inclination, altitude, and shape of the orbit.

3.4.2. Atmospheric Drag

Next to Earth's oblateness, atmospheric drag is the other perturbation that most strongly affects the motion of satellites near Earth. Drag, the force exerted by atmospheric particles on the surface of the satellite, slows down its motion, causing the satellite to decay and eventually re-enter the atmosphere. The denser the atmosphere is, the stronger the effect: at high altitudes, the atmosphere is very rarefied, and the effect of drag is negligible; closer to Earth, the atmosphere is denser, and in the last phases of a satellite's life, drag is the strongest perturbation and altitude decreases rapidly.

Atmospheric acceleration is a combination of the effects of the shape, orientation, and aerodynamic properties of the satellite, the atmospheric density at the satellite's altitude ρ, and the relative velocity of the satellite with respect to the atmosphere v_{air}, and can be expressed according to the following equation

$$\boldsymbol{f}_{drag} = -\left(\frac{1}{2}\frac{C_D A}{m}\rho\, v_{air}^2\right)\hat{\boldsymbol{v}}_{air} \tag{3.14}$$

The ballistic coefficient $B = C_D A/m$ quantifies the ability of the satellite to overcome drag: objects with a large mass m and small cross-sectional area A keep their inertia much better, while lighter and larger satellites decelerate faster. The drag coefficient C_D is a dimensionless number that absorbs all other aerodynamic properties of the satellite, such as its shape and material, and is often approximated to 2.2. The velocity term in eq. (3.14) is taken relative to the atmosphere, which co-rotates with Earth and carries winds of its own. Assuming neither, v_{air} simplifies directly to the satellite's inertial orbital speed.

[1] For a full derivation refer to Vallado and McClain (2013).

In general, the atmospheric density ρ depends on altitude, geographical position, diurnal variations, seasonal variations, solar activity and the solar cycle, the atmosphere's rotation, winds, magnetic storms, and tides, making it by far the most difficult parameter to model accurately. Several atmospheric models exist at varying degrees of sophistication. The simplest is the exponential one, which models the change of density with altitude h according to

$$\rho = \rho_0 \exp\left(-\frac{h - h_0}{H}\right) \tag{3.15}$$

where ρ_0 is the density at a reference altitude h_0, and H is the fractional change in density with height, also called *scale height*. This model provides a very coarse approximation of the atmosphere, but its accuracy can be significantly improved by separating the atmosphere into layers with corresponding tabulated h_0, ρ_0 and H values, and applying the formula to each layer separately. This *layered* version of the exponential model is purely analytical, depends only on altitude, and is fast to compute, still producing moderately accurate results, and is the one adopted in this work. More sophisticated models, such as the Jacchia-Bowman 2008 (JB2008) model adopted by both Jang et al. (2024) in MOCAT-MC and Wang et al. (2025) in NESSY, use statistical fits to atmospheric data collected over decades, allowing the modeling of the time-evolution of the atmosphere under the effects of solar weather and geomagnetic storms, among others.

Because energy is lost due to friction, drag is a non-conservative perturbation and hence cannot be treated via the Lagrange planetary equations. Instead, applying the Gauss planetary equations, and assuming a non-rotating atmosphere with no winds and a circular orbit, the rate of decay of the orbit is expressed as[2]

$$\dot{a} = -\rho B a v = -\rho B \sqrt{a \mu_\oplus} \tag{3.16}$$

3.4.3. Osculating and Mean Elements

Under the perturbations of section 3.4, the Keplerian elements are no longer constant. At any instant the state of the object still identifies a unique Keplerian orbit, the one tangent to the true trajectory at that point, and its elements are called *osculating*. They describe the real motion exactly, but they carry the full spectrum of its variation: short-period oscillations, with the period of one revolution, long-period ones, with the much longer period of the rotation of the line of apsides, and the secular drift.

When the interest lies in the evolution of an orbit over many revolutions, the periodic content is of little use, and averaging it out leaves a smoother and slower description of the dynamics. This is the idea behind the first-order satellite theory of Brouwer (1959): a canonical transformation removes the short and long-period terms and leaves a purely secular set of equations of motion. The elements that satisfy them are the *mean* elements. The short-period terms so removed have an amplitude of order J_2 relative to the element itself, which amounts to a few kilometers in semi-major axis for a LEO orbit. The original formulation degenerates for near-circular and near-equatorial orbits, a singularity later removed by Lyddane (1963).

The mean elements are not an alternative parametrization of the same trajectory, but the variables of a different, averaged, dynamical problem, and the map connecting them to the osculating set is a canonical *contact transformation* (Martinusi et al., 2015). For instance, an osculating and a mean semi-major axis of equal numerical value describe different orbits. All orbital elements in this work are treated as mean elements, averaged according to Brouwer's theory. The analytical propagator adopted in chapter 4 advances the averaged equations directly, the initial populations of chapter 8 are interpreted in the same sense, and no contact transformation is applied in either direction.

[2] For a full derivation refer to Vallado and McClain (2013).

4

Model Formulation

This work introduces the Plane-Resolved Integrated Stochastic Model (PRISM), which is a network model adapted from NESSY (Wang et al., 2025). Each node is identified by an orbit bin and a specie. In general, an orbit bin can be any set of orbital parameters but, in this work, it is defined by a combination of altitude h, inclination i and RAAN Ω, as these are the parameters that most straightforwardly map to the planes of a Walker constellation. In symbols, if the (h, i, Ω) phase-space is divided in N_h, N_i and N_Ω bins respectively, for any $\ell = 1, \dots, N_h$, $m = 1, \dots, N_i$ and $n = 1, \dots, N_\Omega$, an orbit bin is

$$I_h^{(\ell)} \times I_i^{(m)} \times I_\Omega^{(n)} = [h_\ell, h_{\ell+1}) \times [i_m, i_{m+1}) \times [\Omega_n, \Omega_{n+1})$$

where $I_h^{(\ell)}$ is the ℓ-th altitude bin, $I_i^{(m)}$ is the m-th inclination bin, and $I_\Omega^{(n)}$ is the n-th RAAN bin. Geometrically, the two angular intervals give a band on the sphere, the region swept by every orbit whose inclination and RAAN fall inside them, and the altitude interval gives that band a radial thickness. Figure 4.2 shows several such bands, each one the angular extent of a single bin. To simplify notation, the triplet index (ℓ, m, n) is flattened into a single index j via

$$i = \ell N_i N_\Omega + m N_\Omega + n$$

with $i = 1, \dots, J = N_h N_i N_\Omega$. Five species of objects are considered: *non-maneuverable* satellites (N), *fragments* (F), and *active payloads*, subdivided into *station-keeping* (A), *rising* (R), and *de-orbiting* (D) satellites. Therefore, the total number of nodes in the network is $5J$, and the notation S_i will be used to indicate a node of species $S \in \mathcal{S} = \{A, R, D, N, F\}$ in the i-th orbit bin. The number of objects $x_{S_i}(t_k)$ within node S_i at time t_k can vary because of the combination of different events such as collisions, explosions, natural dynamics, orbit-raising and de-orbiting maneuvers, PMD strategies, and new launches: some of these processes are continuous, like the orbit variations induced by drag and Earth oblateness, while others are better represented with discrete jumps in the number of objects, like collisions and explosions. In the network model, links between a node S_i and other nodes (or itself) represent the rate at which objects transition between the nodes due to the events listed above.

The reminder of this chapter focuses on the mathematical characterization of such events and the definition of the stochastic equations governing the network dynamics. Specifically, section 4.1 presents a general overview of the assumptions made and the interactions between nodes. Section 4.2 illustrates the collision rate formulation and its numerical implementation. Finally, sections 4.3 to 4.5 explicate the details about the proposed orbital dynamics, launch traffic, and thrust models, respectively.

4.1. The Network's Evolutionary Model

Let $\boldsymbol{X}_k = (x_{A_1}, x_{R_1}, x_{D_1}, x_{N_1}, x_{F_1}, \ldots, x_{A_J}, x_{R_J}, x_{D_J}, x_{N_J}, x_{F_J})^T$ be the vector that contains all the nodes S_i at step k. Then, the set of equations governing the evolution of the network can be written in the compact form

$$\boldsymbol{X}_{k+1} = \boldsymbol{X}_k + \boldsymbol{Y}(\boldsymbol{X}_k) + \boldsymbol{g}(t_k) \tag{4.1}$$

where $\boldsymbol{Y}(\boldsymbol{X}_k)$ contains the interactions among nodes and satellite failures, while $\boldsymbol{g}(t_k)$ contains all the time dependent terms. Eq. (4.1) can be understood as a stochastic version of the source-sink model from eq. (2.1). The key difference with respect to previous work (D'Ambrosio et al., 2023), is that just as for NESSY (Wang et al., 2025), the evolutionary equations are derived directly from the structure of the network, allowing to flexibly balance computational complexity and accuracy: nodes can represent groups of objects of different physical properties, belonging to sites of different volumes, and the links between pairs of objects are orders of magnitude less than the pairs of objects. More importantly, it is possible to directly follow the evolution of the relationships between nodes and extract information that identifies the interdependency of objects and orbit bins (Wang et al., 2025). Distinctly from NESSY, which only considers bins of altitude and inclination, the proposed model also groups objects by their RAAN, such that each node can capture a constellation's orbital plane, and the interaction between planes can be studied, without assuming the uniform distribution of RAAN. Furthermore, evolutionary models (such as NESSY, MOCAT and DELTA) are generally applied to century-scale scenarios with year-long time steps, and can hence safely ignore short-term processes like orbit raising and de-orbiting, which for a LEO constellation usually last 1 to 2 months. Here, because the model is intended to be applied to a mission lifetime of a few decades, with time steps as short as 1 day, these processes are fully accounted for, hence the presence of the R and D species, allowing to capture the risk introduced by maneuvering objects which cross the constellation's shells during their spiraling motion. On the other hand, rocket bodies and the background population are not included, isolating the effect of endogenous interactions within the constellation.

4.1.1. Constellation Dynamics as a Poisson Point Process

As mentioned at the beginning of the chapter, the evolution of the system can be seen as the combination of slow continuous processes, like natural dynamics and electric orbit raising and de-orbiting, and sudden jumps in the population, like launches, collisions, and random failures. While launches are deterministic, as a launch schedule is determined a priori, collisions and failures are random events that may or may not happen according to some probability. By discretizing the simulation time domain into small steps and freezing the system at the beginning of every step, the event rate can be treated as a constant, independent parameter for the duration of that interval. Since the constellation consists of a large, discrete number of satellites, and the probability of a collision or a catastrophic failure to occur during any tiny fraction of time is extremely low, by the law of rare events, the probability of such events converges to a Poisson distribution. Therefore, the evolution of the system can be seen as a series of successive, step-wise homogeneous Poisson point processes (see Appendix A for a brief review of the probability theory and governing equations of the Poisson point process).

Specifically, the general Poisson jump process $\Delta_{S_iS_j}$ defining the number of collisions between nodes S_i and S_j is

$$\Delta_{S_iS_j} = \int_{t_k}^{t_{k+1}} \mathrm{d}\mathcal{N}\left(\alpha_{\text{CAM}}\, \dot{\eta}_{S_iS_j}\, \Delta t_k\right) \tag{4.2}$$

with

$$\alpha_{\text{CAM}} = \begin{cases} (1 - s_{\text{CAM}})^2 & \text{if both } S_i \text{ and } S_j \text{ are} \in \mathcal{A} \\ (1 - s_{\text{CAM}}) & \text{if either } S_i \text{ or } S_j \text{ are} \in \mathcal{A} \\ 1 & \text{otherwise} \end{cases} \tag{4.3}$$

where s_{CAM} is the CAM success rate, or the fraction of payloads that perform a successful collision avoidance maneuver, $\dot{\eta}_{S_iS_j}$ is the mean collision rate between the two nodes, and $\mathcal{A} = \{A, R, D\}$

is the subset of the active species, comprising station-keeping, rising and de-orbiting satellites, which are the only ones that can perform collision avoidance. Notably, α_{CAM} is smaller for active vs active encounters, for which collision avoidance can be planned and coordinated, leading to a higher success rate compared to active vs non-active conjunctions.

The jump process due to failed PMD maneuvers is

$$\Delta^{\gamma}_{A_i} = \int_{t_k}^{t_{k+1}} \mathrm{d}\mathcal{N}\left(\gamma \Delta^{EOL}_{A_i}\right) \tag{4.4}$$

where γ is the failure rate in performing post-mission disposal. Random failures during operations, such as during orbit transfer maneuvers or station keeping, can be modeled similarly, but in the current work, only PMD failure is modeled for simplicity.

Not every collision destroys both objects. The energy-to-mass ratio EMR of an encounter, defined in section 4.5, is a deterministic function of the masses of the objects involved in the collision and of their relative velocity, so the collision count splits into a catastrophic and a non-catastrophic part with no further sampling,

$$\Delta^{cat}_{S_iS_j} = \sum_{s=1}^{\Delta_{S_iS_j}} \mathbb{1}\{\text{EMR}_s \geq 40\,\text{J/g}\} \qquad\qquad \Delta^{nc}_{S_iS_j} = \Delta_{S_iS_j} - \Delta^{cat}_{S_iS_j}$$

and $\mathbb{1}\{\cdot\}$ is the indicator function, equal to 1 when the condition inside the braces holds and 0 when it does not. A catastrophic collision destroys both objects. A non-catastrophic one destroys only the lighter of the two, while the heavier survives the impact damaged: if it was an active satellite it loses control and becomes non-maneuverable, and if it was already non-maneuverable or a fragment it is unaffected. Adopting this rule, the number of objects a node loses to collisions over the step is denoted by $\Delta^{col}_{S_i}$ and the number that survive damaged, which is non-zero only for the active species since they are the only ones with a status left to lose, by $\Delta^{dam}_{S_i}$. The latter is the mechanism that demotes an active satellite to a faulty one on impact, and it replaces the small-fragment channel of NESSY, in which the demotion is caused by untracked debris below 10 cm entering at a fixed multiple of the tracked collision rate (Wang et al., 2025). Here fragments below that size are not modeled at all, so that channel is absent and the demotion is carried entirely by non-catastrophic collisions between tracked objects, as discussed in chapter 7.

4.1.2. Definition of the Stochastic Dynamic Equations

The five equations below are the bookkeeping of the processes above, written one specie at a time: each event enters the balance of one node as a loss and of another as a gain. Their terms belong to four families, detailed throughout the chapter. The collision terms are the ones defined in the previous section and the fragment source is eq. (4.45). The transitions of the maneuvering species, Δ^{EOL} at end of life and Δ^{dep} on deposition into the target orbit, are eqs. (4.39) and (4.40). The transport terms $\varepsilon^{\pm}$ of natural dynamics follow from the occupancy law of eq. (4.34), and their superscript is the sign the term carries in the balance and not a direction in altitude: ε^{+} is what the node gains from its neighbours and ε^{-} what it loses to them, whether those neighbours lie above or below (or beside) it. The launch terms Δ^{inj} are exogenous, read from a schedule fixed a priori. Of all of these only the collision terms and the PMD failures $\Delta^{\gamma}_{S_i}$ are random.

Active Payloads Station-keeping satellites are lost to collisions and to retirement, and gained from rising satellites reaching their target orbit and from launches injected directly into the operational orbit,

$$x_{A_i}(t_{k+1}) = x_{A_i}(t_k) - \Delta^{col}_{A_i} - \Delta^{dam}_{A_i} - \Delta^{EOL}_{A_i} + \Delta^{dep}_{R_i} + \Delta^{inj}_{A_i}(t_k) \tag{4.5}$$

The two collision terms together account for every collision the node takes part in: the first removes the satellites that are destroyed outright, the second those that survive but are no longer maneuverable, and which leave the node just the same. A station-keeping satellite holds its orbit, so no transport term appears.

De-orbiting satellites are gained from the payloads that retire and dispose successfully, and cross orbit bins as they spiral down,

$$x_{D_i}(t_{k+1}) = x_{D_i}(t_k) - \Delta_{D_i}^{col} - \Delta_{D_i}^{dam} - \varepsilon_{D_i}^{-} + \varepsilon_{D_i}^{+} + \left(\Delta_{A_i}^{EOL} - \Delta_{A_i}^{\gamma}\right) \tag{4.6}$$

Rising satellites are gained from launches and lost when they deposit into their target orbit,

$$x_{R_i}(t_{k+1}) = x_{R_i}(t_k) - \Delta_{R_i}^{col} - \Delta_{R_i}^{dam} - \varepsilon_{R_i}^{-} + \varepsilon_{R_i}^{+} - \Delta_{R_i}^{dep} + \Delta_{R_i}^{inj}(t_k) \tag{4.7}$$

The same $\Delta_{R_i}^{dep}$ enters eq. (4.5) as a gain and eq. (4.7) as a loss, and the same $\Delta_{A_i}^{EOL}$ leaves eq. (4.5) and enters eq. (4.6) less the fraction that fails, so neither transition creates or destroys a satellite.

Faulty Objects Non-maneuverable satellites are gained from every active satellite that loses control, whether by failing its PMD maneuver or by surviving a non-catastrophic collision, and once uncontrolled they decay and drift freely,

$$x_{N_i}(t_{k+1}) = x_{N_i}(t_k) - \Delta_{N_i}^{col} + \sum_{S_i \in \mathcal{A}} \Delta_{S_i}^{dam} + \Delta_{A_i}^{\gamma} + \varepsilon_{N_i}^{+} - \varepsilon_{N_i}^{-} \tag{4.8}$$

The damage sum runs over the three active species at the same orbit bin, since a demoted satellite becomes faulty where it was hit. No corresponding $\Delta_{N_i}^{dam}$ appears, because a faulty satellite that survives a non-catastrophic impact simply stays in its own node.

Fragments Fragments are lost to collisions and gained from every collision in the network, wherever it takes place,

$$x_{F_i}(t_{k+1}) = x_{F_i}(t_k) - \Delta_{F_i}^{col} + \varepsilon_{F_i}^{+} - \varepsilon_{F_i}^{-} + \sum_{j=1}^{J} \sum_{\ell=1}^{J} \sum_{(S_j, S_\ell)} \sum_{s=1}^{\Delta_{S_j S_\ell}} \zeta_{S_j S_\ell, s}^{F_i} \tag{4.9}$$

where (S_j, S_ℓ) is the unordered pair taken from the set of all 15 possible multi-combinations (combinations with repetitions allowed) of the network's species[1], and $\zeta_{S_j S_\ell, s}^{F_i}$ is the share of the cloud of the s-th such collision that lands in node F_i. The generating pair need not involve bin i at all: a fragment is placed by where its own orbit takes it, so a breakup anywhere in the lattice can feed any fragment node its cloud reaches. Fragments are the one specie with no launch term and, being untracked and unmaneuverable, the only one whose population is bounded from below only by re-entry.

4.2. Collision Rate Formulation

The formulation that follows starts from the kinetic impact-rate integral of eq. (2.8), in the phase-space form of Frey and Colombo (2021), and generalizes it from a single target to a second population. That generalization is not taken from the literature: the kinetic integral is written for one object at a given position and velocity, whereas what is needed here is the rate between two node populations, so the target's state is replaced by a distribution and integrated over as well. Setting the second population to a single object at $(\boldsymbol{r}^*, \boldsymbol{v}^*)$ recovers eq. (2.8) exactly, which is the check that the extension is consistent with its origin.

In its most general form, the *collision rate density* per unit volume between two independent and uncorrelated populations of objects, whose densities in the Cartesian phase-space are $n_{\boldsymbol{r},\boldsymbol{v}}^{(1)}$ and $n_{\boldsymbol{r},\boldsymbol{v}}^{(2)}$, respectively, is

$$\iint_{\mathbb{R}^6} n_{\boldsymbol{r},\boldsymbol{v}}^{(1)}(\boldsymbol{r}, \boldsymbol{v}_1) n_{\boldsymbol{r},\boldsymbol{v}}^{(2)}(\boldsymbol{r}, \boldsymbol{v}_2) A_c \|\boldsymbol{v}_1 - \boldsymbol{v}_2\| \, \mathrm{d}^3\boldsymbol{v}_1 \, \mathrm{d}^3\boldsymbol{v}_2$$

[1] (A, A), (A, R), (A, D), (A, N), (A, F), (R, R), (R, D), (R, N), (R, F), (D, D), (D, N), (D, F), (N, N), (N, F), and (F, F).

which measures the number of collisions per unit time per unit volume at a given position $\boldsymbol{r}$. Here A_c is the impact cross-sectional area of a pair. Since a node holds many objects of different sizes, the value used throughout is the mean over the pairs of the two nodes, derived in eq. (4.31). Each phase-space density is now separated into its position and velocity contributions, $n_{\boldsymbol{r},\boldsymbol{v}} = n_{\boldsymbol{r}}\, f_{\boldsymbol{v}|\boldsymbol{r}}$, where $n_{\boldsymbol{r}}(\boldsymbol{r})$ is the spatial number density of the population and $f_{\boldsymbol{v}|\boldsymbol{r}}(\boldsymbol{v}|\boldsymbol{r})$ is the distribution of velocity conditional on position, normalized to unit integral over $\boldsymbol{v}$, as introduced in eq. (2.8). With that substitution the expression above becomes

$$n_{\boldsymbol{r}}^{(1)} n_{\boldsymbol{r}}^{(2)} A_c \underbrace{\iint_{\mathbb{R}^6} f_{\boldsymbol{v}|\boldsymbol{r}}^{(1)}(\boldsymbol{v}_1|\boldsymbol{r}) f_{\boldsymbol{v}|\boldsymbol{r}}^{(2)}(\boldsymbol{v}_2|\boldsymbol{r}) \|\boldsymbol{v}_1 - \boldsymbol{v}_2\| \,\mathrm{d}^3\boldsymbol{v}_1 \,\mathrm{d}^3\boldsymbol{v}_2}_{v_{12}(\boldsymbol{r})}$$

where $v_{12}(\boldsymbol{r})$ is the relative speed between the two populations, marginalized over the velocity phase-space. Assuming that the impact cross-sectional area A_c is constant in every direction, the total collision rate is

$$\dot{\eta}_{12} = A_c \int_V n_{\boldsymbol{r}}^{(1)} n_{\boldsymbol{r}}^{(2)} v_{12} \,\mathrm{d}^3\boldsymbol{r} \tag{4.10}$$

where $V \subset \mathbb{R}^3$ is the region of the physical space reachable by both populations. The above expression is the same result obtained by Kessler (1981), already shown in eq. (2.6).

Let $\mathcal{U}_{ijk} \subseteq \mathbb{R}^4$ be a hyper-volume in the (h, i, Ω, u) phase-space corresponding to the i-th altitude shell, j-th inclination bin, and k-th RAAN bin, or

$$\mathcal{U}_{ijk} \equiv I_h^{(i)} \times I_i^{(j)} \times I_\Omega^{(k)} \times [0, 2\pi) \subseteq \mathbb{R} \times [0, \pi] \times [0, 2\pi) \times [0, 2\pi)$$

Any point $\boldsymbol{u} \in \mathcal{U}_{ijk}$ can be mapped to a position vector in the Cartesian space according to the transformation

$$\boldsymbol{r}(\boldsymbol{u}) = \boldsymbol{r}(h, i, \Omega, u) = (R_\oplus + h) \begin{pmatrix} \cos u \cos\Omega - \sin u \cos i \sin\Omega \\ \cos u \sin\Omega + \sin u \cos i \cos\Omega \\ \sin u \sin i \end{pmatrix} \tag{4.11}$$

where $R_\oplus$ is the Earth's radius, $h \in I_h^{(i)}$, $i \in I_i^{(j)}$, $\Omega \in I_\Omega^{(k)}$ and $u \in [0, 2\pi)$. Eq. (4.11) traces a great circle on the orbital plane of inclination i and RAAN Ω. Therefore, $\mathcal{U}_{ijk}$ is the union of all great circles obtained by varying altitude, inclination and RAAN in the respective ranges.

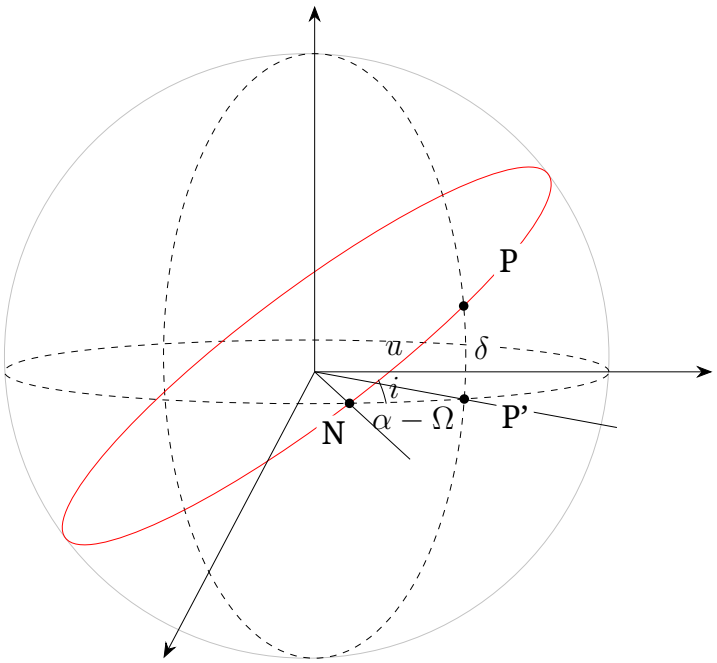

Figure 4.1: Great circle defined by $\boldsymbol{r}(\boldsymbol{u})$.

The transformation $\boldsymbol{r}(\boldsymbol{u}) : \mathcal{U} \subseteq \mathbb{R}^4 \to V \subseteq \mathbb{R}^3$ can be seen as the composition of a transformation from (h, i, Ω, u) to inertial spherical coordinates (r, δ, α), and then one from spherical coordinates to Cartesian coordinates (x, y, z): given the right spherical triangle with vertices at the ascending node N, position P, and projection P' on the equator, edges u, $\alpha - \Omega$, and δ as in fig. 4.1, angle i at N and right angle at P', Napier's rules give

$$\sin\delta = \sin i \sin u \tag{4.12}$$

for the declination, and

$$\cos(\alpha - \Omega) = \frac{\cos u}{\cos\delta} \qquad \tan(\alpha - \Omega) = \frac{\cos i \sin u}{\cos u} = \cos i \tan u \tag{4.13}$$

for the right ascension. The radial distance is

$$r = R_\oplus + h \tag{4.14}$$

The spherical coordinates are then mapped to Cartesian coordinates as usual

$$\begin{cases} x = r\cos\delta\cos\alpha \\ y = r\cos\delta\sin\alpha \\ z = r\sin\delta \end{cases}$$

Crucially, $\boldsymbol{r}(\boldsymbol{u})$ is non-invertible, i.e., multiple combinations of altitude, inclination, RAAN and argument of latitude can yield the same Cartesian position (x, y, z).

Let $n_{\boldsymbol{u}}$ be the objects' distribution in $\mathcal{U}_{ijk}$. Then, the corresponding spatial distribution $n_{\boldsymbol{r}}$ can be obtained via the Dirac generalized function (Frey & Colombo, 2021)

$$n_{\boldsymbol{r}} = \int_{\mathcal{U}_{ijk}} n_{\boldsymbol{u}} \delta[\boldsymbol{r}(\boldsymbol{u}) - \boldsymbol{r}] \,\mathrm{d}^4\boldsymbol{u} \tag{4.15}$$

where the solution of the above integral requires finding all roots of the system $\boldsymbol{r}(\boldsymbol{u}) - \boldsymbol{r} = \boldsymbol{0}$, as reviewed in appendix B.2. Substituting eq. (4.15) in eq. (4.10) yields the collision rate between two nodes

$$\begin{aligned} \dot{\eta}_{12} &= A_c \int_V \left(\int_{\mathcal{U}_{ijk}} n^{(1)}_{\boldsymbol{u}} \delta[\boldsymbol{r}(\boldsymbol{u}_1) - \boldsymbol{r}] \,\mathrm{d}^4\boldsymbol{u}_1 \right) \left(\int_{\mathcal{U}_{lmn}} n^{(2)}_{\boldsymbol{u}} \delta[\boldsymbol{r}(\boldsymbol{u}_2) - \boldsymbol{r}] \,\mathrm{d}^4\boldsymbol{u}_j \right) v_{12}(\boldsymbol{r}) \,\mathrm{d}^3\boldsymbol{r} = \\ &= A_c \int_{\mathcal{U}_{lmn}} \int_{\mathcal{U}_{ijk}} n^{(1)}_{\boldsymbol{u}} n^{(2)}_{\boldsymbol{u}} \delta[\boldsymbol{r}(\boldsymbol{u}_1) - \boldsymbol{r}(\boldsymbol{u}_2)] v_{12}(\boldsymbol{r}(\boldsymbol{u}_1)) \,\mathrm{d}^4\boldsymbol{u}_1 \,\mathrm{d}^4\boldsymbol{u}_2 \end{aligned} \tag{4.16}$$

where the final expression is obtained by applying the sifting property of eq. (B.1), and $\boldsymbol{u}_1 \in \mathcal{U}_{ijk}$ and $\boldsymbol{u}_2 \in \mathcal{U}_{\ell mn}$. Solving eq. (4.16) requires finding all the values of $\boldsymbol{u}_1$ for which the great circle identified by $\boldsymbol{u}_1$ intersects with the one identified by $\boldsymbol{u}_2$ at $\boldsymbol{r}(\boldsymbol{u}_2)$, as $\boldsymbol{u}_2$ varies in $\mathcal{U}_{\ell mn}$. The Dirac delta collapses 3 dimensions of the inner integral, leaving one of $(h_1, i_1, \Omega_1, u_1)$ free to vary in its respective domain, provided that a solution exists.

The roots in, say, (h_1, i_1, Ω_1) given u_1 and $\boldsymbol{u}_2 = (h_2, i_2, \Omega_2, u_2)^T$ are found as follows: first, from eq. (4.14) it must follow $h_1 = h_2$. Remarkably, if the altitude bins do not overlap, this condition cannot be satisfied and the Dirac delta collapses the whole integral to zero. Because the lattice is constructed from disjoint altitude intervals, the integral is non-zero only for those node pairs that share the same altitude shell.

Then, let

$$s_\delta \equiv \sin\delta = \sin i_2 \sin u_2$$

Since $0 \leq \sin i_1 \leq 1$ everywhere (because $0 \leq i_1 \leq \pi$ by definition), it must hold

$$\begin{cases} \mathrm{sgn}(s_\delta) = \mathrm{sgn}(\sin u_1) \\ |s_\delta| \leq |\sin u_1| \end{cases} \tag{4.17}$$

which also satisfies the condition $|\cos(\alpha - \Omega_1)| \leq 1$ dictated by eq. (4.13). In general, for any u_1 satisfying the above condition, there are two possible inclinations in $[0, \pi]$ that solve eq. (4.12)

$$i_1^+ = \arcsin\left(\frac{s_\delta}{\sin u_1}\right) \qquad i_1^- = \pi - \arcsin\left(\frac{s_\delta}{\sin u_1}\right) \tag{4.18}$$

corresponding to a prograde branch and a retrograde branch, respectively. From eq. (4.13)

$$\alpha = \Omega_2 + \arctan_2(\sin u_2 \cos i_2, \cos u_2) \pmod{2\pi} \tag{4.19}$$

Hence, each inclination $i_1^\pm$ determines one value of RAAN in $[0, 2\pi)$ given by

$$\Omega_1^\pm = \alpha - \arctan_2(\sin u_1 \cos i_1^\pm, \cos u_1) \pmod{2\pi}$$

In conclusion, the system $\boldsymbol{r}(\boldsymbol{u}_1) - \boldsymbol{r}(\boldsymbol{u}_2) = \boldsymbol{0}$ admits exactly two solutions, corresponding to a prograde branch and a retrograde branch. In reality, however, i and Ω are restricted between bin boundaries, further reducing the values of u_1 for which a solution is possible, or entirely removing one of the solutions. For instance, if $I_i \subseteq [0, \pi/2)$ only the prograde branch is reached, while if $I_i \subseteq (\pi/2, \pi]$ only the retrograde branch is. Both branches can only be reached if the inclination bin straddles $\pi/2$. Therefore, the domain of integration of the surviving 1-dimensional inner integral in u_1 is

$$\mathcal{I}(s_\delta) = \{u_1 : i_1^\pm(u_1) \in I_i^{(j)}, \Omega_1^\pm(u_1) \in I_\Omega^{(k)}, \operatorname{sgn}(s_\delta) = \operatorname{sgn}(\sin u_1), |s_\delta| \leq |\sin u_1|\}$$

which depends on s_δ and in turn on $\boldsymbol{u}_2$. Consequently, in the general case in which both branches are reached, eq. (B.3) reduces eq. (4.16) to

$$\dot{\eta}_{12} = A_c \int_{\mathcal{U}_{\ell mn}} \int_{\mathcal{I}(s_\delta)} n_{\boldsymbol{u}}^{(2)} \sum_{\sigma=\pm} \frac{n_{\boldsymbol{u}}^{(1)}(\boldsymbol{u}_1^\sigma)\, v_{12}(\boldsymbol{r}(\boldsymbol{u}_1^\sigma))}{|\det \boldsymbol{J}(\boldsymbol{u}_1^\sigma)|} \, \mathrm{d}u_1 \, \mathrm{d}^4\boldsymbol{u}_2 \tag{4.20}$$

where $\boldsymbol{u}_1^\pm = (h_2, i_1^\pm, \Omega_1^\pm, u_1)^T$, and $\det \boldsymbol{J}(\boldsymbol{u}_1^\pm)$ is the *minor* of the Jacobian of $\boldsymbol{r}(\boldsymbol{u})$ corresponding to the transformation from (h, i, Ω) to (x, y, z)

$$\det \boldsymbol{J} = \left|\frac{\partial(x, y, z)}{\partial(r, \delta, \alpha)}\right| \left|\frac{\partial(r, \delta, \alpha)}{\partial(h, i, \Omega)}\right| = r^2 \cos\delta \cos(\alpha - \Omega)$$

evaluated at the roots. After a few manipulations

$$|\det \boldsymbol{J}(\boldsymbol{u}_1^\pm)| = r^2 \cos\delta |\cos(\alpha - \Omega_1^\pm)| = (R_\oplus + h_2)^2 \sqrt{\sin^2 u_1 - s_\delta^2} \tag{4.21}$$

which is the product of a term that depends only on h_2 and can therefore be pulled out of the inner integral, and a term that depends on u_1 and $\boldsymbol{u}_2$ through s_δ. Finally, eq. (4.20) becomes

$$\dot{\eta}_{12} = A_c \int_{\mathcal{U}_{\ell mn}} \frac{n_{\boldsymbol{u}}^{(2)}}{(R_\oplus + h_2)^2} \sum_{\sigma=\pm} \int_{\mathcal{I}(s_\delta)} \frac{n_{\boldsymbol{u}}^{(1)}(\boldsymbol{u}_1^\sigma)\, v_{12}(\boldsymbol{r}(\boldsymbol{u}_1^\sigma))}{\sqrt{\sin^2 u_1 - s_\delta^2}} \, \mathrm{d}u_1 \, \mathrm{d}^4\boldsymbol{u}_2 \tag{4.22}$$

which is the general formula for the collision rate between two nodes in the lattice phase-space.

4.2.1. Overlap Integral

In general, the above integral admits no closed-form solution. However, if the lattice is sufficiently refined, it is reasonable to assume that the objects' distributions and the relative velocity are constant over one bin. These approximations allow pulling the distributions and the relative velocity outside of the integrals, greatly simplifying the expression

$$\dot{\eta}_{12} = A_c n_{\boldsymbol{u}}^{(1)} n_{\boldsymbol{u}}^{(2)} v_{12} \int_{\mathcal{U}_{\ell mn}} \frac{1}{(R_\oplus + h_2)^2} \underbrace{\sum_{\sigma=\pm} \int_{\mathcal{I}(s_\delta)} \frac{\mathrm{d}u_1}{\sqrt{\sin^2 u_1 - s_\delta^2}}}_{\mathcal{J}(s_\delta)} \mathrm{d}^4\boldsymbol{u}_2$$

This expression can be interpreted as the product of a collision rate density in the (h, i, Ω, u) phase-space times the shared hyper-volume of two overlapping nodes. Thanks to this simplification, the 4-dimensional outer integral can be split into an integral of altitude only and a 3-dimensional integral of $\mathcal{J}(s_\delta)$ (the inner integral over u_1)

$$\int_{I_h^{(\ell)}} \frac{\mathrm{d}h_2}{(R_\oplus + h_2)^2} \int_{I_i^{(m)} \times I_\Omega^{(n)} \times [0,2\pi)} \mathcal{J}(s_\delta)\, \mathrm{d}i_2\, \mathrm{d}\Omega_2\, \mathrm{d}u_2 = \left(\frac{1}{R_\oplus + h_\ell} - \frac{1}{R_\oplus + h_{\ell+1}} \right) \iiint \mathcal{J}(s_\delta)\, \mathrm{d}i_2\, \mathrm{d}\Omega_2\, \mathrm{d}u_2$$

The next step is finding a closed-form solution for $\mathcal{J}(s_\delta)$, and what makes one available is the structure of the domain it is taken over. That structure is the subject of Proof 4.2.1, which establishes three things. First, $\mathcal{I}(s_\delta)$ can be empty for some combinations of the outer and inner nodes' inclination bins. Second, wherever it is non-empty, the domain is a finite union of connected, pairwise disjoint arcs

$$\mathcal{I}(s_\delta) = \bigcup_\kappa [a_\kappa, b_\kappa] \qquad \text{with} \quad [a_\kappa, b_\kappa] \cap [a_\iota, b_\iota] = \emptyset \,\forall\, \kappa \neq \iota$$

Third, the endpoints a_κ and b_κ of those arcs are available in closed form. The three are used in that order below: the first bounds the outer integration, the second is what allows the integral to be split into a sum, and the third supplies the limits of each term in it.

Proof 4.2.1: Structure of the inner domain $\mathcal{I}(s_\delta)$

The vector $\boldsymbol{u}_2$ identifies a specific position on a great circle of radius $R_\oplus + h_2$ on an orbital plane of inclination i_2 and RAAN Ω_2. As the argument of latitude u_2 travels along the circle, s_δ spans all values from $-\sin(i_2)$ in the southern hemisphere to $\sin(i_2)$ in the northern one. Let

$$M_1 = \max(\sin i_1 \,\forall\, i_1 \in I_i^{(j)}) \qquad m_1 = \min(\sin i_1 \,\forall\, i_1 \in I_i^{(j)})$$

and

$$M_2 = \max(\sin i_2 \,\forall\, i_2 \in I_i^{(m)}) \qquad m_2 = \min(\sin i_2 \,\forall\, i_2 \in I_i^{(m)})$$

which are always positive and less than 1. Node (1) can only span declinations from $-\arcsin M_1$ to $\arcsin M_1$. This means that for any $|s_\delta| > M_1$, the roots from eq. (4.18) lay outside $I_i^{(j)}$, violating $\mathcal{I}(s_\delta)$. In practice, this condition restricts the domain of u_2 to only those values that satisfy $|s_\delta| \leq M_1$. If $M_1 > M_2$ then all s_δ are reachable.

For those values of s_δ for which at least one solution exists, since $m_1 \leq \sin i_1 \leq M_1$, it must hold

$$\frac{|s_\delta|}{M_1} \leq \sin u_1 \leq \min\left(\frac{|s_\delta|}{m_1}, 1 \right)$$

which satisfies $|s_\delta| \leq |\sin u_1|$ since $M_1 \leq 1$. Remembering that from eq. (4.17) it must be $\operatorname{sgn}(s_\delta) = \operatorname{sgn}(\sin u_1)$, the above condition restricts the domain of u_1 to a set of one or two connected arcs

$$\mathcal{I}_i(s_\delta) = \begin{cases} \left[\pi + \arcsin \dfrac{|s_\delta|}{M_1}, 2\pi - \arcsin \dfrac{|s_\delta|}{M_1} \right] & \text{if } s_\delta \leq -m_1 \\ \left[\pi + \arcsin \dfrac{|s_\delta|}{M_1}, \pi + \arcsin \dfrac{|s_\delta|}{m_1} \right] \cup \left[2\pi - \arcsin \dfrac{|s_\delta|}{m_1}, 2\pi - \arcsin \dfrac{|s_\delta|}{M_1} \right] & \text{if } -m_1 < s_\delta < 0 \\ \left[\arcsin \dfrac{s_\delta}{M_1}, \arcsin \dfrac{s_\delta}{m_1} \right] \cup \left[\pi - \arcsin \dfrac{s_\delta}{m_1}, \pi - \arcsin \dfrac{s_\delta}{M_1} \right] & \text{if } 0 < s_\delta < m_1 \\ \left[\arcsin \dfrac{s_\delta}{M_1}, \pi - \arcsin \dfrac{s_\delta}{M_1} \right] & \text{if } s_\delta \geq m_1 \end{cases}$$

This is the *inclination-controlled* domain, i.e., the allowed values of u_1 given a restriction on i_1 when Ω_1 is unrestricted. The *RAAN-controlled* domain, i.e., the allowed values of

u_1 given a restriction on Ω_1 when i_1 is unrestricted, is more involved. Start by defining $\psi \equiv \alpha - \Omega_1$. Recall Napier's rule from eq. (4.13)

$$\cos u_1 = \cos\delta \cos\psi$$

which is identical on both inclination branches: a branch is distinguished only by the sign of $\sin\psi$, which is by which side of the meridian through α the ascending node sits. Because $-\pi/2 \leq \delta \leq \pi/2$ by definition, $\cos\delta \geq 0$ always. Inverting the above expression, and once again imposing $\mathrm{sgn}(s_\delta) = \mathrm{sgn}(\sin u_1)$, yields

$\mathrm{sgn}(s_\delta)$	branch	ψ valid on	$u_1(\psi)$	u_1 valid on
$+1$	prograde	$[0, \pi]$	$\arccos(\cos\delta\cos\psi)$	$[\delta, \pi - \delta]$
$+1$	retrograde	$[-\pi, 0]$	$\arccos(\cos\delta\cos\psi)$	$[\delta, \pi - \delta]$
-1	prograde	$[-\pi, 0]$	$2\pi - \arccos(\cos\delta\cos\psi)$	$[\pi + \lvert\delta\rvert, 2\pi - \lvert\delta\rvert]$
-1	retrograde	$[0, \pi]$	$2\pi - \arccos(\cos\delta\cos\psi)$	$[\pi + \lvert\delta\rvert, 2\pi - \lvert\delta\rvert]$

The same intervals are identified by resolving condition (4.17) explicitly. The function $u_1(\psi)$ is even, increasing monotonically in $|\psi|$ on each ψ half-circle. Geometrically, every value of RAAN within a half-circle corresponds to exactly one u_1 along the branch. Now, since Ω_1 is restricted between bin boundaries $[\Omega_j, \Omega_{j+1}]$, also $\psi \in [\alpha - \Omega_{j+1}, \alpha - \Omega_j]$. Normalizing this interval into $(-\pi, \pi]$ produces one connected arc if the RAAN bin does not straddle $\alpha + \pi$, or two arcs if it does. Call it I_ψ. The values of ψ that are allowed are obtained by intersecting I_ψ with the branch's valid ψ half-circle producing

$$I_\psi^+ = \begin{cases} I_\psi \cap [0, \pi] & \text{for } s_\delta > 0 \\ I_\psi \cap [-\pi, 0] & \text{for } s_\delta < 0 \end{cases} \qquad I_\psi^- = \begin{cases} I_\psi \cap [-\pi, 0] & \text{for } s_\delta > 0 \\ I_\psi \cap [0, \pi] & \text{for } s_\delta < 0 \end{cases}$$

for the prograde and retrograde branches, respectively. Either of these may be empty (the RAAN bin does not reach that branch), or one connected arc, or two. Because $u_1(\psi)$ is monotonic in ψ on each half-circle, each connected sub-interval of $I_\psi^\pm$ maps to a connected u_1 sub-interval. Once again, each branch can only be reached if the inclination bin allows it. In conclusion, the RAAN-controlled domain $\mathcal{I}_\Omega(s_\delta)$ is made up of the images of each sub-interval of I_ψ^+, or I_ψ^-, or both. Finally, the inclination constraint and the RAAN constraint are combined to obtain

$$\mathcal{I}(s_\delta) = \mathcal{I}_i(s_\delta) \cap \mathcal{I}_\Omega(s_\delta)$$

which is a set of connected and disjoint intervals.

The property of additivity of integrals states that the integral over the union of disjoint sets is equal to the sum of the integrals over each set. Since Proof 4.2.1 establishes that the arcs are disjoint and gives their endpoints, the inner integral becomes

$$\mathcal{J}(s_\delta) = \int_{\mathcal{I}(s_\delta)} \frac{\mathrm{d}u_1}{\sqrt{\sin^2 u_1 - s_\delta^2}} = \sum_\kappa \int_{a_\kappa}^{b_\kappa} \frac{\mathrm{d}u_1}{\sqrt{\sin^2 u_1 - s_\delta^2}}$$

Applying the substitutions $t = \cos u_1 / \sqrt{1 - s_\delta^2}$ and $\sin\varphi = t$ results in

$$\int_{a_\kappa}^{b_\kappa} \frac{\mathrm{d}u_1}{\sqrt{\sin^2 u_1 - s_\delta^2}} = \int_{t(b_\kappa)}^{t(a_\kappa)} \frac{\mathrm{d}t}{\sqrt{(1-t^2)[1-(1-s_\delta^2)t^2]}} =$$

$$= \int_{\varphi(b_\kappa)}^{\varphi(a_\kappa)} \frac{\mathrm{d}\varphi}{\sqrt{1-(1-s_\delta^2)\sin^2\varphi}} = F(\varphi(a_\kappa), k^2) - F(\varphi(b_\kappa), k^2) \tag{4.23}$$

where

$$F(\varphi, m) = \int_0^{\varphi} \frac{\mathrm{d}\varphi}{\sqrt{1 - m \sin^2 \varphi}}$$

is the *incomplete elliptic integral of the first kind* with modulus $k^2 = 1 - s_\delta^2$ and amplitude

$$\varphi(a_\kappa) = \operatorname{sgn}(s_\delta) \arcsin\left(\frac{\cos a_\kappa}{k^2}\right)$$

and similarly for b_κ. Therefore, $\mathcal{J}(s_\delta)$ is the algebraic sum of incomplete elliptic integrals of the first kind

$$\mathcal{J}(s_\delta) = \sum_i \left| F(\varphi(a_\kappa), k^2) - F(\varphi(b_\kappa), k^2) \right| \tag{4.24}$$

whose modulus and amplitude both depend on i_2, Ω_2, and u_2 as shown in Proof 4.2.1.

Numerical Quadrature

The problem of finding the shared hyper-volume of the two nodes reduces to computing the 3-dimensional integral of $\mathcal{J}(s_\delta)$ over (i_2, Ω_2, u_2). This does not have a closed-form solution and must be integrated numerically. The scheme used here is a tensor-product Gauss-Legendre rule over the outer variables, reviewed in appendix C.1, with Q_i, Q_Ω and Q_u points respectively, while the inner integral over u_1 is evaluated exactly through eq. (4.24). The argument of latitude is given twice as many points as either angle of the orbital plane, $Q_u = 2Q_i = 2Q_\Omega$, because it is the fastest-varying of the three: the inclination and the RAAN of the outer node are confined to one bin each, a few degrees or tens of degrees wide, whereas u_2 sweeps a full revolution, over which the integrand crosses the equator twice and reaches the highest declination the plane allows in between. The two nodes of a pair therefore play different numerical roles: whichever node is the outer one is sampled on the product grid, whichever is the inner one is computed exactly in closed-form. In principle, swapping the two nodes should not change the result in the continuum. However, the finite-resolution scheme breaks a symmetry the integral itself possesses, and the amount by which it does so is a lever on the quadrature error.

Two singularities of the integrand break the symmetry. The first is *polar*: when the outer node straddles $\pi/2$, at the poles $|s_\delta| = 1$ and the Jacobian $|\det \boldsymbol{J}|$ vanishes in the denominator. In that configuration an overlap is only possible if the inner node's inclination also straddles $\pi/2$, and then only at the poles, which collapses the inclination-controlled domain to a single point, since $|s_\delta| \leq M_1$ requires $M_1 = 1$, so the outer quadrature never sees it. For points that land close to it, where the inner integral does not collapse, the elliptic integral resolves the limit cleanly: as $|s_\delta| \to 1$ the modulus $k^2 \to 0$, the amplitude $\varphi \to \pi/2$, and the integral converges to its amplitude, $F(\pi/2, 0) = \pi/2$. The second feature is *equatorial* and is genuinely singular. At the equator $s_\delta \to 0$, $k^2 \to 1$ and $F(\varphi, 1) = \operatorname{artanh}(\sin\varphi)$, which diverges at the ascending and descending nodes $u_1 = 0$ and $u_1 = \pi$, as is also visible from $|\det \boldsymbol{J}| = |\sin u_1|$ in the denominator. Written in outer coordinates through $\sin\varphi = \sin i_2 \sin u_2$, the divergence sits on the lines $u_2 \in \{0, \pi, 2\pi\}$ for every i_2, so its location in the outer grid is known in advance.

Gauss-Legendre converges exponentially on a smooth integrand but only algebraically across a logarithmic singularity interior to the interval, and this is what limits the scheme. The remedy is to split the outer u_2 integration at $u_2 = \pi$, so that the three singular abscissae fall on sub-interval endpoints, where a Gauss-Legendre rule never places a node, and each sub-interval is smooth on its interior. The effect is measured on the six probe pairs of table 4.1, chosen to span the range of difficulty, from a pair of nodes sharing an inclination bin to one that touches both singular inclinations at once.

Table 4.1: The six probe pairs, on a lattice of six inclination bins of 30° and 72 RAAN bins of 5°. Node 1 is the inner node of the quadrature and node 2 the outer one.

Node 1		**Node 2**		
$I_i^{(j)}$	$I_\Omega^{(k)}$	$I_i^{(m)}$	$I_\Omega^{(n)}$	
$[0, 30)$	$[0, 5)$	$[0, 30)$	$[5, 10)$	same inclination, adjacent RAAN
$[0, 30)$	$[0, 5)$	$[30, 60)$	$[0, 5)$	adjacent inclination, equatorial
$[30, 60)$	$[0, 5)$	$[60, 90)$	$[0, 5)$	adjacent inclination, interior
$[0, 30)$	$[0, 5)$	$[90, 120)$	$[0, 5)$	both singular inclinations
$[60, 90)$	$[50, 55)$	$[120, 150)$	$[250, 255)$	generic mixed
$[30, 60)$	$[0, 5)$	$[150, 180]$	$[15, 20)$	prograde with retrograde

Figure 4.2 draws the six of them. Each node is the region swept on the sphere by every orbital plane whose inclination and RAAN lie in its bin, so it appears as a band, and the pair's shared hyper-volume is what the two bands have in common. The range of difficulty is visible in the geometry: the first pair differs by one RAAN bin and its two bands almost coincide, whereas the pair touching both singular inclinations meets nearly at right angles and shares very little.

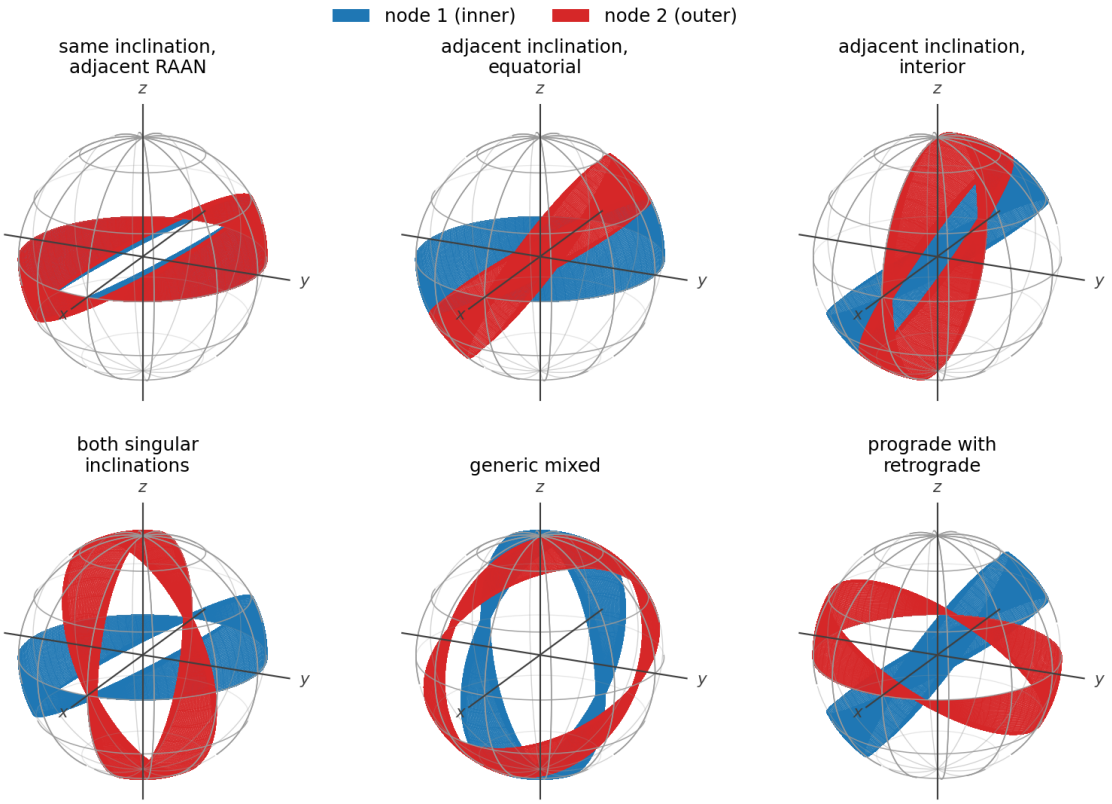

Figure 4.2: The six probe pairs of table 4.1 on the unit sphere. Each band is the region swept by the orbital planes of one node. The globe is drawn as a wireframe, so the far side of each band remains visible.

Figure 4.3 shows both rules on those pairs. The error plotted there is the relative deviation of the swap-averaged value from a single common reference, taken to be the swap-averaged split rule at $Q_i = 96$, the finest resolution computed and therefore the best available estimate of the exact hyper-volume; the same reference is used for both rules, so the two panels are directly comparable. Because both the plotted value and the reference are swap-averaged, the part of the error that changes sign when the two members of a pair are exchanged has already cancelled,

and what remains is the error common to both orientations. At $(Q_i, Q_\Omega, Q_u) = (48, 48, 96)$ the split reduces the worst-pair error from 26 % to 0.99 %, and on the pair touching both singular inclinations the improvement is more than an order of magnitude at fixed cost.

The accuracy of this integral is therefore governed by where the nodes are placed as much as by how many there are. With the equatorial singularity in the interior of the grid, no resolution on the ladder brings the worst pair near 1 %; with the same node count and the singularity moved onto a sub-interval endpoint, every pair is at or below it. That one per cent on the worst pair is the accuracy the overlap tensor is built to, and the split is what makes it affordable.

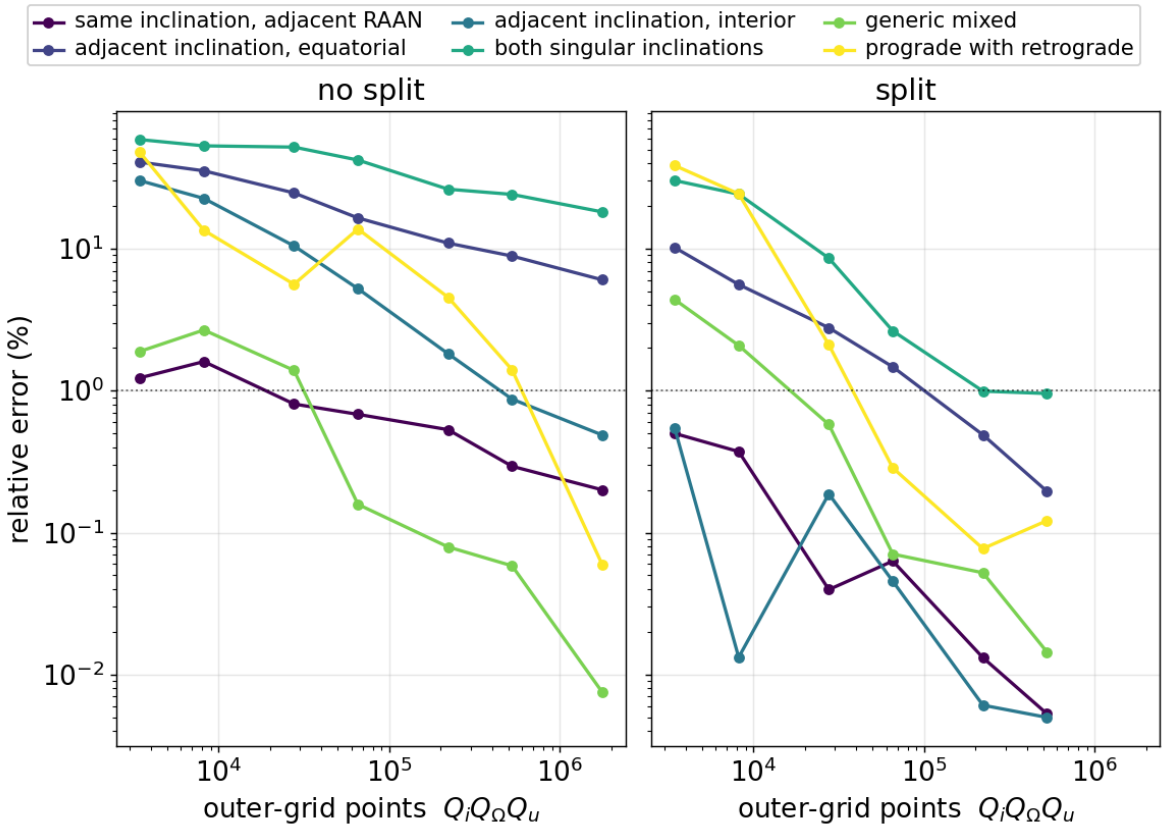

Figure 4.3: Effect of splitting the outer u_2 integration at the equatorial logarithmic singularity, on the six probe pairs of table 4.1. Left, the un-split rule, with the singularity interior to the grid. Right, the split rule, with the singularity on a sub-interval endpoint. The dotted line marks 1 %.

The split does not restore the symmetry between the two orientations of a pair, because it is the placement of the outer grid, not the choice of split points, that differs between them. Rather than pick an ordering, which would be arbitrary since neither is the better one, both orientations are computed and averaged,

$$\Psi_{\text{sym}} = \frac{1}{2}\left(\Psi_{12} + \Psi_{21}\right)$$

where Ψ_{12} denotes the hyper-volume computed with node 1 as the inner node and node 2 as the outer one, and Ψ_{21} the same pair with the roles exchanged. Writing each orientation as the truth plus an error, $\Psi_{12} = \Psi + \varepsilon_{12}$ and $\Psi_{21} = \Psi + \varepsilon_{21}$, the average carries the symmetric part $(\varepsilon_{12} + \varepsilon_{21})/2$ and cancels the antisymmetric part $(\varepsilon_{12} - \varepsilon_{21})/2$ exactly. The observable asymmetry $|\Psi_{12} - \Psi_{21}|$ is proportional to the antisymmetric part alone, so it measures precisely the component that averaging removes and is blind to the error that survives it. It is therefore not an error estimate for the averaged value, and it is not used as one. The error of Ψ_{sym} has to be established by other means, which is what the remainder of this section does.

That the averaging is affordable, and indeed free, follows from a symmetry of the integral.

The integrand of eq. (4.16) depends on the two orbital planes only through their relative geometry, so it is invariant under a rotation of both planes by a common angle about the polar axis, an isometry of the sphere under which Ω_1 and Ω_2 shift equally and no other variable is touched. The RAAN bins are uniform, so a rotation by a whole number s of bin widths carries bin boundaries onto bin boundaries, and for the node pair with inclination bins m_1, m_2 and RAAN bins n_1, n_2,

$$\Psi_{12}(m_1, n_1 + s; m_2, n_2 + s) = \Psi_{12}(m_1, n_1; m_2, n_2) \qquad \forall\, s \in \mathbb{Z} \tag{4.25}$$

with the RAAN indices taken modulo N_Ω. At fixed inclinations the overlap is therefore circulant in (n_1, n_2), depending on the two RAAN bins only through their difference. Choosing $s = -n_1$ in eq. (4.25) and writing $\Delta n = (n_2 - n_1) \bmod N_\Omega$, the whole table is generated by the third-order tensor

$$\Phi_{m_1 m_2 \Delta n} = \Psi_{12}(m_1, 0; m_2, \Delta n) \qquad \Longrightarrow \qquad \Psi_{12}(m_1, n_1; m_2, n_2) = \Phi_{m_1 m_2 [(n_2 - n_1) \bmod N_\Omega]} \tag{4.26}$$

The reduction is exact, and exact at finite quadrature rather than only in the continuum, because the Gauss-Legendre abscissae of a shifted bin are the shifted abscissae of the original one. On a lattice of $N_i = 6$ inclination bins and $N_\Omega = 72$ RAAN bins, the one the timings below are quoted for, a naive enumeration of ordered bin pairs would require $(N_i N_\Omega)^2 = 186\,624$ evaluations of the outer integral, while the tensor of eq. (4.26) needs $N_i^2 N_\Omega = 2592$, a reduction by a factor of exactly N_Ω. The swap-average then becomes an operation internal to that tensor,

$$\Phi^{\text{sym}}_{m_1 m_2 \Delta n} = \frac{1}{2}\left(\Phi_{m_1 m_2 \Delta n} + \Phi_{m_2 m_1 [(-\Delta n) \bmod N_\Omega]}\right)$$

so both orientations of every physical pair are entries that have already been computed, and the averaging costs no additional evaluations. A further factor of two improvement in computational cost would be available from the inclination mirror $i \to \pi - i$, and it is deliberately not taken, as that identity is exact only in the continuum and holds to about 1×10^{-8} at finite precision, independently of the quadrature order, because the sign flip routes the inner integral through different branches of the elliptic integral.

A quadrature scheme of this kind converges with a definite order p, which is the exponent of the leading term in the asymptotic expansion of the error with respect to the number of quadrature points. If $\Psi_{\text{sym}}(Q_i)$ is the swap-averaged value computed at order Q_i, then for sufficiently large Q_i it can be expanded as

$$\Psi_{\text{sym}}(Q_i) \simeq \Psi + C Q_i^{-p} \tag{4.27}$$

for some constant C and some order p, where Ψ is the value the scheme would return with the grid refined without limit, which is the exact hyper-volume. Fitting the three unknowns of eq. (4.27) to the sequence of values computed on a ladder of resolutions therefore estimates that limit without ever reaching it, and the recovered p measures how fast the scheme converges. This is Richardson extrapolation, and it supplies the second of the two error measures of table 4.2, both taken over the same six probe pairs. The first, err_{self}, is the deviation of $\Psi_{\text{sym}}(Q_i)$ from the finest resolution actually computed, at $Q_i = 96$, and is therefore a self-convergence measure. The two agree to within a factor of about one and a half everywhere.

Table 4.2: Cost and accuracy of the swap-averaged overlap tensor against quadrature order, at $(Q_i, Q_\Omega, Q_u) = (Q, Q, 2Q)$. Errors are the worst of the six probe pairs of table 4.1.

Q	**outer points**	**per integral**	**full tensor**	err_{self}	err_{extr}
12	3456	5.2 ms	1.7 s	38.5 %	38.5 %
16	8192	12 ms	3.9 s	24.3 %	24.4 %
24	27 648	40 ms	13 s	8.6 %	8.2 %
32	65 536	93 ms	30 s	2.6 %	2.2 %
48	221 184	320 ms	104 s	0.99 %	0.67 %
64	524 288	790 ms	255 s	0.95 %	0.51 %
96	1 769 472	2600 ms	840 s	n/a	0.45 %

Every one of the six probe pairs of table 4.1 converges monotonically, and by $Q_i = 48$ all of them are at or below 1 %, with five of the six below 0.1 %. The fitted orders separate by which singular inclination a pair touches: the two pairs that touch neither saturate the bound placed on the fit, at $p \approx 8$, and so have no resolvable algebraic order, the pair touching $\pi/2$ sits at $p \approx 4.4$, and the pairs whose inclination bins reach down to the equator are the slowest, at $p = 2.1$ and $p = 1.0$. The limiter in absolute terms is nevertheless the pair touching both singular inclinations, which is also the pair with the smallest overlap. It floors at about 1 % and barely improves from $Q_i = 48$ to $Q_i = 64$, which is the signature of a residual cusp-proximity effect that outer resolution cannot remove. Closing it would require adaptive refinement near the bin boundaries, which is not pursued here. Figure 4.4 plots the same data as accuracy against build cost and shows a clear knee at $Q_i = 48$: accuracy improves steeply up to that point and flattens beyond it, since past $Q_i = 48$ the only pair still improving is the cusp-limited one. Going to $Q_i = 64$ buys a reduction of the worst error from 0.99 % to 0.95 % for 2.4 times the build time. The production tensor is consequently built at $(48, 48, 96)$ with the outer split enabled.

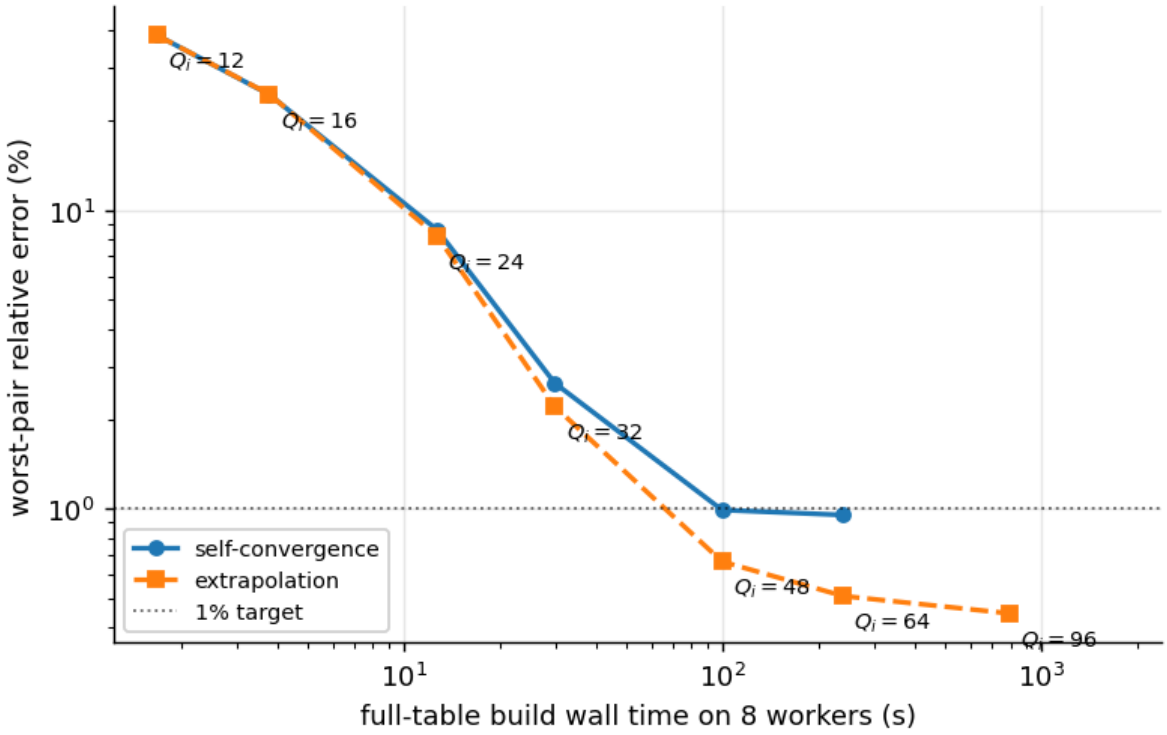

Figure 4.4: Worst-pair accuracy of the swap-averaged overlap tensor against the wall time of a full tensor build on 8 workers. The self-convergence curve stops at $Q_i = 64$ because it is identically zero at its own reference.

4.2.2. Relative Velocity

The relative velocity between two nodes can be approximated with a representative value at the bins' center (Giudici, 2021) or by an averaged value (Wang et al., 2025). In general, the relative velocity between two orbits with orbital velocities v_1 and v_2, respectively, is

$$v_{rel} = \sqrt{v_1^2 + v_2^2 - 2v_1 v_2 \cos\theta_{12}} \tag{4.28}$$

where θ_{12} is the angle between the two velocity vectors. If the orbits are circular, this only depends on the orientation of the orbital planes, and not on the specific position of the objects along the orbit, and

$$\cos\theta_{12} = \sin i_1 \sin i_2 \cos(\Omega_1 - \Omega_2) + \cos i_1 \cos i_2 \tag{4.29}$$

To evaluate eq. (4.28) at the bins' center, first recall that each node is the union of great circles with altitude, inclination, and RAAN varying between bins' edges, and that the collision rate is

non-zero only for those node pairs that share the same altitude shell. Therefore, v_1 and v_2 are identical, and equal to the circular velocity

$$v_1 = v_2 = v^* = \sqrt{\frac{\mu}{R_\oplus + h^*}}$$

where the $*$ indicates that the quantity is evaluated at the center of the bin. Expression (4.28) simplifies to

$$v_{rel} = 2v^* \sin\frac{\theta_{12}^*}{2} \tag{4.30}$$

One consequence of this approach is that all the collisions between objects belonging to the same node are suppressed, because the relative angle θ_{12}^*, and hence the relative velocity, of a node and itself is always zero. This is an accepted limitation since all objects within the same sufficiently narrow bin have nearly tangential velocities in the same direction and hence very low risks of collisions. On the other hand, the largest relative velocity is found for those nodes whose objects travel tangentially in opposite directions, posing a risk of head-on collisions. The effect of this approximation on the simulation results is assessed in detail in section 6.2.3

4.2.3. Average Impact Cross-sectional Area

The average impact cross-sectional area between the objects belonging to two distinct nodes S_i and S_j is

$$A_c = \frac{\pi}{4x_{S_i}x_{S_j}} \sum_{q=1}^{S_i}\sum_{p=1}^{S_j}(d_{S_i}^p + d_{S_i}^q)^2 =$$

$$= \frac{\pi}{4x_{S_i}x_{S_j}} \left[x_{S_j} \sum_{q}^{S_i}(d_{S_i}^q)^2 + x_{S_i} \sum_{p}^{S_j}(d_{S_j}^p)^2 + 2\sum_{q}^{S_i} d_{S_i}^q \sum_{p}^{S_j} d_{S_j}^p \right] \tag{4.31}$$

where $d_{S_i}^q$ is the diameter of the q-th object in node S_i and $d_{S_j}^p$ is the diameter of the p-th object in node S_j, as given by (Wang et al., 2025). The collisions between objects belonging to the same node are suppressed, since the relative velocity of a node and itself is always zero, as discussed in the previous section, so no expression for that case is provided.

4.3. Natural Dynamics

Natural dynamics are applied to each derelict and fragment object individually. Every object carries its own semi-major axis, eccentricity, inclination, argument of periapsis and RAAN, while the nodes are a view of the population rather than its state, rebuilt from the propagated elements at the beginning of every step, so no orbital information is ever stored at bin resolution. The propagation uses the Martinusi analytical propagator (Martinusi et al., 2015), which provides a fully time-explicit closed-form solution for the mean orbital elements under the combined effect of J_2 and atmospheric drag, valid for near-circular orbits, and which connects to the first-order secular rates of the Lagrange planetary equations under J_2 alone.

The closed form assumes that the atmospheric density is constant. Within one averaging orbit the assumption is immediate, since a near-circular orbit stays at near-constant altitude, and across a step it becomes a piecewise-constant scheme: the density is read at each object's mean altitude at the beginning of the step from a tabulated exponential model, and held fixed for the duration of the step. Two approximations are therefore in play. The first is the piecewise-constant treatment itself, which is the same scheme adopted by Jang et al. (2024); there, the analytical propagator was validated against direct numerical integration of the same drag and J_2 dynamics over a population spanning 200 km to 2000 km and a propagation of one year, comparing the number of objects surviving, the decay in semi-major axis, the secular change in right ascension

and the time to decay, each as a function of initial altitude. The distributions agree closely, with the analytical solution underestimating the decay in semi-major axis and overestimating both the secular change in right ascension and the time to decay. That validation is adopted here rather than repeated, since the configuration is the same. The second approximation is the choice of a static mean-activity density profile, which carries no dependence on epoch and so does not represent the 27-day solar rotation or the 11-year cycle. This is a deliberate restriction and not a consequence of the propagator: the same scheme accommodates a time-varying profile by simply refreshing the density at each step from a solar-activity-driven model, as is done by Jang et al. (2024). The impact of different atmospheric models on the simulation is deferred to future work.

Let

$$\alpha_k = \frac{e_k}{\sqrt{a_k}} \qquad \beta_k = \frac{\sqrt{3}}{2} e_k \qquad C_k = \frac{1}{2}\frac{C_D A}{m}\rho_k \qquad c = \cos i \qquad k_2 = \frac{\mu_\oplus J_2 R_{eq}}{2}$$

Then, assuming step-wise constant air density ρ_k and sufficiently small eccentricity such that $\mathcal{O}(e^4) = 0$

$$a_{k+1} = \frac{a_k}{\beta_k^2}\tan^2[\arctan\beta_k - \beta_k n_k a_k C_k(t_{k+1} - t_k)] \tag{4.32a}$$

$$e_{k+1} = \frac{2}{\sqrt{3}}\tan[\arctan\beta_k - \beta_k n_k a_k C_k(t_{k+1} - t_k)] \tag{4.32b}$$

$$\omega_{k+1} = \omega_k + \frac{3k_2(5c^2-1)}{16\mu_\oplus}\frac{1}{C_k}\left[\frac{5\alpha_k^2}{2}\frac{1}{\tau^2} + \frac{4}{3\tau^3}\right|_{\tau=a_k}^{\tau=a_{k+1}} \tag{4.32c}$$

$$\Omega_{k+1} = \Omega_k - \frac{3k_2 c}{8\mu_\oplus}\frac{1}{C_k}\left[\frac{5\alpha_k^2}{2}\frac{1}{\tau^2} + \frac{4}{3\tau^3}\right|_{\tau=a_k}^{\tau=a_{k+1}} \tag{4.32d}$$

Eqs. (4.32a) and (4.32b) only depend on drag, via the term C_k. If this term approaches 0, the classical averaged variational equations for J_2 only are recovered (Martinusi et al., 2015). The propagated elements are then mapped back onto the lattice. Inclination is unchanged by either perturbation, so an object keeps its inclination bin throughout, whereas drag lowers the semi-major axis and J_2 causes RAAN to drift, and both of these eventually carry an object across a bin boundary. Objects whose perigee falls below the re-entry altitude are removed from the population, the threshold being set at 100 km.

An object on an eccentric orbit does not belong to a single altitude shell. Its orbital period is at most a few hours, orders of magnitude shorter than the time step, so over one step it sweeps every shell between perigee and apogee many times: it is genuinely present in all of them, and the physically meaningful quantity is the fraction of the period it spends in each. For an object q of semi-major axis a_q and eccentricity e_q, the fraction of its period spent in the ℓ-th altitude shell is

$$w_{q,\ell}^{(k)} = \frac{1}{T_q}\int_0^{T_q} \mathbb{1}\{r_q(t) \in [r_\ell, r_{\ell+1})\}\,\mathrm{d}t = \frac{M(\theta_{q,\ell+1}) - M(\theta_{q,\ell})}{\pi} \tag{4.33}$$

where T_q is the orbital period, $r_\ell = h_\ell + R$ is the radius of the ℓ-th shell boundary, $\theta_{q,\ell}$ is the true anomaly at which the orbit crosses that boundary,

$$\cos\theta_{q,\ell} = \frac{1}{e_q}\left(\frac{p_q}{r_\ell} - 1\right) \qquad p_q = a_q\left(1 - e_q^2\right)$$

taken as 0 for boundaries at or below perigee and π for boundaries at or above apogee, and $M(\theta) = E(\theta) - e_q \sin E(\theta)$ is the mean anomaly at true anomaly θ (Wang et al., 2025). The normalization is π rather than 2π because the orbit is symmetric about the line of apsides: each shell is swept twice per revolution, once ascending and once descending, and the two passages are of equal duration. For a circular orbit the expression collapses to a single unit weight in

the shell containing a_q, which is the case for every station-keeping payload, so the fractional treatment only ever engages where it is needed, on fragments and on decayed derelicts.

Inclination and RAAN carry no analogous weight: an object occupies exactly one inclination bin and one RAAN bin at any instant. The occupancy of the node identified by the orbit bin (ℓ, m, n) and the specie S is therefore

$$x_{S_i}(t_k) = \sum_{q \in S} w_{q,\ell}^{(k)}\, \delta_{m\, m_q(t_k)}\, \delta_{n\, n_q(t_k)} \tag{4.34}$$

where $m_q(t_k)$ and $n_q(t_k)$ are the inclination and RAAN bin indices the object falls in at t_k, and δ is the Kronecker delta, 1 when $m_q(t_k) = m$ or $n_q(t_k) = n$ and 0 otherwise. This is the quantity the ε terms of eqs. (4.6) to (4.9) account for, and it takes the same form for every specie. The residence weight is the only specie-dependent ingredient: eq. (4.33) holds for any object coasting on an unpowered arc, while the rising and de-orbiting satellites, whose semi-major axis is being driven by thrust, carry a weight of their own, derived in the next section and substituted into eq. (4.34) unchanged. The above expression is a sum over all individual contributions of each object q to node i

$$c_{q,i}^{(k)} = w_{q,\ell}^{(k)}\, \delta_{m\, m_q(t_k)}\, \delta_{n\, n_q(t_k)}$$

Let $\Delta_{q,i}^{(k)} = c_{q,i}^{(k+1)} - c_{q,i}^{(k)}$ be the change of that term when the object's elements are advanced by eq. (4.32), or by a powered transfer, alone. Then

$$\varepsilon_{S_i}^{+} = \sum_{q} \max\left(0, \Delta_{q,i}^{(k)}\right) \qquad\qquad \varepsilon_{S_i}^{-} = \sum_{q} \max\left(0, -\Delta_{q,i}^{(k)}\right)$$

where both sums run over the objects of specie S that are present at both epochs and keep their specie across the step. An object that is launched, generated as a fragment, destroyed in a collision, or moved to another specie by a PMD maneuver or a failure, enters and leaves the node counts through the corresponding ζ or Δ terms of eqs. (4.5) to (4.9), and is excluded here so that it is not counted twice. Atmospheric re-entry is the one exception and is retained, with $c_{q,i}^{(k+1)} = 0$ for an object whose perigee drops below the re-entry altitude: eqs. (4.5) to (4.9) carry no separate term for it, and it is by nature a natural-dynamics loss. Consequently $x_{S_i}(t_{k+1}) - x_{S_i}(t_k)$ reduces to $\varepsilon_{S_i}^{+} - \varepsilon_{S_i}^{-}$ only when no other event fires on that node during the step. Summed over the nodes of a specie, $\sum_i \Delta_{q,i}^{(k)}$ vanishes for every object whose orbit remains within the lattice, so natural dynamics redistribute objects between nodes without creating or destroying them, and the only unmatched outflow is the one at the lower boundary.

The treatment of the altitude axis is the one point where the model departs from NESSY, which assigns each object to exactly one shell, sampled from the categorical distribution whose weights are the residence times (Wang et al., 2025). The two treatments agree in the mean, since $w_{q,\ell}^{(k)}$ is precisely the probability of that draw, but they do not agree in variance. Writing L_q for the sampled shell index, eq. (4.33) is the conditional expectation $w_{q,\ell}^{(k)} = \mathbb{E}\left[\mathbb{1}\{L_q = \ell\} \mid a_q, e_q\right]$, so by the law of total variance

$$\operatorname{Var}\left(\sum_q \mathbb{1}\{L_q = \ell\}\right) = \operatorname{Var}\left(\sum_q w_{q,\ell}^{(k)}\right) + \mathbb{E}\left[\operatorname{Var}\left(\sum_q \mathbb{1}\{L_q = \ell\} \,\middle|\, a_q, e_q\right)\right]$$

The sampled count therefore has a variance no smaller than the fractional one, and strictly larger as soon as a single object straddles a shell boundary. Over an ensemble of Monte Carlo runs the sampled assignment converges to the same occupancy the fractional treatment carries exactly, at the cost of one extra source of scatter that has nothing to do with the physics being estimated. While it is true that the fractional occupancy model is strictly less noisy, the variance removed is only that of the shell assignment: collisions, failures and fragment generation remain sampled, and they dominate the residual spread of the ensemble. Moreover, the fractional occupancy is compatible with the collision rate only because objects in distinct shells are never paired, so splitting an object across shells never lets it collide with itself.

That last property is exactly what forbids extending the same treatment to the RAAN axis. Because of J_2, a derelict at 550 km and 53° drifts through roughly 4.5° of RAAN per day, so with planes resolved at 5° it crosses about one bin per day and the drift is no longer slow compared to the time step. A fractional count would be the natural remedy, and it is not admissible: the fractions would sit in different RAAN bins at the same altitude, which is precisely the configuration that does produce a non-zero collision rate, and the object would be modeled as colliding with itself per that rate. Sampling the bin from a categorical distribution is not admissible either, for a reason that is physical rather than numerical: an object does not have a chance of being in any of several planes, it is in exactly one, and randomizing over them would manufacture encounter geometries that never occur. The bin assignment is consequently kept exact, and the model resolves the drift at the granularity of its time step, which is the residual approximation. All the weights are computed under a frozen-at-start assumption: within one step the orbit, and with it the drift and decay rates, are held at their values at the beginning of the step.

A related limitation affects the decay of the lowest objects. An object under strong enough drag can cross several altitude shells within a single step, which the one-step change of eq. (4.34) attributes as a single transition and so under-resolves, and which is a distinct effect from the multi-shell occupancy of eq. (4.33). The objects concerned are those already low enough to be within a step or two of re-entry, so the misattributed flux is a small fraction of a population that is about to leave the simulation entirely, but the magnitude has not been quantified here.

4.4. Orbit Rising and De-orbiting

Often, a satellite does not begin its life in its operational orbit. The launcher delivers it to a lower injection orbit, from which it raises itself under its own propulsion over a period of weeks to months, and at the end of its operational life a satellite that disposes of itself lowers itself again until it re-enters. Both transfers sweep through every shell in between, and both are resolved explicitly here rather than treated as instantaneous, which is the purpose of the R and D species. The treatment follows MOCAT-6RD (Gusmini et al., 2024) and rests on three assumptions. Satellites are injected into near-circular orbits and dispose from circular operational orbits, so $e \approx 0$ throughout the transfer. They are injected at the operational inclination and RAAN, or acquire them impulsively at once, so the transfer varies the semi-major axis alone and there is no drifting phase at a parking orbit. And the thrust is continuous and purely tangential, with neither coasting nor plane change, which is the strategy that yields the largest rate of change of semi-major axis and under which the secular variations of eccentricity and RAAN both vanish.

Under those assumptions the semi-major axis evolves as

$$a(t) = \left(\frac{1}{a_{\text{inj}}} + C_1 t^2 - \sigma\, C_2 t \right)^{-1} \tag{4.35}$$

where a_{inj} is the semi-major axis at the start of the maneuver and, for $e = 0$,

$$C_1 = \frac{\epsilon^2}{\mu_\oplus} \qquad C_2 = \frac{2\epsilon}{\sqrt{a_{\text{inj}}\mu_\oplus}} \qquad \sigma = \begin{cases} +1 & \text{rising} \\ -1 & \text{de-orbiting} \end{cases}$$

with ϵ the thrust acceleration, taken constant at 5×10^{-8} km/s^2. That value is the one adopted by Gusmini et al. (2024), who obtained it by fitting the historical TLE record of the OneWeb and Starlink orbit-raising maneuvers, so it is a constellation-representative effective acceleration rather than a manufacturer figure for any one propulsion system. Let τ be the time elapsed since the maneuver started, so that

$$\frac{1}{a(\tau)} = \frac{1}{a_{\text{inj}}} + C_1 \tau^2 - \sigma\, C_2 \tau \qquad \text{and} \qquad \dot{a}(\tau) = a^2(\tau)(\sigma\, C_2 - 2 C_1 \tau) \tag{4.36}$$

For a rising satellite ($\sigma = 1$) the semi-major axis grows only while $\tau < \tau_{\text{crit}} = C_2/(2C_1)$, and shrinks again beyond it. That branch is an artifact of extrapolating the expression far past its

intended range rather than a physical turning point: for a typical LEO injection at $a_{\text{inj}} = 6800\,\text{km}$ it sits at $\tau_{\text{crit}} \approx 4.8$ years, two orders of magnitude beyond the one to four months a constellation transfer actually takes, so $a(\tau)$ is monotonic over every interval of interest. For a de-orbiting satellite ($\sigma = -1$) it decreases monotonically for all τ.

Because $\dot{a}$ is not constant, the time an object spends in a shell cannot be read off as a shell width divided by a nominal rate, and eq. (4.35) has to be inverted. Solving the quadratic for the time at which a given semi-major axis is reached gives

$$\tau(a) = \sigma \frac{C_2 - \sqrt{C_2^2 - 4C_1 D}}{2C_1} \qquad \text{where} \quad D = \frac{1}{a} - \frac{1}{a_{\text{inj}}} \tag{4.37}$$

A maneuvering satellite enters the network through the occupancy law of eq. (4.34) exactly as any other object does, and only the definition of the residence weight changes. What is averaged over is the step rather than an orbital period, since a spiraling object crosses each shell once on its way through instead of sweeping it twice per revolution, but the weight plays the identical role and is substituted into eq. (4.34) in place of eq. (4.33). The fraction of the step Δt_k that object q spends in the shell $[h_\ell, h_{\ell+1})$ is

$$w_{q,\ell}^{(k)} = \frac{1}{\Delta t_k} \int_{\tau_k}^{\tau_{k+1}} \mathbb{1}\{a(\tau) \in [a_\ell, a_{\ell+1})\}\, \mathrm{d}\tau$$

where τ_k is the time elapsed since the beginning of the maneuver by step k and $a_\ell = R_\oplus + h_\ell$. Since $a(\tau)$ is monotonic, the object enters and leaves each shell exactly once and the integral collapses to the overlap of $[\tau_k, \tau_{k+1}]$ with $[\tau(a_\ell), \tau(a_{\ell+1})]$ for a rising satellite, or with $[\tau(a_{\ell+1}), \tau(a_\ell)]$ for a de-orbiting one,

$$w_{q,\ell}^{(k)} = \frac{\max(0, \min(\overline{\tau}_\ell, \tau_{k+1}) - \max(\underline{\tau}_\ell, \tau_k))}{\Delta t_k} \tag{4.38}$$

where $\overline{\tau}_\ell$ and $\underline{\tau}_\ell$ are respectively the larger and the smaller of $\tau(a_\ell)$ and $\tau(a_{\ell+1})$. The weights sum to unity whenever the step is entirely contained in the maneuver, and to less than unity only on the step in which the maneuver ends, since a target altitude is in general not one the object reaches exactly at a step boundary: the missing share is the tail of the step over which the satellite was no longer in transit. A satellite launched part-way through a step produces no such deficit, because it is not counted at all until the following step boundary, by which point its transfer clock is already running and the whole of the next step falls inside the maneuver. What that costs instead is the fraction of the ascent performed between the launch epoch and the end of the step containing it, which is attributed to no node, and which is a discretization of the launch epoch onto the step grid rather than a loss of residence within a step.

The end of a rising maneuver is what moves an object from an R node to an A node. Writing a_q^* for the target semi-major axis of object q and $\tau_q^* = \tau(a_q^*)$ for the instant it is reached, both from eq. (4.37), the object deposits during step k if $\tau_k \le \tau_q^* < \tau_{k+1}$. By the second assumption above it deposits at its own inclination and RAAN, so the destination node is fully determined by its target orbit, and the number of satellites completing their ascent into node i over the step is

$$\Delta_{R_i}^{dep} = \sum_{q \in R} \mathbb{1}\left\{\tau_k \le \tau_q^* < \tau_{k+1}\right\} \delta_{\ell\, \ell(a_q^*)}\, \delta_{m\, m_q}\, \delta_{n\, n_q} \tag{4.39}$$

where $\ell(a)$ denotes the index of the shell containing the semi-major axis a. This is the term that appears with a positive sign in eq. (4.5) and with a negative one in eq. (4.7). Unlike the collision and failure terms it carries no randomness: given the launch schedule and the thrust level, the deposit epoch of every satellite is fixed at injection. The transition takes effect at the end of the step in which it occurs, so the satellite is counted as station-keeping in its target node from the following step onwards, which is also why its residence weights over the completing step sum to less than unity.

The reverse transition is triggered by time rather than by geometry. Every payload carries an end-of-operations epoch T_q^{EOL} fixed at launch, and the number of station-keeping satellites

in node i reaching it during step k is

$$\Delta_{A_i}^{EOL} = \sum_{q \in A} \mathbb{1}\left\{t_k \leq T_q^{EOL} < t_{k+1}\right\} \delta_{\ell\,\ell_q}\, \delta_{m\,m_q}\, \delta_{n\,n_q} \tag{4.40}$$

which is again deterministic, and which enters eq. (4.5) as the loss of retiring payloads. What happens next is not deterministic: of those satellites, $\Delta_{A_i}^{\gamma}$ fail their PMD maneuver, drawn as the jump process of eq. (4.4) and truncated at $\Delta_{A_i}^{EOL}$, since no more satellites can fail than retire. The failures become non-maneuverable at the orbit they occupied and enter eq. (4.8) through that same term. The remaining $\Delta_{A_i}^{EOL} - \Delta_{A_i}^{\gamma}$ satellites enter eq. (4.6) and begin a de-orbit spiral described by the same eqs. (4.35) to (4.38), with $\sigma = -1$, a_{inj} equal to the operational semi-major axis and a target at the re-entry altitude. A disposing satellite is therefore tracked all the way down and removed only on reaching that target, rather than being released to natural decay at some intermediate altitude, so the risk it poses to the shells it crosses on the way down is accounted for in full.

4.5. Break-up Model

When a collision occurs, the two objects involved in the collision can be partially or completely destroyed, producing a number of fragments. The break-up model describes the number of fragments produced, their size distribution, and their initial velocity distribution. While several break-up models have been proposed in the literature, the NASA SBM (Johnson et al., 2001) is the most widely used and accepted model for space debris generation. It is, for instance, the one used by all evolutionary models described in section 2.2 as well as NESSY. The SBM is based on empirical data from hypervelocity impact experiments and provides a statistical description of the fragment size distribution and velocity distribution. The model provides empirical formulas for both explosions and collisions, and separate coefficient sets for payloads and rocket bodies; only collisions between payloads are considered in this work, since every object in the population originates from a satellite. For the remainder of this section, the larger of the two objects involved will be called the *target*, of mass m_t, and the smaller the *projectile*, of mass m_p. Specifically, the number of fragments produced in a collision is given by

$$Z(\geq L_c) = 0.1 M_{tot}^{0.75} L_c^{-1.71} \tag{4.41}$$

where L_c is the characteristic length of the debris in meters, defined as the diameter of a sphere with the same volume as the debris, and M_{tot} is the combined mass of the target and projectile if the collision is catastrophic, or the mass of the projectile scaled by the collision velocity squared if the collision is non-catastrophic (Krisko, 2011)

$$M_{tot} = \begin{cases} m_t + m_p & \text{if EMR} \geq 40\,\text{J/g} \\ m_p v_{rel}^2 & \text{if EMR} < 40\,\text{J/g} \end{cases}$$

with v_{rel} in km/s in the second branch, where the expression is an equivalent mass rather than a physical one. EMR is the *energy-to-mass ratio* of the collision, defined as

$$\text{EMR} = \frac{m_p v_{rel}^2}{2 m_t}$$

that is, the kinetic energy of the projectile per unit mass of the target, conventionally expressed in J/g so that the threshold above takes its familiar value of 40 J/g. Since $m_p \leq m_t$ by construction, a collision between two intact objects is catastrophic for all but the most extreme mass ratios, whereas an impact by a small fragment generally is not.

Eq. (4.41) is cumulative, so the number of fragments generated by a single event is the count within the tracked size band

$$\zeta_{S_j S_\ell} = \left\lfloor Z(\geq L_c^{min}) - Z(\geq L_c^{max}) \right\rfloor \tag{4.42}$$

rounded down to an integer. The band used throughout this work is $L_c \in [10\,\mathrm{cm}, 1\,\mathrm{m}]$, the lower limit being the conventional threshold for trackable debris, and the individual characteristic lengths are drawn from the density $\propto L_c^{-2.71}$ that eq. (4.41) implies over that interval. The distribution function of the area-to-mass ratio then depends on the fragment size. For $L_c \geq 11\,\mathrm{cm}$, which covers all but the smallest of the fragments considered here, it is bimodal

$$f_{A/m}(\lambda_c, \chi) = w(\lambda_c)N(\mu_1, \sigma_1, \chi) + (1 - w(\lambda_c))N(\mu_2, \sigma_2, \chi) \tag{4.43}$$

where $\lambda_c = \log_{10}(L_c)$ and $\chi = \log_{10}(A/m)$ is the variable of the normal distribution

$$N(\mu, \sigma, \chi) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(\frac{-(\chi - \mu)^2}{2\sigma^2}\right)$$

so that A/m itself is log-normally distributed, and the parameters are given by

$$\mu_1 = \begin{cases} -0.6 & \lambda_c \leq -1.1 \\ -0.6 - 0.318(\lambda_c + 1.1) & -1.1 < \lambda_c < 0 \\ -0.95 & \lambda_c \geq 0 \end{cases} \qquad \mu_2 = \begin{cases} -1.2 & \lambda_c \leq -0.7 \\ -1.2 - 1.333(\lambda_c + 0.7) & -0.7 < \lambda_c < -0.1 \\ -2.0 & \lambda_c \geq -0.1 \end{cases}$$

$$\sigma_1 = \begin{cases} 0.1 & \lambda_c \leq -1.3 \\ 0.1 + 0.2(\lambda_c + 1.3) & -1.3 < \lambda_c < -0.3 \\ 0.3 & \lambda_c \geq -0.3 \end{cases} \qquad \sigma_2 = \begin{cases} 0.5 & \lambda_c \leq -0.5 \\ 0.5 - (\lambda_c + 0.5) & -0.5 < \lambda_c < -0.3 \\ 0.3 & \lambda_c \geq -0.3 \end{cases}$$

$$w = \begin{cases} 0 & \lambda_c \leq -1.95 \\ 0.3 + 0.4(\lambda_c + 1.2) & -1.95 < \lambda_c < 0.55 \\ 1 & \lambda_c \geq 0.55 \end{cases}$$

Below 8 cm the same distribution collapses to a single mode, of parameters

$$\mu_s = \begin{cases} -0.3 & \lambda_c \leq -1.75 \\ -0.3 - 1.4(\lambda_c + 1.75) & -1.75 < \lambda_c < -1.25 \\ -1.0 & \lambda_c \geq -1.25 \end{cases} \qquad \sigma_s = \begin{cases} 0.2 & \lambda_c \leq -3.5 \\ 0.2 + 0.1333(\lambda_c + 3.5) & \lambda_c > -3.5 \end{cases}$$

and between 8 cm and 11 cm the model bridges the two forms stochastically, choosing the single-mode branch with probability $10(\lambda_c + 1.05)$ and the bimodal one otherwise. The tracked band overlaps that bridge only in its lowest centimeter, so the vast majority of the fragments in this work are drawn from the bimodal form, but the bridge is retained rather than truncated so that the size distribution stays continuous across the tracking threshold.

The average cross-sectional area follows from the characteristic length alone,

$$A = 0.556945 L_c^{2.0047077}$$

for $L_c \geq 1.67\,\mathrm{mm}$, which every tracked fragment satisfies, and the mass is then recovered as $m = A/(A/m)$. Using A/m as the independent variable, the distribution function of the velocity kick imparted to the fragments is

$$f_{\Delta v}(\chi, \nu) = N(\mu, \sigma, \nu) \tag{4.44}$$

where $\mu = 0.9\chi + 2.9$, $\sigma = 0.4$ and $\nu = \log_{10}(\Delta v)$ is the variable of the normal distribution, so that the magnitude of the kick is log-normally distributed as well, with the larger area-to-mass ratios receiving the larger kicks. The kick is assumed to be isotropic, so the velocity increment of each fragment is obtained by sampling a random direction on the unit sphere and scaling it by Δv.

What remains is to place the cloud in space. The collision itself is assumed to occur at a point drawn uniformly in mean anomaly along the target's orbit, which for the near-circular orbits of interest is uniform in position as well, and every fragment inherits the position and velocity of the target at that point. The velocity increment is added to the inherited velocity, and the resulting state is converted back to orbital elements through the standard two-body

transformation (Vallado & McClain, 2013). Fragments whose post-kick orbit is hyperbolic, and therefore escapes, or whose perigee falls below the re-entry altitude, and therefore decays immediately, are discarded rather than added to the population. Neither loss is bookkept separately; both simply make the number of fragments injected fall short of eq. (4.41).

Each of the $\zeta_{S_j S_\ell}$ fragments is sized, kicked and placed as described above, and enters the network exactly as any other object does, through the occupancy law of eq. (4.34). Writing ℓ_q, m_q and n_q for the altitude, inclination and RAAN bins of the q-th fragment of the event, and $w_{q,\ell}$ for its residence weight from eq. (4.33),

$$\zeta^{F_i}_{S_j S_\ell} = \sum_{q=1}^{\zeta_{S_j S_\ell}} w_{q,\ell}\, \delta_{m\, m_q}\, \delta_{n\, n_q} \tag{4.45}$$

so $\zeta^{F_i}_{S_j S_\ell}$ is the share of one event's cloud that node F_i receives, while $\zeta_{S_j S_\ell}$ is the size of the cloud. Three properties follow. It is not in general an integer, since a fragment on an eccentric orbit is spread across the shells its orbit sweeps like any other object. It is a random variable even at fixed M_{tot}, because the collision point, the characteristic lengths, the area-to-mass ratios and the kick directions are all sampled, so two events with identical masses and relative velocity populate the network differently. And summing over all fragment nodes gives

$$\sum_{i=0}^{J} \zeta^{F_i}_{S_j S_\ell} \leq \zeta_{S_j S_\ell}$$

where the deficit is made up of the escaping and immediately re-entering fragments discarded above, together with any residence weight carried outside the altitude range the lattice covers. That inequality is the per-event form of the row deficit of the fragment redistribution table reported in chapter 7.

The SBM does not conserve mass by construction: the count, the size distribution and the area-to-mass distribution are three independent empirical fits, so the masses of the sampled fragments do not add up to M_{tot}, and nothing in the formulation forces them to. Several remediation schemes exist, from redistributing the residual over the sampled fragments to spawning or removing objects until the budget closes (Jang et al., 2024; Krisko, 2011), and the size band over which fragments are drawn governs how large the residual is, the upper limit far more than the lower one, since the count is dominated by the smallest fragments but the mass by the largest. Here the problem is ignored for simplicity: only fragments above 10 cm are tracked at all, and the quantities this work reports are collision rates and fragment counts rather than a mass budget.

5

Software Framework

This chapter describes PRISM, the simulation framework that implements the population evolution model of chapter 4 and the network analysis tools of chapter 7. It is organized around two core topics: the structure and operation of the simulator itself, and the treatment of rare events. Section 5.1 describes the overall design, the flow of a simulation step, and the modular interfaces. Section 5.2 addresses the numerical challenges that arise when collisions are rare, presenting the collision-sampling optimization and the conditioning scheme used to extract converged statistics from sparse events. Finally, section 5.3 then demonstrates that the implementation faithfully reproduces the theoretical predictions.

5.1. Software Architecture

Figure 5.1 shows PRISM's dynamic simulator, the component that advances the population of chapter 4 through the discrete time steps Δt_k introduced there. At every step, the simulator first samples the stochastic events (collisions and failed disposals) from the population as it stood at the beginning of the step, then applies the deterministic transport of that step (launches, orbit propagation, orbit-raising and de-orbiting maneuvers, re-entry) and finally advances the clock and refreshes the bookkeeping the next step's event rates will need. This ordering is deliberate and not arbitrary, since it is what makes each step a well-defined, self-contained Poisson process, as detailed in section 5.2.3.

Each stage of that loop sits behind its own narrow module and does not reach into the internals of the others. At present the orbital dynamics are the analytic drag and secular-drift model of chapter 4, the atmosphere is a tabulated exponential density profile, and the collision handling is the rate formulation of section 4.2, but none of the surrounding framework, the registry, the recording of results, the Monte Carlo and conditioning machinery of section 5.2, needs to know which particular choice is behind any one of these interfaces. A higher-fidelity propagator, a thermosphere model that resolves the solar cycle, or an entirely different collision formulation can each be substituted on its own, allowing for future extensions of the framework.

Within each stage, every operation is expressed once over the whole population it applies to, rather than object by object: propagating a species forward by one step, for instance, is a single array-level computation over every satellite of that species simultaneously, not as many sequential ones as there are satellites. This is possible because the population is held internally as a structure-of-arrays, described in section 5.1.2, and it is not a stylistic preference: measured against an otherwise identical calculation carried out one object at a time, the vectorized form is on the order of forty to sixty times faster over populations of 10^3 to 10^5 objects, and the same principle carries over to the density lookup and to the collision rates of section 5.2.2.

Because PRISM is stochastic, no single run of the simulator is representative on its own; a

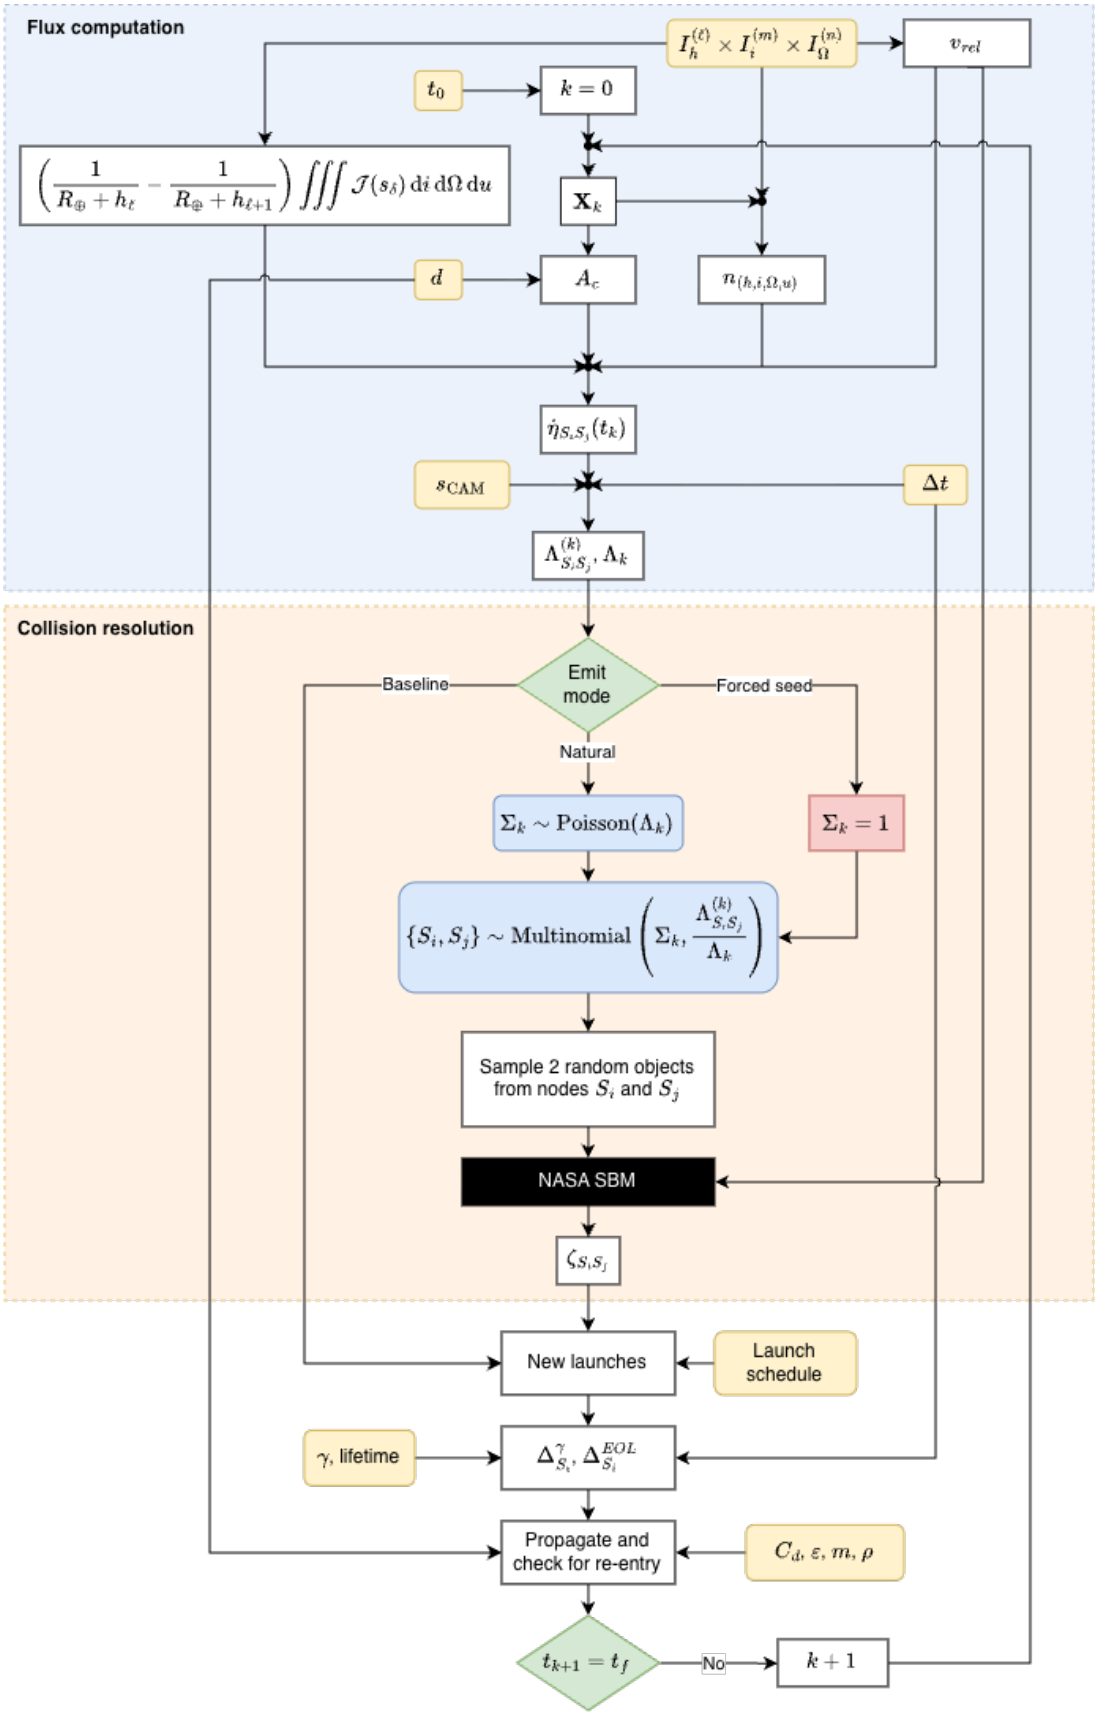

Figure 5.1: Flowchart diagram of the dynamic simulator.

meaningful description of a scenario's future requires repeating the run under many independent random seeds and combining the outcomes, as section 5.2.3 details. Seeds do not depend on one another at any point during their execution and only need to be brought together once every one of them has finished, so a batch of seeds maps directly onto a pool of worker processes, each running its own independent copy of the simulator to completion. Since a single cluster node exposes many processing cores, this lets tens of seeds advance at once rather than one after another, so that the wall-clock time of a Monte Carlo campaign scales with the number of cores made available to it rather than with the number of seeds required.

5.1.1. Initial Population and Data Sources

Regardless of where it comes from, every object handed to the simulator is described by the same schema: its specie, as defined in section 4.1; its orbital elements, semi-major axis, eccentricity, inclination, RAAN, and argument of perigee, at a reference epoch; the physical properties, mass, cross-sectional area, and drag coefficient, needed to evaluate the dynamics and the collision rate; and the epoch at which it enters the simulation. Rising and de-orbiting satellites additionally carry the injection and target orbit of their transfer. This shared representation is what allows every later stage of the software to remain indifferent to which of the following two sources produced the population it is handed.

The first source builds an idealized constellation directly from the conventional Walker-pattern i:T/P/F, provided for each shell: altitude, inclination, total number of satellites, and how they are divided into orbital planes and phased within each. Satellites are placed exactly, evenly spaced within their pattern, with no dispersion in their physical properties, and all already on-station, with no build-up history to speak of. This is the appropriate choice whenever a study wants any scatter in the results to be attributable to a mechanism under the modeler's control, rather than to the incidental details of how a real constellation happened to be deployed.

The second source builds a population from a real constellation's records in the Database and Information System Characterising Objects in Space (DISCOS) database, maintained by ESA. For a named constellation, the catalog entries and their recorded injection and target orbital elements and physical properties are retrieved and converted into the shared schema above, giving a heterogeneous population that reflects an actual deployment, dispersion in orbital elements, mass, and cross-section. Because the catalog only records objects launched up to the day it was queried, this route additionally needs a model of the launch traffic still to come in order to extend the manifest across the whole simulated horizon. For each shell, the historical launch cadence and the cumulative number of deployed satellites are fit, once, to one of a small number of trajectory shapes: a linear ramp toward a target fleet size followed by a replenishment phase, a linear indefinite continuation at the historical deployment rate, a deployment already considered complete that only replenishes retirements, or a closed loop that replaces every retiring satellite one for one, holding the fleet at a constant size. That fit is carried out offline, outside the simulator, on the historical record alone; only the resulting target fleet size, deployment rate, number of satellites per launch and injection altitude, are read in when the population is built, and the simulator itself never re-derives them from the raw catalog. Extending the retrieved catalog with the launches this model predicts produces a manifest indistinguishable, in form, from one where those future launches had already taken place, each carried at its expected epoch. Swapping in a different assumption about future traffic, or an altogether different data source, is consequently a matter of producing a manifest in the shared schema above, without any of the rest of the framework needing to change.

5.1.2. Scenario Creation

The manifest produced by either source above is static: it is the complete catalog of every object that will ever appear over the course of a run, each tagged with its own launch epoch, but it is not by itself the population at any particular simulated time, and it is never modified once built. The state the simulator actually advances step by step is kept separately, in a mutable registry in which every attribute of the schema above (orbital elements, physical properties,

specie) is held as its own array, with a common index running across all of them identifying a given object consistently in every array. This layout, one array per attribute rather than one record per object, is what makes the vectorized, whole-population operations described above possible: acting on one attribute of every currently active object reduces to a single array-level operation rather than one touching each object's own record in turn.

At the start of a run, this mutable state is seeded with every object whose launch epoch in the static manifest precedes the simulation's start date. As the simulated clock advances, whenever it passes the epoch of the next scheduled launch, every object due to enter at that epoch is appended to the mutable state in a single batch operation and is then no longer considered for launch, so that each object enters the simulation exactly once, at the epoch the manifest assigns it. Objects that leave, whether removed by a collision, natural re-entry, or a failed disposal, are dropped from the registry the same way, one batch operation removing every object due to leave at that step rather than one operation per object. Finally, the whole of this registry can be captured and later restored to a run at an earlier point in its evolution, a capability section 5.2.3 relies on to rewind a collision-free run to a chosen step and force a collision there.

5.2. Stochastic Simulation

This section turns the model of chapter 4 into the sequence of operations the simulator actually carries out at every step, and describes the two optimizations that make it tractable at the scales of interest: an aggregate way of drawing collisions that avoids ever enumerating every pair of nodes when it is not needed, and a conditioning scheme that avoids the need for an intractably large ensemble when collisions are rare.

5.2.1. Step Ordering

Section 4.1 treats the population as frozen at the start of a step for the purpose of evaluating event rates. In practice, this requires carrying out every step in three phases. First, the stochastic events (collisions and failed disposals) are drawn from rates evaluated on the population exactly as it stood at the start of the step; nothing produced later in that same step, a new launch, a fragment, an object that has just moved into another node, feeds back into this draw. Second, the deterministic transport of the step is applied: scheduled launches enter, satellites reaching the end of their life or failing an attempted disposal change specie, drag and the low-thrust maneuvers carry every object along its trajectory, and objects that have re-entered or completed a maneuver leave or change specie accordingly. Third, the clock advances and the population's node occupancy is recomputed exactly once, from the state the first two phases produced, ready for the following step's first phase.

This ordering is what makes the per-step Poisson treatment of section 4.1 homogeneous over the whole step rather than only approximately so: the rate governing a step is a function of one well-defined population snapshot, not of a population that is itself changing while that same rate is being evaluated. It also means that an object still part-way through a maneuver at the start of a step needs no special handling relative to collisions, since whatever fraction of the step it spends in a given node is already reflected in the occupancy carried over from the previous step's close; the transport phase later in the same step is what advances it toward its next position, to be picked up in turn by the following step's rate evaluation.

Freezing the population for the length of a step is an approximation, and its error is governed by the step itself: the shorter the step, the less the true population departs from the snapshot the rates were evaluated on. That dependence is measured directly in section 6.3, which quantifies it through the Courant number of eq. (6.18), the number of altitude bins an object crosses in one step, and identifies the active-derelict channel as the one exposed to it. The result of that study is that the step is a second-order parameter for the scenarios studied here and carries the smallest error of any of the discretization choices, so that the approximation is acceptable for the purposes of this work.

5.2.2. Collision Sampling

The collision rate of section 4.2 has, in principle, to be evaluated for every pair of populated nodes at every step, which is a cost that grows with the square of the number of populated nodes. As explained in section 5.2.1, the population is frozen for the duration of a step, so the stochastic processes governing it can be treated as Poisson point processes, homogeneous over that step. Therefore, the number of collision events occurring between a pair of nodes over one time step follows a Poisson distribution

$$\Sigma^{(k)}_{S_iS_j} \sim \text{Poisson}\left(\Lambda^{(k)}_{S_iS_j}\right)$$

where $\Lambda^{(k)}_{S_iS_j} = \alpha_{\text{CAM}}\,\dot{\eta}_{S_iS_j}\Delta t_k$ is the per-step per-pair expected number of collisions. Because the collision events of different pairs of nodes are independent from each other (within the step), the total number of collisions across all node pairs within the same step is also Poisson-distributed

$$\Sigma_k = \sum_{S_iS_j} \Sigma^{(k)}_{S_iS_j} \sim \text{Poisson}(\Lambda_k) \qquad \text{with} \quad \Lambda_k = \sum_{S_iS_j} \Lambda^{(k)}_{S_iS_j}$$

Conditional on Σ_k, the per-pair number of collisions can be recovered from a multinomial distribution

$$\Sigma^{(k)}_{S_iS_j} \mid \Sigma_k \sim \text{Multinomial}\left(\Sigma_k, \frac{\Lambda^{(k)}_{S_iS_j}}{\Lambda_k}\right)$$

Since the number of node pairs is very large, drawing the number of collisions for each pair individually is extremely inefficient, especially since the overwhelming majority of steps produce no collision at all, so almost every one of those draws would yield zero. Aggregating the total number of collisions first and only then decomposing it into the individual pairs reduces the process to one Poisson draw on Λ_k, followed by Σ_k (which is small, typically zero) categorical draws over the pair index. What must therefore be known at every step is only the scalar Λ_k; the full set of per-pair rates $\Lambda^{(k)}_{S_iS_j}$ that the categorical draw needs is only ever assembled on the rare step where Σ_k turns out to be nonzero, since it is only then that they are needed at all.

Obtaining Λ_k without enumerating every pair rests on the same symmetry already exploited in chapter 4 to build the geometric overlap look-up table, eq. (4.26): the rate between a pair of nodes depends on their two RAAN bins n_1, n_2 only through the signed difference $\Delta n = (n_2 - n_1) \bmod N_\Omega$, not on the two bins individually. There, that property was used once, offline, to build the overlap tensor $\Phi_{m_1m_2\Delta n}$ itself at a fraction of the cost. Here it is used again, this time at run time, to fold every RAAN pair into the aggregate rate at once. For a fixed pair of inclination bins, $\dot{\eta}_{S_iS_j}$ is bilinear in quantities carried by each node along the RAAN axis, namely the object counts and diameter moments of section 4.2, so summing it over every pair of RAAN bins takes the form

$$\sum_{n_1,n_2} \Phi_{m_1m_2[(n_2-n_1) \bmod N_\Omega]}\, x[n_1]\, y[n_2] = \sum_{\Delta n} \Phi_{m_1m_2\Delta n}\, (x \star y)[\Delta n]$$

where x and y are the relevant node quantities viewed as length-N_Ω vectors indexed by RAAN bin, and $(x \star y)[\Delta n] = \sum_n x[n]\, y[(n+\Delta n) \bmod N_\Omega]$ is their circular cross-correlation. By the correlation theorem, a circular cross-correlation is the transform of a pointwise product of discrete Fourier transforms, so the sum over Δn is obtained once per inclination pair rather than once per RAAN pair. Appendix C.2 reviews the transform and the correlation identity this rests on.

Evaluating every pair of populated nodes explicitly costs $O(P^2)$, with P the number of populated nodes; obtaining Λ_k through the transform instead costs $O(N_i^2 N_\Omega \log N_\Omega)$, trading a term quadratic in the node count, which grows with the size of the population, for one that is quadratic only in the (fixed, small) number of inclination bins and log-linear in the number of RAAN bins. Measured against enumerating every populated pair, the transform-based aggregate is on the order of a hundred times faster, with the exact factor depending on how densely each node is populated. A more sparsely populated scenario enumerates fewer pairs to begin with, so the two costs move closer together, but the asymptotic argument guarantees the transform-based path is never the more expensive of the two.

Once a pair of nodes has been selected for a collision event, the two colliding objects are drawn from the objects occupying those nodes uniformly. This is a simplification relative to the approach of NESSY (Wang et al., 2025), which instead draws from a distribution favoring larger objects, on the grounds that a larger cross-section makes a collision more likely in the first place; the uniform choice made here forgoes that refinement in exchange for a simpler and less costly draw. How often a collision happens is unaffected by the choice: the cross-section of eq. (4.31) is built from the first and second moments of the diameters held in each node, so a node of larger objects already produces proportionally more events. What the uniform draw changes is only which two objects are named once an event has been drawn, and with them the mass that enters the breakup model. For the intact species that mass is effectively fixed, since a constellation flies one satellite design, and a size-weighted draw would select much the same mass the uniform one does. The fragment nodes are the exception, their size distribution being broad by construction, and there the uniform draw takes a lighter pair on average than a size-weighted draw would. Fragment production from fragment-fragment events is therefore understated.

5.2.3. Conditioning

Let X be any observable (node count, total population, fragment count, etc.) and N be the total number of collisions in a run. If the number of objects in the simulation is low, then the expected number of collisions over the whole run is $\Lambda \ll 1$ and a plain Monte Carlo would require tens or hundreds of thousands of runs to be able to witness even just one collision. Attempting to extract meaningful statistics by brute force is therefore computationally intractable.

A standard remedy for rare-event simulation, importance sampling, was considered and discarded. Importance sampling simulates under a different, tilted probability measure that makes the rare event common, then corrects for the change of measure so the estimator stays unbiased: running the collision process under an inflated intensity $w(t)\lambda(t)$ and re-weighting each realization by the likelihood ratio between the true and the tilted intensity recovers an unbiased estimate of $\mathbb{E}[X]$. This works whenever X is a scalar functional computed after the fact from an otherwise faithful trajectory, so that only the final average needs correcting. It fails here because a collision does not merely register as an outcome to be weighted afterwards: it also injects a cloud of fragments into the population, which then propagate and can collide again for the remainder of the run. Under an inflated rate, those fragments are the byproduct of a collision that was not physically real, yet they feed back into the rest of the same trajectory, through further propagation and further collisions, rather than only into a scalar output tallied at the end. No re-weighting applied after the fact can undo having contaminated the sample path itself with objects that should never have existed at that rate.

One solution that still captures the ripple effect of collisions without altering any rate is to force a collision to occur, and then simulate the evolution of the system conditioned on that collision having occurred. According to the law of total expectation, the expected value of X can be obtained as the sum of the expected value of X conditioned on no collisions occurring, and the expected value of X conditioned on at least one collision occurring, weighted by their corresponding probabilities

$$\begin{aligned}\mathbb{E}[X] &= P(\Sigma = 0)\mathbb{E}[X \mid \Sigma = 0] + P(\Sigma \geq 1)\mathbb{E}[X \mid \Sigma \geq 1] = \\ &= e^{-\Lambda}\mathbb{E}[X \mid \Sigma = 0] + (1 - e^{-\Lambda})\mathbb{E}[X \mid \Sigma \geq 1]\end{aligned} \tag{5.1}$$

The collision process up to the first event is an inhomogeneous Poisson process with time-varying intensity $\lambda(t)$, so the expected number of collisions over the full simulation is

$$\Lambda = \int_0^T \lambda(t)\,\mathrm{d}t \qquad \text{with} \quad \lambda(t) = \sum_{S_i S_j} \alpha\,\dot{\eta}_{S_i S_j}(t) \tag{5.2}$$

The Probability Density Function (PDF) of the first collision to occur exactly at time t conditioned

on observing at least one collision over the full run is

$$f_{T_1|\Sigma(t)\geq 1}(t) = \frac{f_{T_1}(t)}{P(\Sigma \geq 1)} = \frac{\lambda(t)e^{-\Lambda(t)}}{1 - e^{-\Lambda}}$$

where $\Sigma(t)$ is the cumulative number of collisions at time t and $\Lambda(t)$ is the integral of the intensity from 0 to t. Therefore, the expected value of X conditioned on a collision to occur can be estimated by forcing the first collision at a time drawn from the first-arrival PDF, and simulating the evolution of the system from that point onward via plain Monte Carlo. Finally, the unconditional estimate can be recovered from eq. (5.1).

In practice, because the simulation is performed in discrete time steps $[t_k, t_{k+1})$ and collisions are modeled by a step-wise homogeneous Poisson process, the step-wise constant intensity is

$$\lambda_k = \lambda(t_k) = \sum_{S_i S_j} \alpha \, \dot{\eta}_{S_i S_j}(t_k)$$

and the expected number of collisions over each step is $\Lambda_k = \Lambda(t_{k+1}) - \Lambda(t_k) = \lambda_k \Delta t_k$ where

$$\Lambda(t_k) = \sum_{j<k} \int_{t_j}^{t_{j+1}} \lambda_j \, \mathrm{d}t = \sum_{j<k} \lambda_j \Delta t_j = \sum_{j<k} \Lambda_j$$

Now, the probability that the first collision occurs exactly during step k is

$$P(t_k \leq T_1 \leq t_{k+1}) = P(T_1 \leq t_{k+1}) - P(T_1 \leq t_k) = e^{-\Lambda(t_k)} - e^{-\Lambda(t_{k+1})}$$

These are mutually exclusive events, i.e., if the first collision occurs at step k it cannot occur at any other step, therefore summing them gives the total probability of the first collision to occur any time between 0 and T

$$P(\Sigma \geq 1) = \sum_k P(t_k \leq T_1 \leq t_{k+1}) = 1 - e^{-\Lambda}$$

Therefore, the seed step s at which a collision is triggered for each one of the Monte Carlo runs is sampled from a categorical distribution over all steps

$$s \sim \text{Categorical}\left(\frac{P(t_k \leq T_1 \leq t_{k+1})}{1 - e^{-\Lambda}}\right)$$

Then, a pair of nodes is drawn from a categorical distribution over all pair indices

$$\{S_i, S_j\} \sim \text{Categorical}\left(\frac{\Lambda^{(s)}_{S_i S_j}}{\Lambda_s}\right)$$

Finally, the specific pair of objects involved in the collision is sampled as usual. Each realization shares the baseline up to step s and forks afterwards: a realization seeded at $s > k$ still equals the baseline at k, hence

$$\hat{\mathbb{E}}[X_k \mid \Sigma \geq 1] = \frac{1}{M}\left(\sum_{i=1}^{M_a(t_k)} X^{(i)}(t_k) + X_b(t_k)\,(M - M_a(t_k))\right)$$

where M is the number of forced-seed realizations, of which $M_a(t_k)$ have already forked from the baseline by step k, $X^{(i)}(t_k)$ is the value observed in the i-th such realization, and $X_b(t_k)$ is the baseline value. Finally,

$$\hat{\mathbb{E}}[X_k] = e^{-\Lambda} X_b(t_k) + (1 - e^{-\Lambda})\hat{\mathbb{E}}[X_k \mid \Sigma \geq 1]$$

5.3. Verification

PRISM is not validated against historical observations of a real constellation. Doing so convincingly would demand considerable effort assembling and cleaning a historical record, and even then a real operator's decisions are not something the model has any way of anticipating: if an operator lowers its constellation in response to a period of high solar activity, or delays a disposal campaign for operational reasons, the model, driven only by its own drag and disposal assumptions, has no mechanism to reproduce that choice, and the resulting mismatch would say nothing about whether the underlying physics is right. Even a close match would prove only that the model can be tuned to reproduce one particular deployment, not that it is correct in general or that it extrapolates reliably to a constellation, or a future, it has not seen. Instead, following the approach generally adopted in the literature for evolutionary models of this kind, the model is verified by comparing its output against an independently developed one, MOCAT-MC (Jang et al., 2024). MOCAT-MC is chosen because it is open-source[1], has itself been extensively cross-checked against other models, and is also a Monte Carlo model. MOCAT-MC also propagates its population with the same analytical secular propagator PRISM uses (Martinusi et al., 2015), so the comparison does not exercise the collision computation alone: agreement between the two also confirms that this model's own implementation of that propagator is correct.

Three considerations make the full LEO environment, rather than a constellation scenario, the right benchmark. The first is that MOCAT-MC detects collisions with the CUBE method (J. C. Liou et al., 2003), a spatial-density scheme that discretizes the cartesian space into cubes and counts co-located objects at each time step (see section 2.3.2), sharing essentially nothing with the RAAN-resolved rate construction of section 5.2.2. The two models reach their collision counts by structurally unrelated routes, which is what makes MOCAT-MC an independent reference rather than a second implementation of the same idea. That independence is only worth having where CUBE is itself reliable, and Lewis et al. (2019) showed that it is not for constellation geometries, where it predicts collisions between satellites on non-intersecting orbits unless the cube size is driven far below what a tractable run allows, as discussed in chapter 1. Benchmarking on a constellation would consequently be a weaker test rather than a sharper one: the reference would be the less trustworthy of the two, and any disagreement would be uninterpretable, falling in exactly the regime the two models are designed to differ in. The full environment is where CUBE is at its most defensible, and therefore where a disagreement would genuinely implicate PRISM. The second consideration is that matching one of MOCAT-MC's own published scenarios, specifically the one Wang et al. (2025) used to verify NESSY, anchors the comparison to a third party: agreement here also cross-checks that this MOCAT-MC run reproduces the published reference, rather than resting on a configuration set up independently and possibly incorrectly. The third, which will resurface in chapter 6, is that a real, heterogeneous population is markedly more robust to the stochastic noise inherent in a Monte Carlo comparison than an idealized constellation would be, so that whatever differences survive can be attributed to the modeling choices that separate the two models rather than to seed noise.

Concretely, the scenario is an ensemble of 200 seeds evolving, without new launches, the real LEO population as recorded in DISCOS and cross-matched against TLE data at the epoch of 1st January 2023, over a 100 year horizon. Table 5.1 summarizes the parameters shared by the two runs. DISCOS records of every object's mass and cross-sectional area are incomplete, especially for debris, and MOCAT-MC fills the gaps by resampling them, per seed, from a Gaussian-mixture model fit to the population with known attributes; the resulting per-seed population files are read directly into the same population-source abstraction described in section 5.1.2, so that PRISM's own ensemble reproduces, seed for seed, exactly the physical attributes MOCAT-MC itself drew. The species model of section 4.1 has no counterpart for rocket bodies or other mission-related hardware that is neither a payload nor collision or explosion debris, and such objects are consequently removed from the population handed to this model. What remains after that exclusion is 17 715 objects, 5471 of them active, 2440 non-maneuverable, and 9804 fragments, identical across all seeds since only the physical attributes, not the population membership, are

[1] Repository available at https://github.com/ARCLab-MIT/MOCAT-MC.

resampled.

Table 5.1: Configuration of the comparison run.

Parameter	MOCAT-MC	This work
Atmosphere model	JB2008 (dynamic)	layered exponential (static)
Time step	5 d	30 d
Time horizon	100 year	100 year
Number of seeds	200	200
Launches	none	none
PMD failure rate	5 %	5 %
CAM success rate	99.99 %	99.99 %
Disposal / injection	-	instantaneous
Catastrophic threshold	40 J/g	40 J/g
Altitude range	200 km to 2000 km	200 km to 2000 km
Population epoch	1st January 2023	1st January 2023

One structural difference between the two runs could not be eliminated: MOCAT-MC's atmosphere is the dynamic, solar-flux-driven JB2008 model, against the static, layered-exponential one PRISM is restricted to, a limitation already noted in chapter 3. The maneuver model of section 5.1.2, by contrast, is simply disabled for this comparison, so that disposal and injection are treated as instantaneous on both sides, matching MOCAT-MC's own orbit-control convention.

Figure 5.2 shows the two ensembles agree closely on both the cumulative number of collisions and the population by specie, with the shaded band in each case the mean plus or minus one standard deviation across the seed ensemble. The one systematic qualitative difference is the absence, in PRISM, of the roughly month-to-decade oscillations visible in the MOCAT-MC curves, which track the solar cycle through JB2008's time-varying density and have no counterpart in a static atmosphere. PRISM also runs slightly below MOCAT-MC in the cumulative collision count, and the most likely contributor is that the two implementations of the SBM treat the mass budget of a breakup differently. Both draw a fragment count from the same size distribution and assign each fragment a mass through its sampled area-to-mass ratio, but MOCAT-MC then closes the budget explicitly: it compares the total mass of the drawn fragments against the mass of the parents and, if the draw falls short, samples a small number of additional remnant objects carrying exactly the missing mass, or, if it overshoots, removes the largest fragments until the two match. This model performs no such reconciliation. It takes the fragments as drawn, so the mass lying below the lower end of the modeled size range is simply never represented as objects. Since the size distribution is truncated from below, the draw is expected to fall short more often than it overshoots, and the remnant objects MOCAT-MC adds back to absorb that shortfall are few and correspondingly massive, hence large in cross-section and disproportionately likely to be involved in a further collision. The direction of the effect is therefore consistent with the small deficit observed, though the magnitude has not been quantified and the attribution remains unconfirmed. Despite these differences, the agreement is a meaningful demonstration that the implementation is correct. Two models that share their propagator and the breakup model underlying their fragmentation, but nothing of how collisions are detected in the first place, track each other to well within the spread of their own ensembles over a century of evolution, which is a considerably stronger statement than agreement between two implementations of the same algorithm would be.

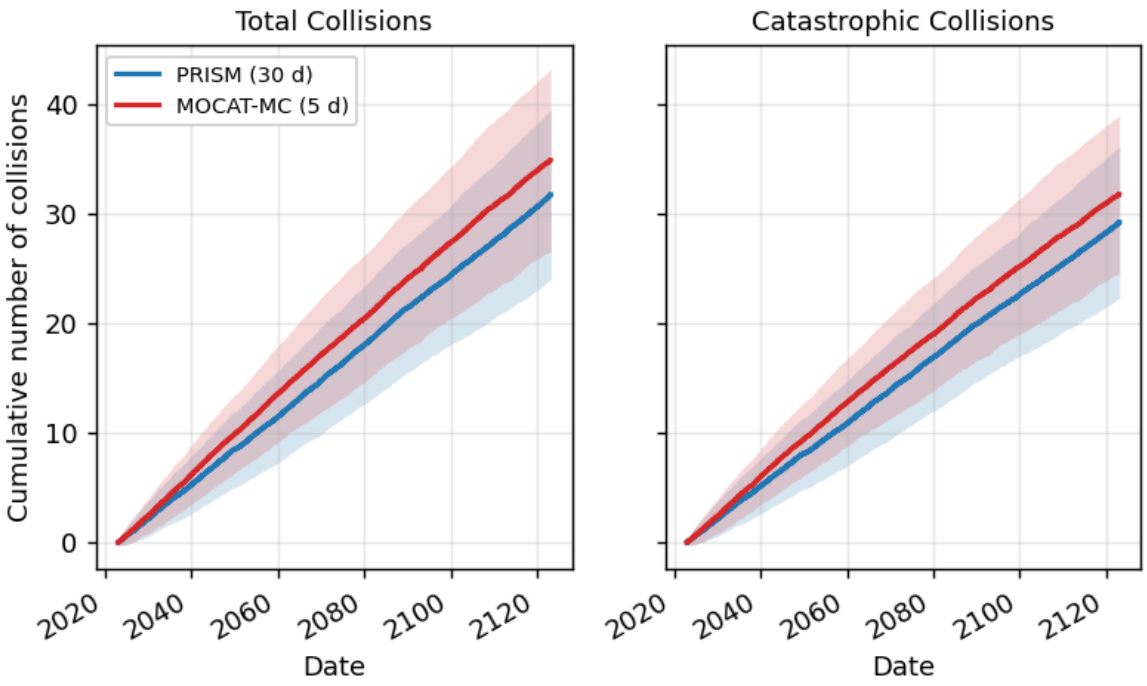

(a) Cumulative number of collisions (left) and catastrophic collisions (right).

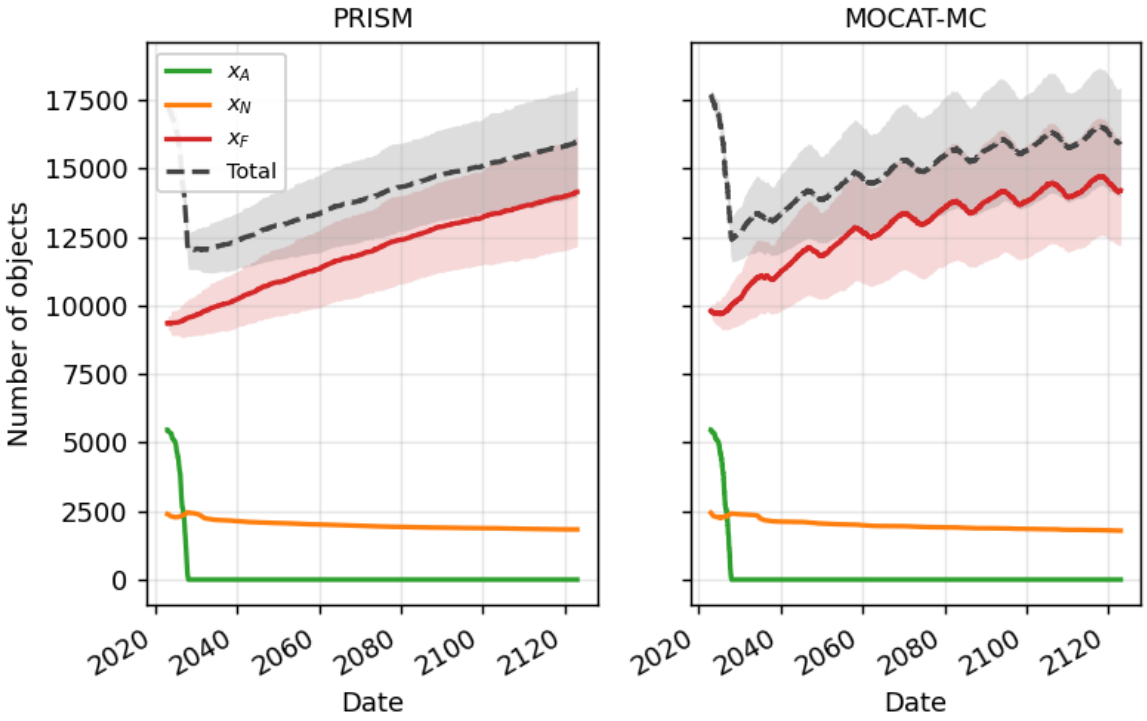

(b) Evolution of the population by specie.

Figure 5.2: Comparison with MOCAT-MC on the full population as of January 1st 2023 averaged across 200 runs.

6

Sensitivity Analysis and Robustness

This chapter quantifies how far PRISM's collision-rate estimates, derived in chapter 4, depend on choices that have no counterpart in the physical system: the widths Δh, Δi and $\Delta\Omega$ of the lattice bins, and the time step Δt that advances the simulation. Section 6.1 screens all four jointly, with a variance decomposition, to establish which parameter matters and in what order. Sections 6.2 and 6.3 then take up the dominant ones in turn, tracing each result back to its cause.

Consider a lattice composed of equally-sized bins in each dimension, of size Δh, Δi, and $\Delta\Omega$ respectively. Section 4.2 assumed that, for a sufficiently refined lattice, the phase-space density of each node and the relative velocity of a node pair can be treated as constant over one bin and pulled out of the collision-rate integral, isolating the geometric overlap factor from the objects' distribution. In other words

$$n_{\boldsymbol{u}} = \frac{x_{S_i}}{2\pi\Delta h\Delta i\Delta\Omega}$$

Under this assumption, the raw collision rate between two nodes scales as

$$\dot{\eta}_{S_iS_j} \propto \frac{x_{S_i}x_{S_j}}{(\Delta h\Delta i\Delta\Omega)^2}\left(\frac{1}{R_\oplus + h_\ell} - \frac{1}{R_\oplus + h_{\ell+1}}\right)\iiint \mathcal{J}(s_\delta)\,\mathrm{d}i\,\mathrm{d}\Omega\,\mathrm{d}u \tag{6.1}$$

The rest of this chapter tests that assumption, and the lattice built on it, against five populations: four idealized Walker-constellation archetypes, and the real, heterogeneous LEO catalog of 17 715 objects already used in chapter 5 to verify PRISM against MOCAT-MC . The four archetypes are surrogates loosely based on Starlink, OneWeb and Iridium-NEXT, chosen to span the parameter space relevant to the discretization problem: a large single-shell constellation, a multi-shell constellation whose shells sit only a few kilometers apart in altitude, a small polar constellation, and a high-altitude polar constellation. Table 6.1 summarizes the orbital and physical parameters of each; the catalog carries no equivalent row, since it is not a constellation but a snapshot of the whole tracked LEO population.

Table 6.1: Parameters of the idealized Walker-constellation archetypes used in this chapter.

Archetype	Altitude	i: T/P/F	Pattern	Mass	Area	EOL
Large single-shell	550 km	53°: 1584/72/1	Delta	260 kg	20 m^2	5 yr
Multi-shell	540/550/560 km	53°:1440/48/1	Delta	800 kg	40 m^2	5 yr
	570 km	43°:1440/48/1				
Small polar	780 km	86.4°: 66/6/1	Star	689 kg	5 m^2	5 yr
High-altitude polar	1200 km	87.9°: 588/12/1	Star	148 kg	5 m^2	5 yr

6.1. Global Sensitivity Analysis

This chapter attributes the spread of an estimate to its inputs with a variance decomposition in the sense of Sobol. The idea is to treat each discretization parameter X_i as a random variable, uniform over the set of values a practitioner might plausibly choose, so that the response Y becomes a random variable with a finite variance of its own. That variance is then partitioned into additive contributions: a first-order term for each parameter acting alone, plus higher-order terms for every combination of parameters acting jointly. Writing $V = \mathrm{Var}(Y)$ and $V_i = \mathrm{Var}(\mathbb{E}[Y|X_i])$, the first-order index of a parameter is

$$S_i = \frac{V_i}{V} \tag{6.2}$$

which is the fraction of the spread that would disappear if that one parameter were pinned and all the others left free. The second-order terms follow from the same conditional variances,

$$S_{ij} = \frac{\mathrm{Var}(\mathbb{E}[Y|X_i, X_j]) - V_i - V_j}{V} \tag{6.3}$$

that is, whatever the pair explains beyond what its two members explain separately, and higher orders are built the same way, each one net of every lower-order term it contains. The decomposition is exhaustive,

$$\sum_i S_i + \sum_{i<j} S_{ij} + \ldots = 1$$

so the orders above the first can always be read as the part of the spread that no single parameter can be held responsible for. In the factorial used here the third-order term is simply the residual left once the main effects and the pairs are subtracted.

The results that follow rest on two different response variables. The first is the instantaneous collision rate $\dot{\eta}$ evaluated on a frozen population: the objects are placed once, the lattice is built around them, and the rate follows in closed form from the overlap integrals, with no propagation or time stepping involved. The second is the expected number of collisions $\mathbb{E}[\Sigma]$ accumulated over a ten-year horizon, obtained by propagating the same population forward and summing the collision rate over each step; no individual collision is drawn, so the result is the expectation of the underlying Poisson process rather than one simulated realization of it. The two answer different questions. The static rate isolates the lattice dependence of the estimator itself, since with the population pinned only the bin widths can move the answer, and at a fraction of a second per cell it is cheap enough to afford a complete factorial over all three spatial bin widths at once. The time-integrated $\mathbb{E}[\Sigma]$ is the only response that depends on the time step Δt, since Δt changes how the population evolves rather than how a fixed population is binned; at minutes per cell it is also expensive enough that its factorials hold Δh fixed per archetype. The two factorials therefore sweep different cubes, $(\Delta h, \Delta i, \Delta\Omega)$ for the static rate and $(\Delta i, \Delta\Omega, \Delta t)$ for $\mathbb{E}[\Sigma]$ with Δh pinned; only the two angular widths are swept in both.

Independently of that split, either response can be reported with or without collision avoidance applied, and both variants appear below. The *raw* rate counts every pair the lattice admits; the *CAM-realized* rate instead weights each pair by the avoidance factor α_{CAM} of eq. (4.3), which suppresses any channel an active satellite takes part in and leaves what survives dominated by derelicts and fragments. These are the same response with one physical effect switched on, not two more responses, and they are reported side by side because the raw rate exposes the estimator's discretization behaviour undiluted while the realized rate is the one an operator would actually experience. Nothing guarantees the two depend on the lattice the same way, and fig. 6.1 shows that they do not.

For any of this to be meaningful, nothing but the lattice may differ between two cells of a factorial, so the two places that could otherwise introduce randomness are controlled directly. The first is satellite failure: each shell sheds derelicts on a fixed schedule, with each arrival inheriting the elements of a specific, index-stepped orbital slot, so failure times and locations are set once and never redrawn. The second is the fragment cloud, present only in compositions that carry one: writing A for active satellites, N for faulty ones, and F for the fragments a breakup

leaves behind, $A + N$ has no cloud while $A + N + F$ adds one, generated by forcing a single breakup inside the archetype's own population so the colliding pair, and with it the altitude and inclination at which the cloud is born, belong to that constellation. This breakup is the one genuinely stochastic ingredient, since PRISM samples the number, mass and velocity of the fragments it produces, and two safeguards keep that single draw from leaking into the factorial: it is generated once, under a fixed seed, on an internal lattice and time step decoupled from the values being swept, so the same draw always produces the same cloud; and the resulting population is saved to disk once per case, with every lattice cell in that factorial reloading the same file rather than regenerating the breakup. Any difference between two cells can therefore only come from the bin widths, or the time step where one is involved, so every variance share below is attributable to the discretization alone; with no run-to-run scatter to average away, one evaluation per cell is not a sample but the answer. The results remain conditional on the one breakup realization each fragment-rich case carries, so what is measured is the sensitivity of the estimate to the lattice, not to which fragmentation happened to occur.

A second, independent source of approximation error would normally come from the Sobol indices themselves, ordinarily estimated by Monte Carlo sampling of the input space with an estimator error of its own. That step is unnecessary here: each parameter takes only a handful of discrete levels and the response is a deterministic function of them, so the complete factorial is evaluated instead, every combination of levels exactly once, and the conditional expectations in eqs. (6.2) and (6.3) reduce to exact averages over the evaluated cells; the indices reported in this chapter are therefore computed outright, with no estimation step involved. Two caveats govern how those shares should be read. A Sobol share depends on the specific values swept for each parameter in this factorial, for instance the 5 km to 100 km altitude ladder or the four inclination settings used throughout, not on some intrinsic importance of the parameter itself: test a wider spread of Δh and its share rises for that reason alone, with nothing about the underlying model having changed. A share is also relative to the total spread, not to the outcome itself: a parameter can own nearly all of a response's variance while that response barely moves at all, in which case the share is large but the effect it describes is not. Every variance figure in this chapter therefore also reports the total swing of the response across its factorial, so the two can be read together.

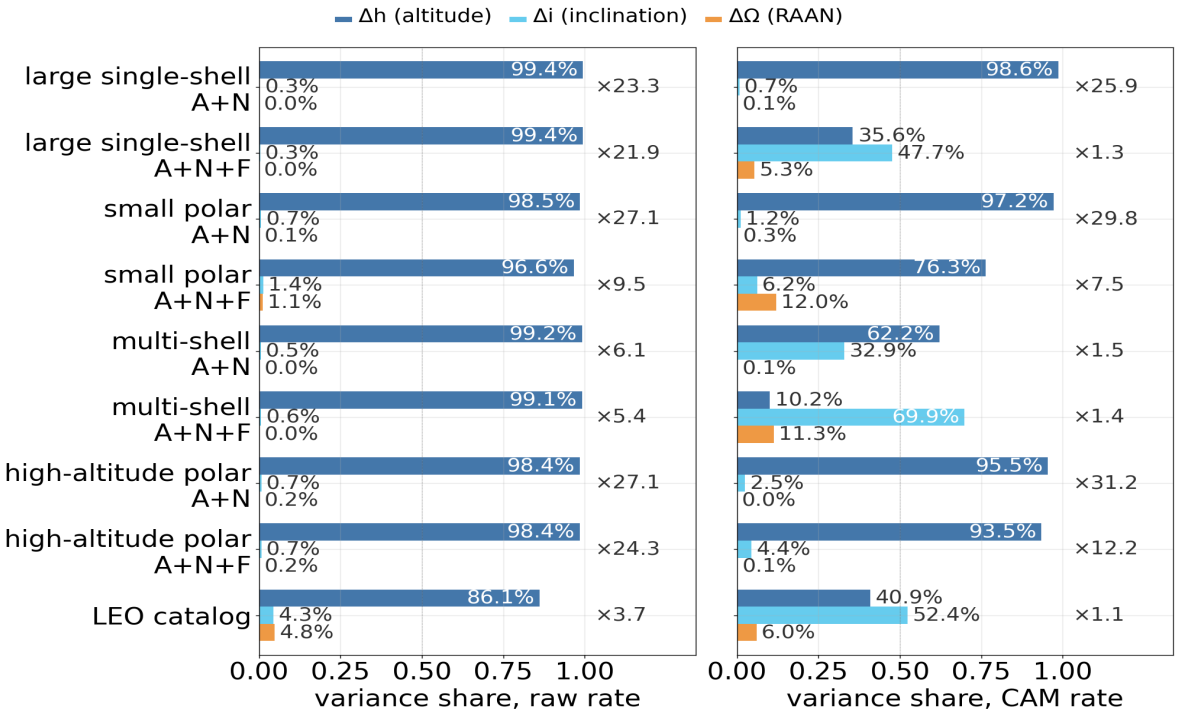

Figure 6.1: First-order variance shares of the joint $(\Delta h, \Delta i, \Delta\Omega)$ cube, for the raw rate (left) and the CAM-realized rate (right). The $\times n$ annotation on each row is the total swing of the response across that cube.

With the machinery in place, fig. 6.1 decomposes the variance of the static rate over the joint spatial cube, raw on the left and CAM-realized on the right, one row per archetype and population composition. The altitude bin width is swept over a ladder from 5 km to 100 km, spanning what a practitioner might plausibly pick, against the same grid of four inclination and four RAAN settings used throughout; the time step is absent by construction, since the static rate has none. For the raw rate the outcome is unambiguous: altitude owns the large majority of the variance in every case, the angular parameters and their interactions never amount to much, and the one partial exception is the heterogeneous catalog. This matches the divergence that section 6.2.1 derives analytically below, and it is the empirical justification for treating Δh first and separately in the rest of this chapter. The CAM-realized rate tells a less uniform story: altitude still dominates in most cases, but for two of the four archetypes carrying fragments, the large single-shell one and the multi-shell one, its share collapses and inclination takes over instead, with interactions also becoming non-negligible. Section 6.2.2 explains that reversal. The catalog row is not an isolated case; it is the extreme end of a pattern already visible among the others. Wherever altitude's share collapses, the total swing collapses with it: to 1.3 and 1.4 for the two fragment-carrying archetypes just discussed, to 1.5 for the multi-shell fleet even without fragments, and to 1.1 for the catalog. Read against that swing rather than its share, each of these four answers is close to insensitive to the lattice altogether, so the shares are not a competing finding so much as a description of an already-flat response. That the same four rows lose their altitude share and their total swing together is a first hint that altitude's dominance depends on how heterogeneous the population's own altitude structure already is, a thread section 6.2.2 picks up next.

The complementary question is what remains once Δh has been settled. Figure 6.2 decomposes the time-integrated response over the $(\Delta i, \Delta\Omega, \Delta t)$ cube, with the altitude bin width held at one value per archetype. Before comparing shares, the swing itself is worth checking: no case moves by more than a factor 1.7 across the whole cube, and six of the eight by less than 1.6, against the factors of 5 to 30 that the altitude ladder produced in fig. 6.1. Whatever is left to decide after Δh is therefore a second-order concern, which is the practical justification for the ordering of this chapter.

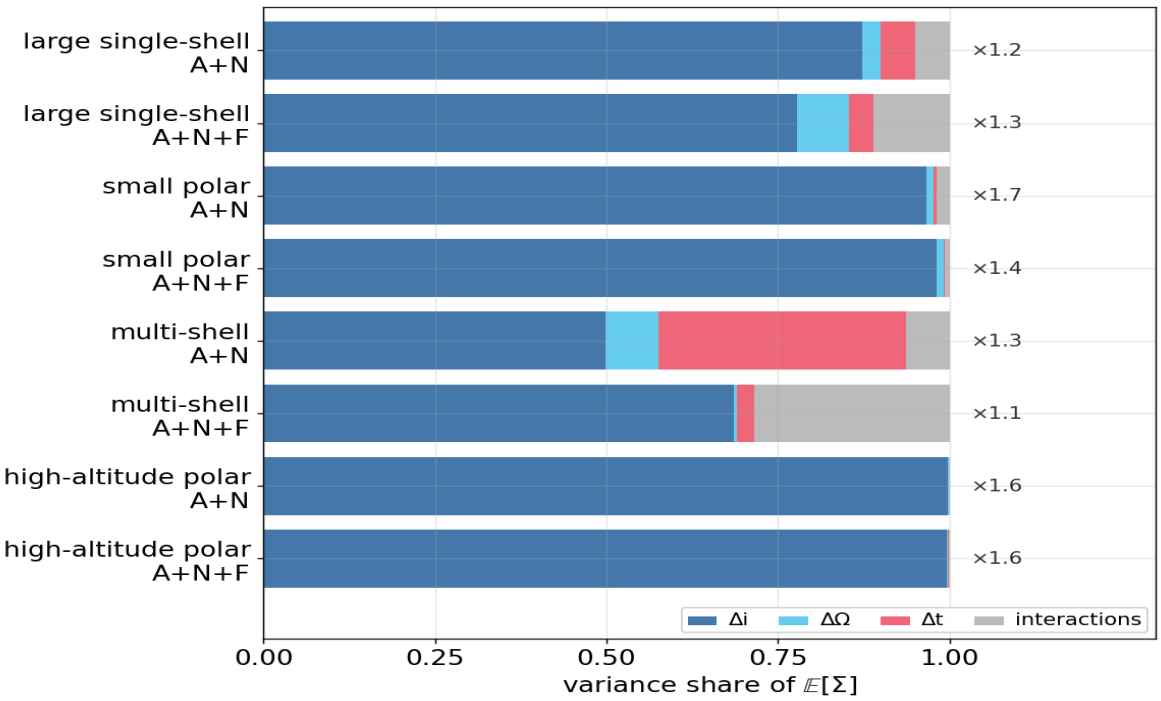

Figure 6.2: Variance shares of the time-integrated $\mathbb{E}[\Sigma]$ over the $(\Delta i, \Delta\Omega, \Delta t)$ cube, with Δh held fixed per archetype. The $\times n$ annotation is the total swing of $\mathbb{E}[\Sigma]$ across the cube.

Within that residual the ordering is unambiguous. Inclination leads in all eight cases, from 49.9 % for the multi-shell fleet without fragments to 99.8 % for the high-altitude polar one, while RAAN never exceeds 7.7 %; section 6.2.3 explains why the two angular axes are so unevenly matched. The time step contributes least almost everywhere, at or below 5.1 % in seven of the eight cases and indistinguishable from zero for both polar archetypes. The one exception is the multi-shell fleet without fragments, where it reaches 36.1 %; section 6.3 picks up that thread. Interaction terms stay modest throughout except for the multi-shell archetype carrying fragments, where they reach 28.5 %, with two thirds of that in the inclination-time step term alone.

6.2. Lattice Discretization

6.2.1. Altitude Binning

At first glance, it might seem that the rate would diverge with Δh. Indeed,

$$\dot{\eta}_{S_i S_j} \propto \frac{x_{S_i} x_{S_j}}{\Delta h^2} \left(\frac{1}{R_\oplus + h_\ell} - \frac{1}{R_\oplus + h_{\ell+1}} \right) \approx \frac{x_{S_i} x_{S_j}}{\Delta h^2} \frac{\Delta h}{(R_\oplus + h^*)^2} \propto \frac{x_{S_i} x_{S_j}}{\Delta h} \tag{6.4}$$

If this were the case, each choice of Δh would be arbitrary and yield different results. Whether the divergence actually occurs, however, depends on how the objects are distributed, since the real distribution enters the rate computation through the numerator too, not just through the explicit Δh in the denominator. For simplicity, consider only node pairs of the same species. The total rate within a given altitude bin is

$$\dot{\eta} \propto \frac{x_S^2}{\Delta h} \tag{6.5}$$

where x_S is the total number of objects of the given species in the bin. Refining Δh splits co-located objects into different altitude bins, and since the cross-bin rate is zero by construction, it effectively determines which objects are allowed to collide with each other. If the objects are well spread in space, refining Δh shrinks x_S proportionately, and the ratio $x_S^2/\Delta h$ stays bounded. If instead the objects are delta-distributed, no refinement separates them: they stay co-located in the same altitude bin regardless of Δh, so x_S never changes and the rate genuinely diverges. In practice this only occurs for active satellites, which are all bunched on the same orbital plane, fixed in time. It is a lesser problem than it might seem, since collisions involving active satellites are already virtually suppressed by the CAM success rate; the choice of Δh should instead be driven by the species that actually cause collision hazard, faulty satellites and fragmentation clouds. Why the same refinement applied to the two angular axes produces nothing comparable is not obvious from the scaling alone and is explored in section 6.2.3.

Figure 6.3 makes the divergence quantitative. The left panel plots the raw rate against the altitude bin width on logarithmic axes, together with a reference line of slope -1. A population lying along that line obeys $\dot{\eta} \propto 1/\Delta h$ exactly, which on log axes is what inverse proportionality looks like, and which in practice means that halving the bin width doubles the reported collision rate with nothing physical having changed. The right panel plots the product $\dot{\eta}\,\Delta h$, which divides that scaling out, so the same behaviour appears as a horizontal line instead. The two panels are therefore two views of a single statement, and the reference line exercises the undesired behavior: lying on slope -1 in (a), equivalently staying flat in (b), is the *worst* outcome, because it means the reported rate is fixed entirely by an arbitrary numerical choice and no bin width can be defended over any other. The desirable behaviour is the opposite, a rate that flattens out in (a) and therefore rises with slope $+1$ in (b), which marks a population whose rate is insensitive to the lattice and can be quoted as a physical number.

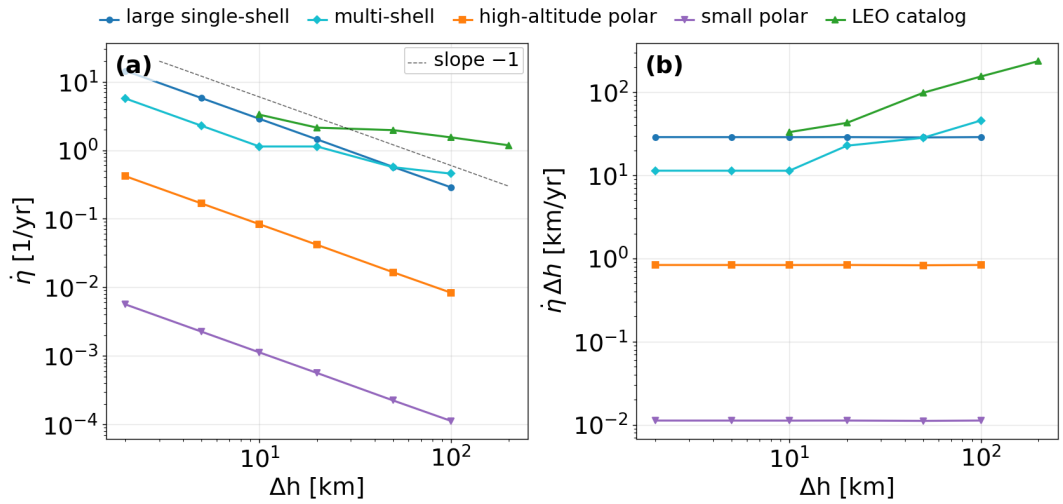

Figure 6.3: Raw collision rate versus altitude bin width for the five archetypes. (a) the rate itself, with a slope -1 reference; (b) the product $\dot{\eta}\,\Delta h$, constant wherever the $1/\Delta h$ law of eq. (6.4) holds exactly.

Read against that reference, the large single-shell, the high-altitude polar and the small polar archetypes are the pathological cases. Their product holds constant to within 1 % while the bin width changes by a factor 50, from 2 km to 100 km, so they follow the divergence over the entire ladder. Each occupies a single thin layer, so no refinement ever separates its members, x_S in eq. (6.5) never falls, and the $1/\Delta h$ law is reproduced exactly. For these three the raw rate is not a physical quantity at all, only a restatement of the bin width used to compute it. The multi-shell archetype tracks the same divergence down to 10 km and departs above it, precisely the separation between its four shells: once a bin is wider than that, it merges shells PRISM should keep apart, and the population per bin grows instead of staying fixed. The heterogeneous catalog is the one well-behaved case, its rate falling by only a factor 2.8 while the bin width grows by a factor 20, well away from the reference slope: spread over the whole altitude range, it gains members in a bin roughly in proportion to the bin width, so the two effects in eq. (6.4) nearly cancel and the rate becomes largely a physical number instead of a lattice artifact. The lesson is therefore that susceptibility to the divergence is a property of the population, not the lattice, so it cannot be fixed by refining alone, and it is the four constellation archetypes, not the real catalog, that need a defensible Δh most.

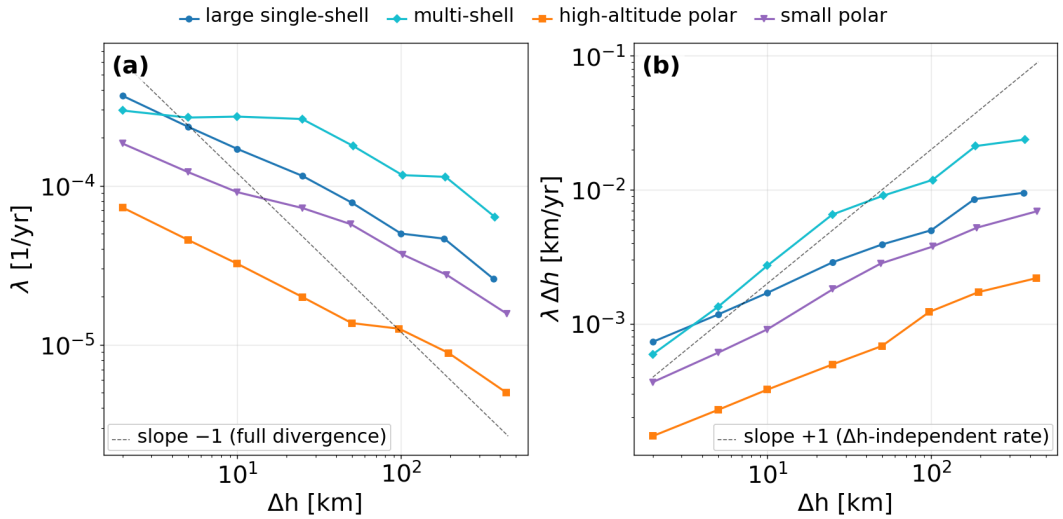

Figure 6.4: (a) The rate realized after avoidance maneuvers against bin width. (b) The product $\lambda\,\Delta h$, flat if the rate diverges as $1/\Delta h$, with slope $+1$ if it is fully independent of the bin width.

The divergence of fig. 6.3 is a statement about a population that occupies one thin layer, and a constellation does not stay that way. Satellites fail and are left behind, collisions produce fragment clouds, and the population that results is spread over altitude even when the operational fleet is not. Figure 6.4 repeats the ladder for each archetype once derelicts and a fragment cloud are present. Unlike the previous figure, this one plots λ, the rate realized after avoidance maneuvers, instead of the raw rate $\dot{\eta}$, so as to isolate the species-pair combinations, or channels (such as N-N or N-F), that avoidance does not suppress. Spread objects such as derelicts and fragments soften the divergence: the curves in the figure follow a slope shallower than the $1/\Delta h$ law, rather than reproducing it.

The single N-N channel needs one further distinction. A fleet that loses its entire population of faulty satellites at once, before any of them has had time to decay and spread across the range of altitudes they occupy after failing, behaves as a delta-distribution and diverges, too. Figure 6.5 compares the N-N rate for such a co-located, simultaneous failure against a staggered one, in which the same number of derelicts fail stochastically over time and are therefore spread over that decay range at any given instant. The co-located case reproduces the $1/\Delta h$ divergence of eq. (6.5) exactly, and is unreliable at any resolution. The staggered case is bounded instead: the rate plateaus over a wide range of bin widths, and only drops off once bins are fine enough to resolve the individual derelicts one per bin, at which point it under-counts rather than diverges. The divergence is therefore not an intrinsic property of derelict populations but an artifact of the idealized simultaneous failure mode. In practice failures are staggered in time, so the co-located case is a pathological limit, not a realistic scenario.

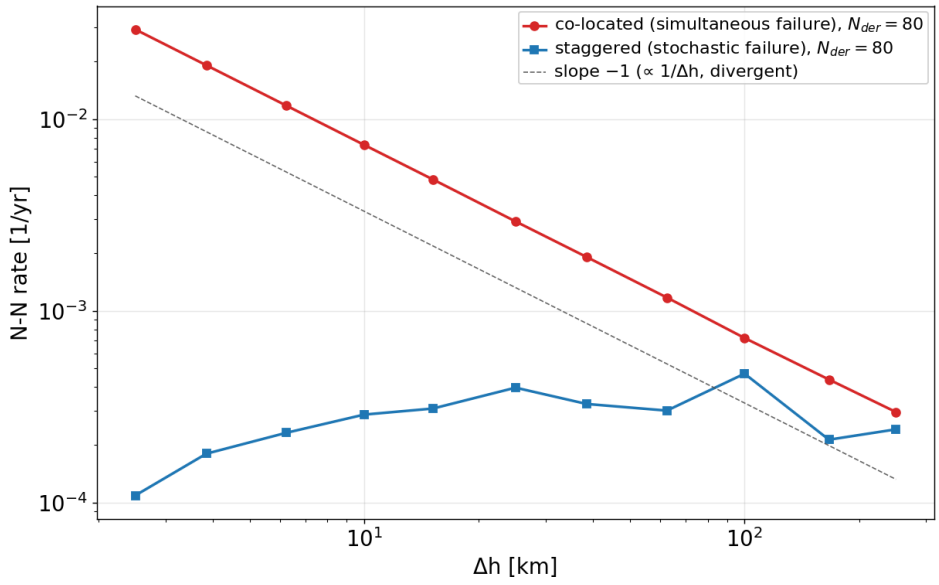

Figure 6.5: Instantaneous N-N rate versus altitude bin width for co-located versus staggered derelict failure, same derelict count and angular diversity in both cases. Co-located failure reproduces the genuine $1/\Delta h$ divergence; staggered failure is bounded over the operational window.

The issue is framed more formally by introducing the effective width, or participation ratio, of the distribution

$$W = \frac{\left(\int n_h \, \mathrm{d}h\right)^2}{\int n_h^2 \, \mathrm{d}h} = \frac{1}{\int f_h^2 \, \mathrm{d}h} \tag{6.6}$$

where $n_h$is the distribution of the objects over altitude, obtained by marginalizing the phase-space density $n_{\boldsymbol{u}}$ over inclination, RAAN and argument of latitude, and f_h is its normalized PDF. Imagine drawing two objects independently from the population. Then the denominator in eq. (6.6) is precisely the probability density of finding the two objects in the same bin. If the distribution were uniform of width L, the above integral would evaluate exactly to $1/L$.

Therefore, W is the width of a uniform population that would give the same probability of finding two objects in the same altitude bin as the real population.

Example 6.2.1: Effective Width of a Uniform Distribution

Consider a uniformly distributed population with density

$$f_h = \begin{cases} \dfrac{1}{\Delta h} & \text{for } h_\ell \le h \le h_{\ell+1} \\ 0 & \text{otherwise} \end{cases}$$

then

$$\int_{-\infty}^{\infty} f_h^2 \, \mathrm{d}h = \frac{1}{\Delta h^2} \int_{-\infty}^{\infty} \mathrm{d}h = \frac{1}{\Delta h} \quad \Rightarrow \quad \boxed{W = \Delta h}$$

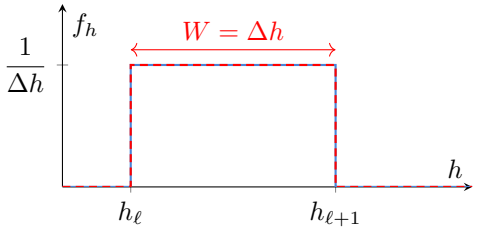

Figure 6.6: The effective width of a uniform distribution coincides with the distribution's own width.

Example 6.2.2: Effective Width of a Delta Distribution

Consider a delta-distributed population $f_h = \delta(h - h^*)$. In this case

$$\int_{-\infty}^{\infty} f_h^2 \, \mathrm{d}h = \int_{-\infty}^{\infty} \delta(h - h^*)\delta(h - h^*) \, \mathrm{d}h = \delta(0) \to \infty \quad \Rightarrow \quad \boxed{W = 0}$$

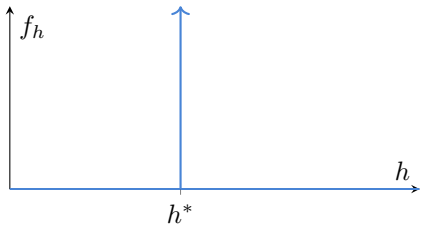

Figure 6.7: The whole population is concentrated in an infinitesimally small width.

Let $X_a \in \{1, \dots, N_h\}$ be a categorical random variable describing the bin where object a is located. The probability of finding object a in the ℓ-th bin is $P(X_a = \ell) = p_\ell$. The co-location probability, or the probability of finding two objects a and b in the same altitude bin, is

$$\mu_2 = P(X_a = X_b) = \sum_\ell p_\ell^2 \tag{6.7}$$

As $\Delta h \to 0$, $\mu_2 \to \Delta h / W$, so W is the fine-bin normalization of the co-location probability, and it is the rate at which co-located objects disappear from a bin as Δh is refined. In other words, eq. (6.5) reads as

$$x_S^2 \propto \frac{\Delta h}{W} \quad \Rightarrow \quad \dot{\eta} \propto \frac{1}{W}$$

and the raw rate scales as the probability density of finding two objects in the same bin. For a delta-distributed population $W = 0$, and no lattice can resolve the distribution, no matter how

refined, making the rate diverge. Contrarily, if the population is well-spread, co-located objects disappear linearly with Δh and the rate plateaus. In practice, once Δh grows comparable to W the objects are merged into a single layer and the information about their distribution is washed out. Assuming the objects are uniform across the whole bin understates their true local density, so the rate is *underestimated*. The amount lost is measured by the *dilution factor*

$$\frac{\dot{\eta}}{\dot{\eta}_\infty} = \mu_2 \frac{W}{\Delta h} \tag{6.8}$$

the ratio of the rate a lattice of width Δh reports to the one an infinitely fine lattice would. In the same limit $\Delta h \to 0$, the dilution converges to 1, and it falls away from unity as soon as the bin starts to swallow structure.

While refining the lattice under the effective width remediates the divergence issue, it has the downside of introducing noise. As Δh decreases, there is no systematic error, but the number of pairs sharing a bin also decreases, causing larger between-seed fluctuations. The number of co-located pairs in the same altitude bin, K, follows from the multinomial occupancy of the lattice, and its variance splits into a shot-noise term and a covariance term:

$$\mathrm{Var}(K) = \underbrace{\binom{x_S}{2}\mu_2(1-\mu_2)}_{\text{shot noise}} + \underbrace{2N\binom{x_S-1}{2}(\mu_3-\mu_2^2)}_{\text{overlapping-pair covariance}} \tag{6.9}$$

Proof 6.2.1

The number of co-located pairs in the same altitude bin is

$$K = \sum_{a<b} \mathbb{1}_{ab} \quad \text{where } \mathbb{1}_{ab} = \begin{cases} 1 & \text{if } a \text{ and } b \text{ are in the same bin} \\ 0 & \text{otherwise} \end{cases}$$

The indicator $\mathbb{1}_{ab}$ is Bernoulli distributed with probability

$$P(\mathbb{1}_{ab}=1) = P(X_a = X_b) = \mu_2 \quad \Rightarrow \quad \mathbb{E}[\mathbb{1}_{ab}] = \mu_2$$

hence, the expected number of co-located pairs is

$$\mathbb{E}[K] = \binom{x_S}{2}\mu_2 = \frac{x_S(x_S-1)\mu_2}{2}$$

while its variance is

$$\mathrm{Var}(K) = \sum \mathrm{Var}(\mathbb{1}_{ab}) + 2\sum \mathrm{Cov}(\mathbb{1}_{ab}, \mathbb{1}_{cd})$$

The first term is simply the variance of Bernoulli indicators

$$\mathrm{Var}(\mathbb{1}_{ab}) = \mu_2(1-\mu_2)$$

The second term accounts for the covariances of all pairs of pairs

$$\mathrm{Cov}(\mathbb{1}_{ab}, \mathbb{1}_{cd}) = \mathbb{E}[\mathbb{1}_{ab} I_{cd}] - \mathbb{E}[\mathbb{1}_{ab}]\,\mathbb{E}[I_{cd}]$$

with $\mathbb{E}[\mathbb{1}_{ab} I_{cd}] = P(X_a = X_b, X_c = X_d)$. There are two possible cases. First, consider two pairs of distinct objects $\mathbb{1}_{12}$ and $\mathbb{1}_{34}$. Because the objects are independent, $\mathbb{E}[\mathbb{1}_{12}\mathbb{1}_{34}] = P(X_1 = X_2)P(X_3 = X_4) = \mu_2^2$. Therefore, the covariance between *disjoint* pairs is

$$\mathrm{Cov}(\mathbb{1}_{12}, \mathbb{1}_{34}) = \mu_2^2 - \mu_2^2 = 0$$

always. Now consider two pairs sharing one object $\mathbb{1}_{12}$ and $\mathbb{1}_{13}$. Then, $\mathbb{1}_{12}\mathbb{1}_{13} = 1$ if and only if all three objects are co-located in the same bin, thus

$$\mathbb{E}[\mathbb{1}_{12}\mathbb{1}_{13}] = P(X_1 = X_2 = X_3) = \sum_\ell p_\ell^3 = \mu_3$$

Therefore, the covariance of *overlapping* pairs is $\mathrm{Cov}(\mathbb{1}_{12}, \mathbb{1}_{13}) = \mu_3 - \mu_2^2$. Overall, the variance is eq. (6.9).

Because of the Cauchy-Schwarz inequality, $\mu_3 \geq \mu_2^2$, so the covariance is always positive and always increases the noise. It vanishes only for an occupancy distribution close to uniform, where $\mu_3 \approx \mu_2^2$, which is not the case for the populations of interest here, which are strongly peaked as will be shown later on. The full term is therefore kept, and the fluctuation of the number of co-located pairs across multiple runs is measured by the *coefficient of variation*

$$\frac{\sqrt{\mathrm{Var}(K)}}{\mathbb{E}[K]} \tag{6.10}$$

A smaller coefficient of variation means less variability and greater consistency between samples. Putting everything together, the relative error due to the discretization in altitude is the sum of a bias term and a noise term

$$\varepsilon(\Delta h) = \underbrace{\left(1 - \frac{\dot{\eta}}{\dot{\eta}_\infty}\right)}_{\text{bias}} + \underbrace{\frac{\sqrt{\mathrm{Var}\,K}}{\mathbb{E}[K]}}_{\text{noise}} \tag{6.11}$$

Both terms follow from the occupancy moments μ_2 and μ_3 alone, with no assumption about the shape of the layer, so eq. (6.11) holds for any altitude distribution. What it does require is those two moments, which is the practical difficulty and the subject of the rest of this section.

Requiring the error to stay below a tolerance, $\varepsilon(\Delta h) \leq \varepsilon_0$, identifies a *reliable window* of admissible bin widths. The two terms of eq. (6.11) pull in opposite directions: refining the lattice removes dilution but starves the bin of pairs, while coarsening it suppresses the noise but merges the layer. Their sum ε is therefore U-shaped in Δh, the window is the interval between its two roots, and its lower edge is set by the noise, its upper edge by the dilution. Consequently, the minimum of ε over Δh is a *noise floor*: an error that no choice of lattice can go below, fixed by the population's own count and concentration. A tolerance tighter than the floor admits no bin width at all, and the only option is to accept a larger error.

Because different species evolve in different ways, they occupy the altitude axis differently, and their occupancy moments differ accordingly. A fragmentation cloud is born narrow and rich, then spreads under differential drag and thins as objects re-enter; faulty satellites start co-located at the shell that shed them and spread only as fast as drag carries them down. Each channel therefore has its own reliable window, and the bin width has to be chosen for the channel that carries the highest collision hazard. The obstacle is circularity: μ_2 and μ_3 are properties of a binned population, so measuring them from a simulation presupposes the very Δh they are meant to fix. What follows removes that circularity by obtaining the moments from the mission profile instead, analytically for the derelict band and from a single breakup sample for a fragment cloud, neither of which requires a lattice to be chosen first.

6.2.2. Selecting Altitude Binning from the Mission Profile

Derelict Band

At the end of their operational lifetime, active satellites attempt PMD with a failure rate γ. Therefore, faulty satellites are supplied at the operational semi-major axis a_0 at a rate of

$$\vartheta = \frac{\gamma x_A}{T_{\mathrm{EOL}}}$$

with units of satellites per year, where x_A is the number of active satellites and T_{EOL} is the lifetime of the satellites. Throughout this chapter $\gamma = 5\,\%$ and $T_{\mathrm{EOL}} = 5\,\mathrm{yr}$, and x_A is the full constellation, summed over shells. The resulting inflows are 15.8 derelicts per year for the large single-shell archetype, 14.4 for the multi-shell one, 5.9 for the high-altitude polar one and 0.66 for the small polar one, and the same numbers are used both in the analytical estimates below and in the simulations they are compared against.

Over time, the failed objects decay due to drag. If new faulty objects keep being supplied at the operational shell at a rate higher than the rate at which they decay, the derelict population forms a band over altitude. Remarkably, the dilution of a drag-swept band depends on the bin width only through $\Delta h / H_\ell$, where H_ℓ is the scale height.

Proof 6.2.2

A circular orbit under drag decays with the secular rate

$$\dot{a} = -\rho(a) B a v = -\rho(a) B \sqrt{a \mu_\oplus}$$

where B is the ballistic coefficient. The residence density of the objects is equal to the flux ϑ over the descent speed

$$n_a = \frac{\vartheta}{|\dot{a}|} = \frac{\vartheta}{\rho(a) B a v}$$

Marginalizing over all altitudes in a band of width $L = a_0 - a$, the number of faulty satellites in the band is

$$x_N = \int_a^{a_0} n_a \, \mathrm{d}a = \vartheta \int_a^{a_0} \frac{\mathrm{d}a}{\rho(a) B a v} = \vartheta \tau$$

where τ is the decay time. The above expression is precisely *Little's law*, which states that the number of individuals is the product of the arrival rate times the time spent in the system. For an exponential atmosphere

$$\rho(a) = \rho_\ell \exp\left(-\frac{a - a_0}{H_\ell}\right)$$

where ρ_ℓ and H_ℓ are the reference density and the scale height at the shell. Since $av = \sqrt{a\mu_\oplus}$ changes much more slowly than the exponential term does, the integral is dominated by the latter, and the former can be kept constant at $a_0 v_0 = \sqrt{a_0 \mu_\oplus}$. The decay time is

$$\tau = \frac{1}{\rho_\ell B \sqrt{a_0 \mu_\oplus}} \int_a^{a_0} \exp\left(\frac{a - a_0}{H_\ell}\right) \mathrm{d}a = \frac{H_\ell}{\rho_\ell B \sqrt{a_0 \mu_\oplus}} \left[1 - \exp\left(-\frac{L}{H_\ell}\right)\right]$$

In the limit, a, which decreases with time, reaches the re-entry value. Then, $L \gg H_\ell$ and

$$\tau \to \tau_\infty = \frac{H_\ell}{\rho_\ell B \sqrt{a_0 \mu_\oplus}}$$

Let $A = \vartheta / (\rho_\ell B \sqrt{a_0 \mu_\oplus})$. The marginal of the density n_a is

$$\int_a^{a_0} n_a \, \mathrm{d}a = \vartheta \tau = A H_\ell \left[1 - \exp\left(-\frac{L}{H_\ell}\right)\right]$$

while the marginal of the density squared is

$$\int_a^{a_0} n_a^2 \, \mathrm{d}a = A^2 \int_a^{a_0} \exp\left(\frac{a - a_0}{H_\ell}\right) \mathrm{d}a = A^2 \frac{H_\ell}{2} \left[1 - \exp\left(-\frac{2L}{H_\ell}\right)\right]$$

Finally, the effective width is

$$W_N = \frac{A^2 H_\ell^2 \left(1 - e^{-L/H_\ell}\right)^2}{A^2 H_\ell / 2 \left(1 - e^{-2L/H_\ell}\right)} = 2 H_\ell \frac{1 - e^{-L/H_\ell}}{1 + e^{-L/H_\ell}} = 2 H_\ell \tanh\left(\frac{L}{2 H_\ell}\right) \tag{6.12}$$

The same residence density also gives the occupancy moments outright, and with them the error of eq. (6.11) at any bin width. Integrating n_a over the ℓ-th bin of a lattice tiling the band gives occupancies in geometric progression, so the sums of their powers are geometric series that telescope. For the second moment,

$$\mu_2 = \frac{\left(1 - e^{-\Delta h / H_\ell}\right)^2}{\left(1 - e^{-L/H_\ell}\right)^2} \cdot \frac{1 - e^{-2L/H_\ell}}{1 - e^{-2\Delta h / H_\ell}} = \tanh\left(\frac{\Delta h}{2 H_\ell}\right) \coth\left(\frac{L}{2 H_\ell}\right) \tag{6.13}$$

capped at unity once $\Delta h \geq L$ and a single bin holds the whole band, and the third moment follows by the same route with the exponents tripled. Substituting eq. (6.13) and eq. (6.12) into the dilution factor of eq. (6.8), the band depth cancels exactly and

$$\frac{\dot{\eta}}{\dot{\eta}_\infty} = \mu_2 \frac{W}{\Delta h} = \frac{\tanh\left(\Delta h / (2 H_\ell)\right)}{\Delta h / (2 H_\ell)} \tag{6.14}$$

Bounding $1 - \dot{\eta}/\dot{\eta}_\infty \leq \varepsilon_0$ fixes the coarsest admissible bin as a pure multiple of the local atmospheric scale height H_ℓ, which at the 69 km scale height of the 550 km shell means 55 km to 163 km across tolerances ε_0 of 5 % to 30 %. Eq. (6.14) assumes a bin edge sits exactly at a_0, so the lattice's tiling matches the band's leading edge. Then, the population's core (concentrated near a_0 where drag is weakest) does not coincide with the bin center, the point at which the rate model evaluates the pairwise relative velocity. his offset can make the realized rate fall short of the value predicted by eq. (6.14) by as much as a factor of two, so eq. (6.14) is an upper limit rather than a sharp threshold. Furthermore, it holds only for a band fed by staggered failures; a cohort that fails simultaneously is genuinely co-located and no bin width resolves it.

The band itself spreads over time, in fact, $L(t) = a_0 - a(t)$. Specifically

$$\frac{\mathrm{d}L}{\mathrm{d}t} = -\dot{a} = -\rho(a) B \sqrt{a\mu_\oplus} \approx -\rho_\ell \exp\left(-\frac{a - a_0}{H_\ell}\right) B\sqrt{a_0\mu_\oplus}$$

Since $L(t = 0) = 0$, separating variables and integrating yields

$$L(t) = -H_\ell \ln\left(1 - \frac{t}{\tau_\infty}\right) \tag{6.15}$$

At $t = 0$, the band has not yet formed, and the population behaves like a delta-distribution of infinitesimal width. Therefore, no Δh can resolve the distribution, but it never needs to since the population carries zero hazard. As time progresses, the band expands. As long as the band stays thin ($L \ll H_\ell$), the effective width is $W_N = L$ because the distribution is nearly uniform over the full band. However, even if the band develops fully down to the re-entry altitude, the effective width saturates to a value of $W_N \to 2H_\ell$. This occurs because faulty objects spend most of their time near the top of the shell, where drag is weakest. Figure 6.8 sketches both regimes.

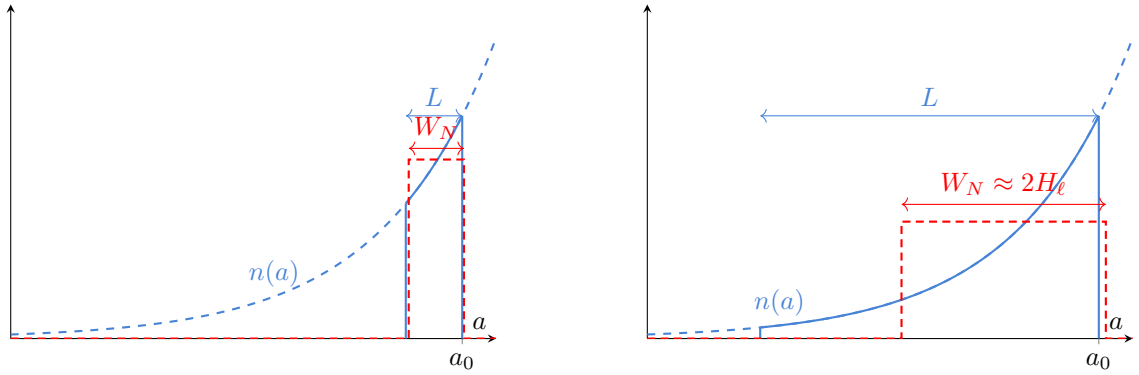

Figure 6.8: Trend of the (un-normalized) altitude distribution of faulty satellites. Early on (left), the objects are bunched near the operational altitude, and the distribution is nearly uniform. As time passes (right), objects at lower altitudes decay faster, and the effective width saturates.

When the band is narrow, it also carries the least collision hazard. The recommended window should instead be evaluated when the band has had time to develop and carries a higher risk. Therefore, the object count and effective width are computed at the minimum between τ_∞ and t_f, the final simulation time, whichever comes first. At low LEO $\tau_\infty < t_f$, the band is fully developed and $W_N \to 2H_\ell$. Above 600 km $\tau_\infty \gg t_f$ (708 years at Iridium's altitude, 4900 years at OneWeb's), so within a mission the band develops a depth $L(t)$ of well under a kilometer, the derelicts are effectively co-located (near-delta), and PRISM returns a sub-km, practically unresolvable window.

Figure 6.9 makes this argument quantitative. The left panel tracks $W_N(t)$ for the four archetypes, each using its own ballistic coefficient and operational altitude: the low-altitude archetypes approach their steady-state $2H_\ell \tanh(L/2H_\ell)$ well within the ten-year analysis horizon, while the high-altitude ones are still essentially undeveloped by the time the horizon is reached. The right panel generalizes this across a family of ballistic coefficients, evaluating the

steady-state width at the fixed horizon as a function of shell altitude: below roughly 600 km the band has enough time to develop regardless of B, whereas above it τ_∞ exceeds the horizon so much that the width stays small no matter how draggy the object is. The threshold is therefore a property of the decay timescale relative to the mission horizon, not of any one archetype's parameters.

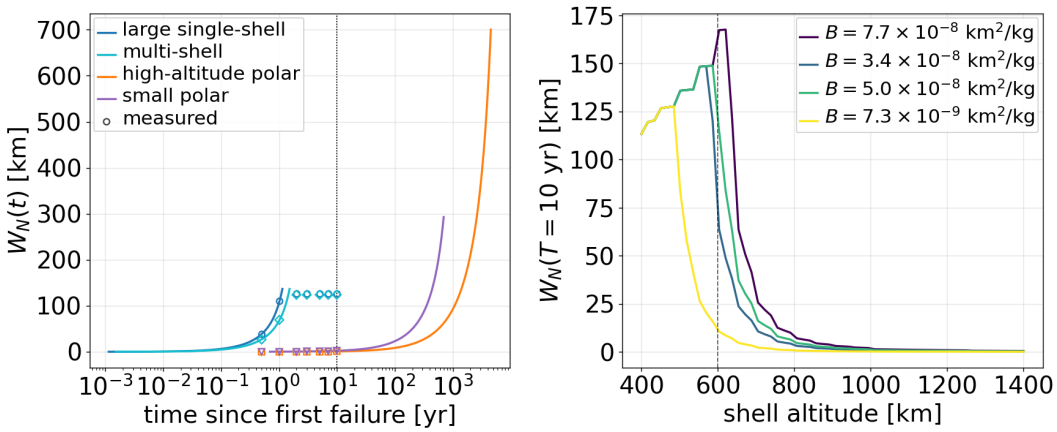

Figure 6.9: (left) Development of the derelict band's effective width $W_N(t)$ over time for each archetype. (right) Steady-state width at the horizon versus shell altitude, for a family of ballistic coefficients:.

Fragmentation Clouds

The other non CAM-suppressed population is the one constituted of fragments. The exact distribution in orbital elements of a debris cloud ejected from a collision can be obtained directly from the output of NASA's SBM adopting the method described by Frey and Colombo (2021) already used in eq. (4.15) to transform a distribution between two phase-spaces via the Dirac generalized function. To a first approximation, a fragment receiving a Δv kick undergoes a change in semi-major axis and eccentricity proportional to the fractional velocity change $\Delta v/v$. While the specific magnitude changes depending on the direction of the kick, impulses along the velocity direction dominate, and the fragments' orbits are found near the parent's one. For instance, a fragment receiving a tangential kick in the direction of flight makes the collision point r_0 its new perigee, while a fragment receiving a kick in the opposite direction makes it its new apogee, and because the kick is isotropic (every direction is equally likely), the altitude range of the cloud is centered in r_0 and extends equally in both directions. However, since Δv follows a log-normal distribution, much of the cloud remains close to r_0, whereas the tail end of the distribution produces few highly-energetic fragments far away (Frey & Colombo, 2021).

As time advances, the cloud spreads due to differential drag, and thins as objects re-enter. In other words, most of the risk associated with fragments-vs-fragments collisions is attributable to the first instants of the cloud's life. Therefore, the recommended window should be computed at cloud formation, where the hazard is highest. The simplest and most direct way to quantify the effective width of the fragment population W_F is to simulate a collision at the operational altitude and empirically measure the spread of the fragments over altitude. Table 6.2 reports the measured values for the four archetypes, along with the root mean square kick Δv_{rms}. The occupancy moments μ_2 and μ_3 of the cloud come from the same sample.

Table 6.2: Measured fragment-cloud widths W_F.

Archetype	a [km]	Δv_{rms} [m/s]	W_F [km]
large single-shell	6921	208.4	217.8
multi-shell	6941	199.9	210.1
high-altitude polar	7571	164.4	267.3
small polar	7151	228.4	221.8

The number of fragments x_F is also inferred directly from the initial population, using eq. (4.42) of the SBM. Because x_F only sets the noise edge, and more fragments only lower the noise and widen the window, sampling a single collision is conservative.

Cross-species Distributions

Until now, only the self-species case was considered. Nothing about the construction changes when the pair straddles two species: the bin-sharing indicator is still Bernoulli, it now just requires one object of each. Writing p_ℓ and q_ℓ for the occupancies of species N and F, the number of co-located cross pairs is

$$K_{NF} = \sum_{\ell} n_h^{(N)} n_h^{(F)}$$

and the moments that play the roles of μ_2 and μ_3 are the cross overlap μ' and the two mixed triples μ_N and μ_F,

$$\mu' = \sum_{\ell} p_\ell q_\ell \qquad \mu_N = \sum_{\ell} p_\ell q_\ell^2 \qquad \mu_F = \sum_{\ell} p_\ell^2 q_\ell$$

so that

$$\mathbb{E}\left[K_{NF}\right] = x_N x_F Q \tag{6.16}$$

$$\operatorname{Var}(K_{NF}) = x_N x_F Q(1-Q) + x_N x_F (x_F - 1)(\mu_N - Q^2) + x_F x_N (x_N - 1)(\mu_F - Q^2) \tag{6.17}$$

The two triple terms are the cross-species counterparts of the overlapping-pair covariance, one for each species supplying the shared object, and the coefficient of variation follows exactly as in eq. (6.10). The noise edge is set by the sparser of the two populations, which for a derelict-fragment channel means the handful of faulty satellites, so the channel prefers coarse bins.

The dilution factor is likewise unchanged in form,

$$\frac{\dot{\eta}}{\dot{\eta}_\infty} = \mu' \frac{W_{NF}}{\Delta h} \qquad \text{with} \quad W_{NF} = \left(\int f_h^{(N)} f_h^{(F)} \, \mathrm{d}h \right)^{-1}$$

A fragment cloud is clustered around the altitude that produced it, i.e., the operational altitude, which is also the starting altitude of the derelict band's development. Hence, W_{NF} depends on how the cloud and the band sample each other, and how much the two have spread over time. Where the band has spread across the cloud, it samples the skirt as much as the spike and the product density is broader than either layer: the large single-shell and multi-shell archetypes carry bands 276 km and 307 km deep and measure $W_{NF} = 293\,\mathrm{km}$ and 456 km against clouds of 217 km and 210 km. Where the band has not developed, it sits instead on the cloud's spike, where $f_h^{(F)}$ is several times its own mean, and the cross width collapses to a fraction of the cloud's: the two high-altitude archetypes, whose bands deepen by under 3 km in a decade, measure 49 km and 52 km against clouds of 268 km and 221 km.

That second case also marks the limit of W_{NF} as a design quantity. An undeveloped band is a near-delta, and the overlap of a delta with anything is the value of the other density at a single point

$$W_{NF} = \left(\int \delta(h - h_0) f_h^{(F)}(h) \, \mathrm{d}h \right)^{-1} = \frac{1}{f_h^{(F)}(h_0)}$$

On the other hand, the mixed moments μ', μ_N and μ_F carry no such ambiguity, so substituting eqs. (6.16) and (6.17) into eq. (6.11) yields the reliable window Δh exactly.

The cross-species moments now explain the anomaly left open in fig. 6.1: the Sobol study found the multi-shell archetype carrying fragments to be the extreme case whose CAM-realized rate barely responded to the altitude bin width, at 10.2 % of the variance against 69.9 % for inclination, with the large single-shell fleet behind it at 35.6 %. Half the explanation is that the CAM-realized rate is dominated by the derelict-fragment channel once avoidance maneuvers

suppress every channel an active satellite takes part in, 92 % of the residual rate here and 95 % for the large single-shell and small polar fleets. Being a cross-species channel, its dilution edge sits at the width of the wider layer, the fragment cloud, rather than the narrow derelict band, so it survives coarsening that the self-species N-N channel of the same run does not. The high-altitude polar fleet confirms this by contrast rather than exception: its derelict band is still only 0.4 km deep after five years, so the same reasoning predicts no flattening, and indeed 68.9 % of its residual rate is the near-delta self-species N-N channel, leaving its response altitude-dominated at 93.5 %. The flat response of the multi-shell case therefore depends on its band having spread, a property of the population, not the lattice. The other half is specific to this archetype: coarsening merges its 10 km-spaced shells, and the pairs gained by the merge very nearly cancel the pairs lost to dilution, a numerical coincidence rather than a licence to coarsen, since satellites in distinct shells never actually collide. The shell separation is therefore a hard ceiling on Δh regardless of what it does to the variance: a lattice that cannot resolve a constellation's individual shells reports collisions that cannot happen.

With the moments in hand the reliable window of every channel follows from eq. (6.11), and table 6.3 evaluates it for the three channels collision avoidance does not suppress. Recall that the window is the set of bin widths whose total error $\varepsilon(\Delta h)$ stays at or below a tolerance ε_0, and that because ε is U-shaped in Δh that set is a single interval. Its two edges are the two competing error terms: the lower edge is where the noise term rises past the tolerance, so a lattice finer than it starves the bin of pairs, and the upper edge is where the dilution term does, so a lattice coarser than it merges the layer. The three rightmost columns are three choices of ε_0, that is, a 10 %, 20 % and 30 % ceiling on the relative error of the rate that channel contributes; a column reading *none* means no bin width at all meets that tolerance, and a lower edge of 0 means the noise term never reaches it, so refinement is unconstrained from below. The floor column is the minimum of ε over all bin widths and it runs from a fraction of a percent for the undeveloped bands, which hold tens of objects in a layer barely a kilometre deep, to 21 % for the derelict-fragment channel of the high-altitude polar archetype, whose overlap is sampled by 59 derelicts against a spike. Wherever the floor exceeds the tolerance the window is empty and no refinement recovers it. A 10 % target is unreachable for seven of the twelve channels, including every channel that dominates the realized rate of a low-altitude archetype. The chapter therefore designs to 30 %, the only one of the three columns under which every archetype has a non-empty window.

Table 6.3: Reliable window of the three unsuppressed channels of each archetype, at three tolerances ε_0. W is the effective width and N the object count of the sparser species. ‡ marks a window narrower than a factor 1.5; *none* marks a tolerance no bin width meets.

Archetype	**Chan.**	W [km]	x	**Floor** [%]	$[\Delta h_{\text{noise}}, \Delta h_{\text{dilution}}]$ [km]		
					$\varepsilon_0 = 10\%$	20%	30%
large single-shell	N-N	136.6	18	15.1	none	[20.3, 95.6]	[9.22, 146]
	F-F	217.8	462	10.5	none	[0, 12.1]	[0, 36.2]
	N-F	283.0	18	14.8	none	[103, 111]‡	[101, 115]‡
multi-shell	N-N	136.5	22	13.1	none	[13.9, 102]	[6.33, 149]
	F-F	209.1	1092	6.7	[0, 3.76]	[0, 20.3]	[0, 40.6]
	N-F	206.7	22	19.0	none	[19.9, 20.0]‡	[0, 27.0]
high-altitude polar	N-N	0.8	59	0.1	[0.05, 0.92]	[0.01, 1.04]	[0, 1.19]
	F-F	267.3	337	12.7	none	[0, 9.90]	[0, 37.4]
	N-F	51.2	59	21.2	none	none	[3.29, 3.32]‡
small polar	N-N	2.1	7	0.2	[1.81, 2.39]‡	[1.24, 2.69]	[0.81, 3.07]
	F-F	221.8	1031	7.3	[1.94, 2.75]‡	[0, 16.1]	[0, 36.8]
	N-F	51.0	7	9.9	[4.38, 4.38]‡	[4.38, 4.38]‡	[4.32, 4.45]‡

Read down the columns, the table separates three regimes. A rich cloud is one-sided: its floor is low, and every one of its F-F windows opens at a lower edge of 0, meaning the noise term never binds and no lattice is ever too fine for it, so the only active constraint is the dilution edge

on the right of the bracket, which sits between 36 km and 41 km at the design tolerance. A ceiling in the tens of kilometres is a permissive constraint, because it is coarser than the constraints the same population already faces from elsewhere: coarser than the 10 km shell separation of the multi-shell archetype, and coarser than the lower edge of the derelict band N-N in the rows above it. Wherever a cloud and a band coexist, the band's noise edge therefore binds before the cloud's dilution edge does, and the cloud alone never forces a choice of lattice. A developed derelict band is the opposite, two-sided and narrow, with both edges active: it admits bins from a few tens of kilometres up to about 150 km and nothing finer, because refinement starves it of pairs faster than it buys resolution. An undeveloped band is neither: it is a near-delta whose entire window sits at the kilometre scale, one to two orders of magnitude below any lattice a mission-length run would use, so its self-rate is a co-location artifact rather than a resolved quantity. The cross channel is the most demanding of the three, inheriting the cloud's narrow dilution edge and the band's small count at once, which is why most of its entries are knife edges or empty.

Table 6.4 collects what the rule prescribes once the channels are ranked and the shell-separation ceiling is imposed, and these are the altitude lattices used for the angular study of the next subsection. The percentage beside each binding channel says how much of the population's collision hazard that channel is expected to carry. It is obtained from the structure of the population alone, before any simulation: each channel is given a score, its pair count divided by its effective width and weighted by the avoidance factor α_{CAM}, and the percentage is that score over the sum of the scores of the channels competing with it. With derelicts alone the binding channel is N-N in every case. The large single-shell archetype admits 9.2 km to 146 km, wide enough to be comfortable but with an upper half coarser than any refinement intuition would suggest, and the lattice chosen inside it is 36.6 km, eleven bins over the whole altitude span. Where the band has not developed the window asks for bins of a kilometre or less; the lattice used there is 5 km, deliberately coarser, and those runs are dilution-limited by construction and marked as such. Once a fragment cloud is present the binding channel becomes F-F and the constraint flips from a lower edge to an upper one, between 36 km and 41 km for every archetype, so any lattice fine enough to resolve the shells satisfies it. The multi-shell archetype is where two constraints meet without conflicting: its shells are 10 km apart, its band admits from 6.3 km upward, and 8 km satisfies both. The exception is the high-altitude polar archetype, where the cross channel and the cloud both clear the relevance gate and their windows do not intersect, so the 5 km lattice used there is a declared fallback rather than a recommendation.

Table 6.4: Recommended altitude bin width per archetype and population composition, at the 30 % design tolerance. All widths in km. Window is the reliable range of the binding channel; ceiling is the smallest inter-shell separation. † marks a Δh outside its own window.

Archetype	Comp.	Binding channel	Window	Ceiling	Δh
large single-shell	A+N	N-N (98%)	[9.2, 145.6]	-	36.6
	A+N+F	F-F (94%)	[0.0, 36.2]	-	5.0
multi-shell	A+N	N-N (99%)	[6.3, 149.0]	10	8.0
	A+N+F	F-F (96%)	[0.0, 40.6]	10	5.0
high-altitude polar	A+N	N-N (100%)	[0.0, 1.2]	-	5.0†
	A+N+F	-	-	-	5.0†
small polar	A+N	N-N (100%)	[0.8, 3.1]	-	5.0†
	A+N+F	F-F (95%)	[0.0, 36.8]	-	5.0

The derelict band and the fragment cloud are both non-stationary, so a width is a property of a scenario at a time, while the lattice is fixed once for the whole run. Both tables price each channel at a single epoch, the cloud at formation and the band at the horizon, so neither of them answers whether the bin width they recommend is still admissible in between. That is checked by evaluating eq. (6.11) at eight epochs, with the band at its depth $L(t)$ and count ft, the cloud at the residence density its propagated fragments actually have, and the cross channel

at the overlap of the two, then intersecting the channels that clear the relevance gate. The two layers move in opposite directions and both move as predicted: the band of the large single-shell archetype widens from 39 km at six months to 126 km and then freezes, since the first derelict reenters after 1.1 years and inflow balances outflow from then on, and the measured widths track the analytic staggered-derelict curve of panel (a) of fig. 6.9 over the whole climb, while the fragment cloud spreads and thins, from 220 km and 462 objects at formation to 706 km and 120 at ten years. Because the pair rate scales as N^2/W, that costs the cloud a factor 48 of its contribution over the decade, and the multi-shell cloud a factor 58: the hazard is front-loaded at the breakup, so formation is the epoch worth sizing for.

The intersection is consequently the cloud's own window for most of the horizon, since the cloud never falls below 66 % of the rate proxy in any archetype and the derelict self channel never reaches 10 %, and it widens as the cloud thins rather than closing. Taking the tightest edges over all epochs leaves 4.0 km to 32 km for the large single-shell archetype, 1.3 km to 10 km for the multi-shell one where the shell separation binds from above rather than any channel, and everything up to 37 km for the small polar one. The lattices of table 6.4 sit inside those ranges at every epoch, as does the 36.6 km lattice of the derelict-only case against its own 11 km to 55 km. The high-altitude polar archetype is the exception the cross-species discussion anticipates: its band deepens by only 0.8 km in ten years while the count climbs to 59, and that near-delta is the sparse side of the derelict-fragment channel, which passes the relevance gate after one year and reaches 58 % at ten. Its window is a knife edge near 2.1 km that never intersects the range the cloud admits, so for nine of the ten years no bin width serves the dominant channel and its lattice is reported as a declared fallback. The small polar archetype has an equally undeveloped band and escapes only because that band stays under 5 % of its rate throughout.

6.2.3. Inclination and RAAN Binning

Figure 6.10 refines each lattice axis in turn, holding the other two fixed, for the four archetypes including derelicts and a fragment cloud. In the altitude column the angular lattice is held at $N_i = 6$ and $N_\Omega = 72$; in the two angular ones Δh is held at the per-archetype value of table 6.4.

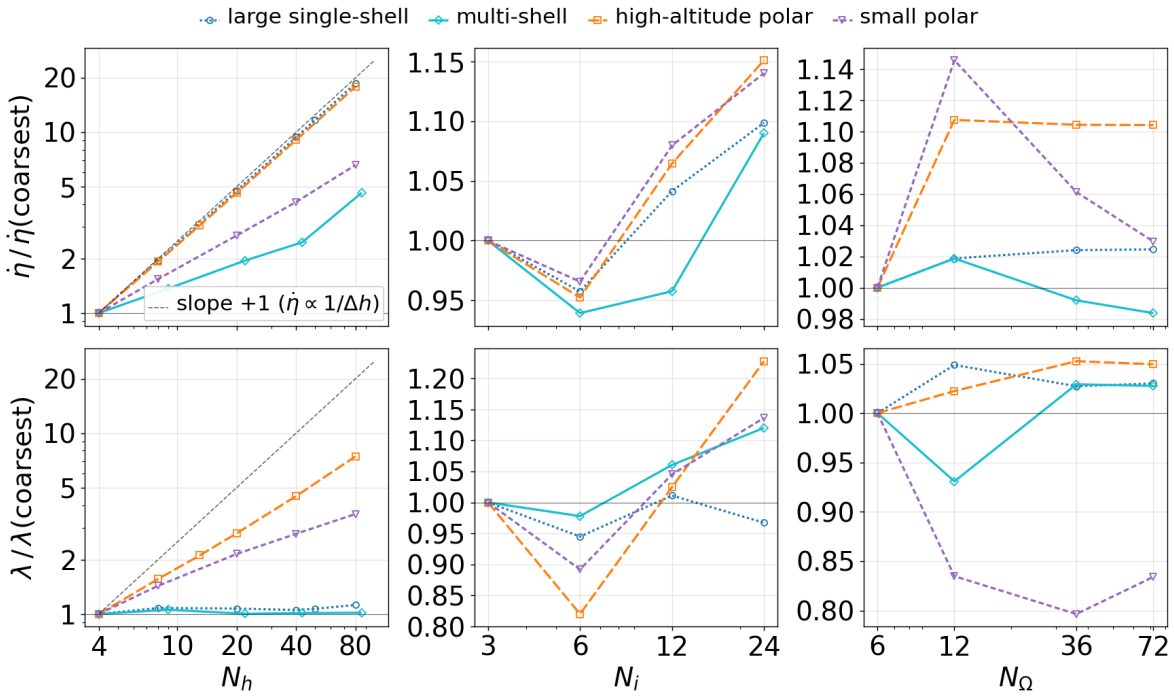

Figure 6.10: Sensitivity of the rate estimate to each lattice axis in turn, for the four archetypes carrying derelicts and a fragment cloud. Top row: the raw rate $\dot{\eta}$; bottom row: the rate realized after avoidance maneuvers, λ.

The contrast with the altitude axis is immediate: over its ladder the raw rate moves by a factor between 4.6 and 18.6 depending on the archetype, while the inclination ladder never moves it by more than 1.21 and the RAAN ladder never by more than 1.15; on the rate realized after avoidance maneuvers the corresponding bounds are 1.50 and 1.26. The angular axes therefore do not diverge, and the first half of this section explains why not. What they do instead is visible in the same figure, and is the subject of its second half: not one of the four angular panels is monotone in the bin count. Every inclination ladder dips below its coarsest value at $N_i = 6$ before rising past it, most sharply for the high-altitude polar archetype, whose realized rate falls to 0.82 and then overshoots to 1.23, and the small polar archetype runs 1.00, 0.84, 0.80 and 0.83 across the four RAAN settings. A refinement that neither converges nor diverges but wanders inside a bounded band is the signature of an error whose sign depends on where the bin centres happen to fall relative to the population, and it is that error, rather than any bin-width scaling, which governs the two angular axes.

That the angular axes do not diverge is not obvious from eq. (6.1) alone: the binned rate between a node pair carries $(\Delta i \Delta\Omega)^{-2}$ exactly as the altitude term carries Δh^{-2}. The crucial difference is what multiplies it: while nodes at different altitude shells carry zero rate by construction, every pair of inclination and RAAN bins within a shell is allowed to interact, with a coupling strength given by the true geometric overlap between the two finite orbital bands, i.e. the exact double integral of the angular overlap over the full $(\Delta i, \Delta\Omega)$ extent of *both* bins. For small bins this overlap therefore scales as the product of both bins' angular areas,

$$\iiint \mathcal{J}(s_\delta)\,\mathrm{d}i\,\mathrm{d}\Omega\,\mathrm{d}u \propto \Delta i_{S_i}\Delta\Omega_{S_i}\Delta i_{S_j}\Delta\Omega_{S_j} = (\Delta i \Delta\Omega)^2$$

Substituting this into the pairwise rate eq. (6.1) shows that the two powers of $\Delta i\,\Delta\Omega$ contributed by the densities are canceled, node by node, by the two equal powers supplied by the angular overlap itself, leaving the pair rate independent of the angular bin size to leading order. What remains after this cancellation is only the error of representing v_{rel} (omitted in the above equation) by the single value that the two bin centres imply, instead of by its true pairwise value, and that error is bounded rather than divergent. The error is then the offset between the true inclination and the centre of whichever bin happens to contain it, bounded by $\Delta i/2$ and therefore first order in the bin width, and both its size and its sign jump whenever a refinement moves the objects from one bin centre to another. This is the origin of the non-monotone behaviour of fig. 6.10: the estimate stays inside a bounded band, but refining the angular lattice does not make it settle within that band. This is possible precisely because the angular overlap is a genuinely pairwise, doubly bin-integrated quantity, unlike the altitude overlap, which is defined on a single shell and can therefore only ever supply one power of Δh against the two powers forced into the density product by the same-shell pairing rule, the origin of the divergence discussed above. Consequently, refining the angular lattice trades a bounded discretization bias for computational cost, while refining the altitude lattice trades a bounded bias for an unbounded one whenever the underlying population is not itself resolved on the scale of Δh.

In each archetype's mixed population, pairs are drawn uniformly among the derelicts and fragments sharing an altitude bin of 20 km up to 20 000 per archetype. A bin is selected with probability proportional to its own pair count and two distinct members uniformly within it, which samples the full pair set uniformly without enumerating it. Each pair's relative speed is then evaluated twice in closed form at its mutual line of nodes. The reference is eq. (4.28) applied to the two objects themselves,

$$v_{rel}^{\mathrm{osc}} = \sqrt{v_1^2 + v_2^2 - 2v_1 v_2 \cos\theta_{12}}$$

with v_1, v_2 and θ_{12} computed from their own osculating elements. This is a per-pair value, not an average over the contents of a bin. Against it stands the speed the rate formulation actually assigns that same pair, which differs from the osculating one by four substitutions: the inclination and the RAAN are each replaced by their bin centre, the eccentricity is set to zero, and the semi-major axis is replaced by the shell radius that both objects then share. With all

four applied the expression collapses to the bin-centre form of eq. (4.30),

$$v_{rel} = 2\sqrt{\frac{\mu_\oplus}{R_\oplus + h^*}}\sin\frac{\theta^*_{12}}{2}$$

The quantity under study is therefore the difference between these two speeds, taken pair by pair. To attribute it to its four causes, all sixteen on and off combinations of the substitutions are evaluated, a complete deterministic factorial running from the osculating speed at one corner to PRISM's own speed at the other.

Figure 6.11 reports the outcome at the campaign baseline angular lattice, $N_i = 6$ and $N_\Omega = 72$, one panel per substitution. What is plotted is a relative error: the part of the difference $v_{rel} - v^{\text{osc}}_{rel}$ that each substitution accounts for, as a fraction of $\langle v^{\text{osc}}_{rel}\rangle$, the population's own mean osculating speed. A violin's width is then the density of that substitution's contribution across the sampled pairs, and the open marker is its mean. A violin centered on zero is noise: positive and negative contributions of similar weight cancel in the total. A violin offset from zero is bias: it shifts every pair the same way and does not cancel. The last violin in each panel is a control population of 4000 objects whose elements are drawn rather than inherited, uniform in inclination, log-uniform in eccentricity and spread over six shells: having no operational orbit for a bin centre to align with, it is the falsification check for every claim that follows.

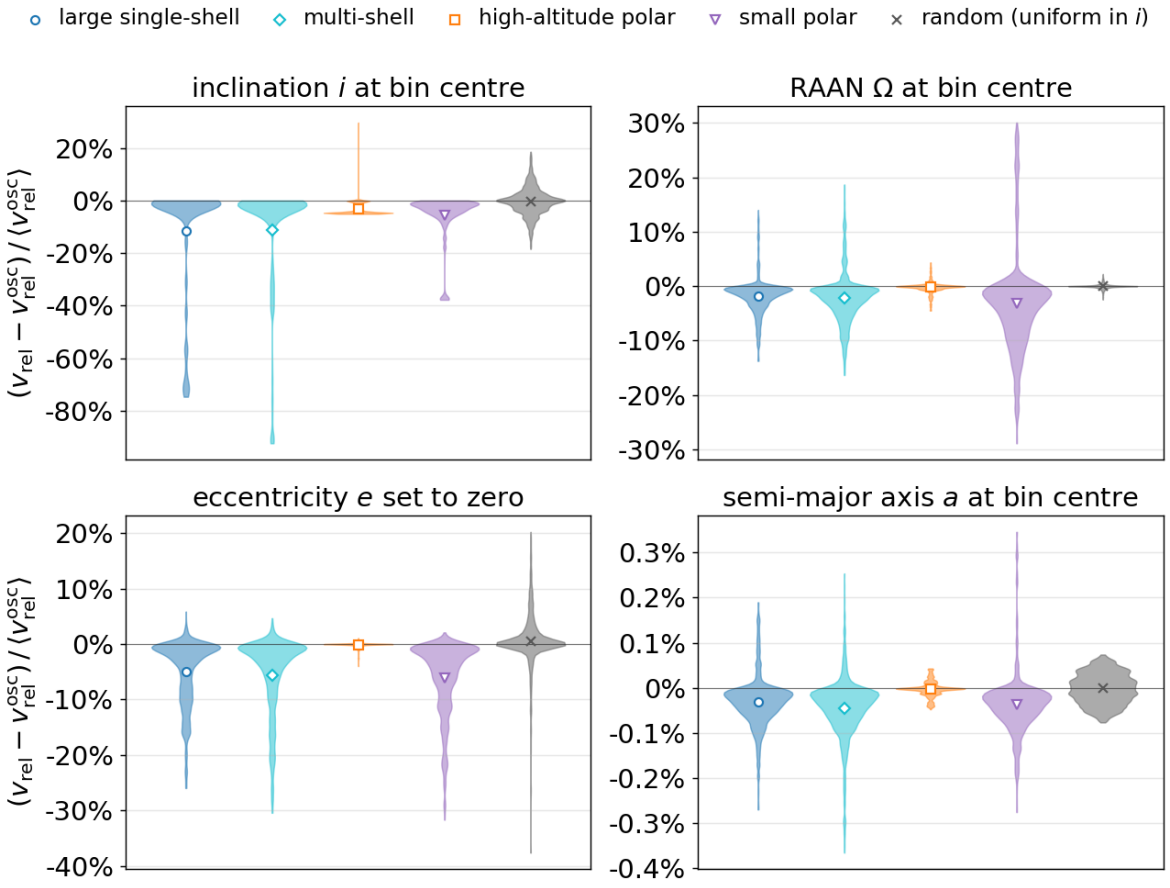

Figure 6.11: Cost of each element discarded by the bin-centre relative velocity formulation, for the four archetypes carrying $A + N + F$ and for a control population uniform in inclination.

Every archetype's violins sit offset from zero, hence it is biased, while the control's sit centered on zero: PRISM is systematically low on the weighted mean relative speed for every real archetype, and only for those. The sign is not an accident of the sample: collapsing two clustered inclinations onto a shared bin centre shrinks the angle between the planes, and zeroing

the eccentricity removes the radial and out-of-plane velocity a real pair carries at a common node, so both substitutions make the modelled encounter more co-orbital than the real one. The mean speed error (the center of the violins) is dominant for eccentricity and inclination, while the RAAN and semi-major-axis substitutions are negligible throughout, and interactions between substitutions stay below 1.9 %, so the four contributions can be read as independent. Additionally, error is not spread evenly across pair types, and which pairs carry it changes what it means. Fragments retain, to within the spread a break-up imparts, the plane of their parent, so a fragment pair generally falls into a single inclination and RAAN bin, at which point the model's own convention that a node has no relative velocity with itself assigns the pair zero speed. For the large single-shell archetype this deletion covers 91.7 % of the sampled pairs, yet it contributes only 7.7 % of the total error, because the deleted pairs are the slow ones: a fragment self-pair moves at 0.19 km/s to 0.24 km/s, against 8 km/s to 12 km/s for a pair involving a derelict, so the rate it carries is small in the same proportion. On the pairs that dominate physical risk, those with at least one derelict, the velocity error is 2.7 % to 11.7 %.

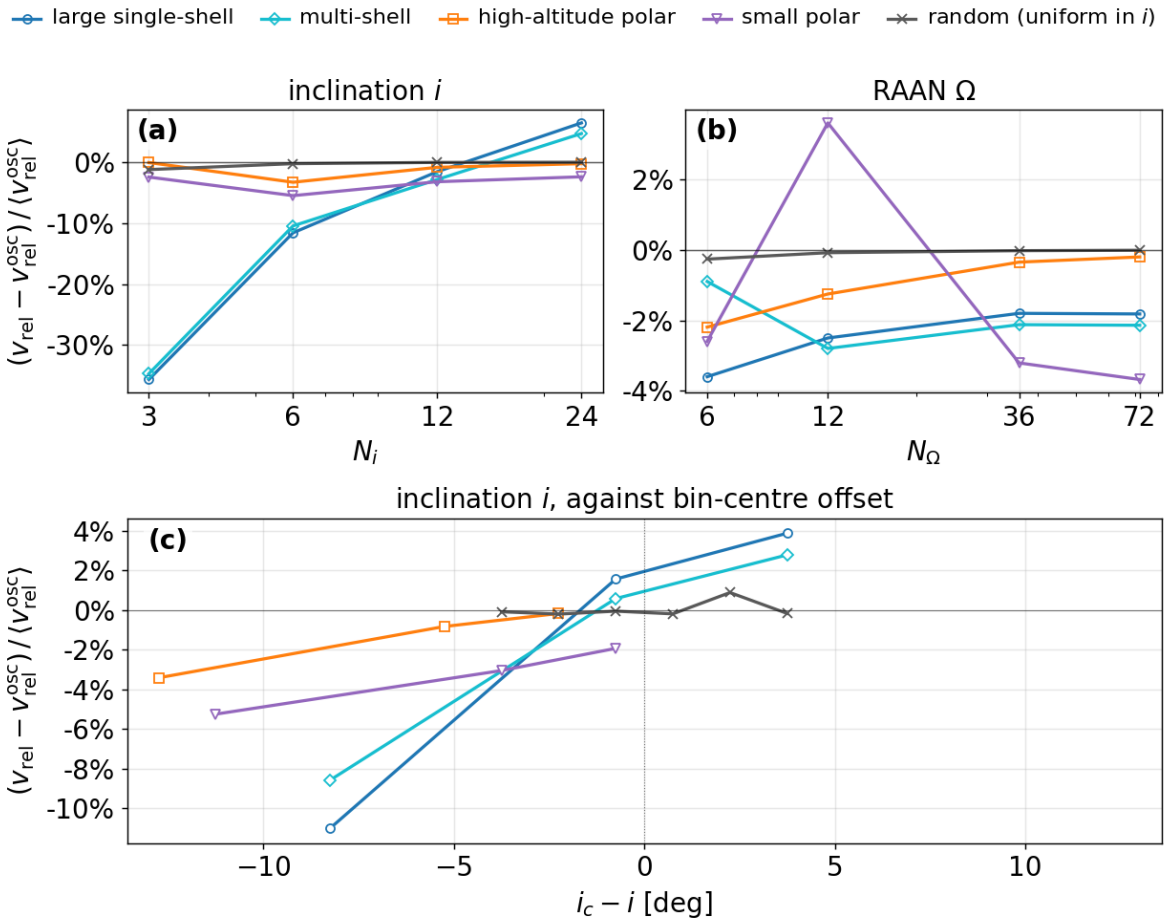

Figure 6.12: The inclination (a) and RAAN (b) contributions against their own bin counts. (c) The inclination contribution against the signed offset between the bin centre and the population's operational inclination.

Figure 6.12 tracks the same contributions as N_i and N_Ω increase, holding the other axis fixed. Nothing on the inclination axis converges. The two mid-inclination archetypes improve steeply at first, from roughly −35 % at $N_i = 3$ to under 3 % at $N_i = 12$, then overshoot to 6.5 % and 4.8 % at $N_i = 24$. The two polar ones stay within a few percent throughout, yet neither settles either: each worsens before it recovers, to −3.2 % and −5.4 % at $N_i = 6$. Every curve turns at least once and none decays monotonically, so a finer inclination lattice is not reliably a more accurate one. The RAAN axis behaves differently. There the contribution never exceeds 3.7 % for any archetype, and refinement neither reliably helps nor hurts: it eases towards zero for the large

single-shell and high-altitude polar archetypes, settles near −2 % for the multi-shell one, and only for the small polar archetype does it swing in sign, from −2.6 % at $N_\Omega = 6$ to 3.6 % at 12 and back to −3.7 % at 72. The control population is the only case that converges monotonically on both axes, having no single operational inclination for a bin centre to sit near or miss. Therefore, the two axes share a common failure to converge: the error is first order in the bin width, and its sign depends on which side of a bin centre the population happens to fall, so refining relocates that offset with no guarantee of shrinking it.

Panel (c) isolates the actual driver: pooling every archetype and inclination bin count onto the signed offset, in degrees, between the bin centre and the population's operational inclination, the inclination contribution becomes very nearly a function of that offset alone, falling monotonically to zero as the offset closes. The control stays flat across the same range, which is the falsification check confirming the effect is placement, not resolution.

This turns the leading term of the angular discretization error from a resolution limit into a design choice. Derelicts and fragments both sit at their parent constellation's operational inclination rather than spread over the axis, so a lattice whose bin edges are placed to land a bin centre on that value removes most of the leading error for free, while a merely finer lattice pays for every bin and accepts whatever offset the arithmetic hands it. The lattices used throughout this campaign are archetype-blind, a uniform division of $[0, \pi]$, which is why the large single-shell archetype sits 8.3° off its nearest bin centre at $N_i = 6$ and pays 11.6 % for it. The alignment is exact for only one inclination at a time, so a multi-shell operator can align to its dominant population and no further, and the residual after alignment is not zero: what remains is the other three substitutions, whose relative weight grows once the inclination term is removed.

Both the dominance of the inclination substitution and its failure to converge under refinement follow directly from the crossing-angle equation, for every population. For a pair of planes at a common inclination i with RAAN separation $\Delta\Omega$ uniform on the circle, eq. (4.28) reduces to

$$\cos\theta = \sin^2 i \cos\Delta\Omega + \cos^2 i,$$

whose sensitivity to an inclination offset is

$$\frac{\partial}{\partial i}\cos\theta = \sin 2i\,(1 - \cos\Delta\Omega),$$

which vanishes at $i = 90°$ and is largest near 45°. This is the entire reason the inclination substitution carries 11.6 % and 10.4 % for the two mid-inclination archetypes and only 3.2 % and 5.4 % for the two polar ones: an inclination offset enters the crossing angle at second order, through $\sin 2i \to 0$, precisely where the archetype sits near the pole. The same equation also explains why refinement fails to settle in fig. 6.12: moving the shell to its inclination bin centre commits an error set by the distance from the true inclination to the nearest centre, which a refinement relocates rather than shrinks monotonically.

The RAAN axis behaves differently for a structural reason. Inclination is nearly a delta for a real constellation, so every object shares the same offset from its bin centre, and alignment removes it for the whole population at once. RAAN, by contrast, enters the crossing angle only through $\Omega_1 - \Omega_2$ on a periodic domain, and a Walker constellation is a uniform comb of planes spanning the full circle: each plane's offset from its bin centre is different, and they largely cancel in the pair sum rather than share a sign. RAAN refinement therefore buys little accuracy, the contribution staying within 3.7 % regardless of N_Ω. What it does buy is resolution in a different sense: attributing risk to individual planes requires each plane to occupy its own bin, so N_Ω should be set by that requirement rather than by the velocity error.

Wang et al. (2025) take the opposite choice and average the crossing angle over the orbital elements. Because the cross term $\sin i_1 \sin i_2 \cos(\Omega_1 - \Omega_2)$ integrates to zero over a uniform RAAN difference, their expected cosine factorizes into a product over the two inclinations alone: RAAN does not enter it at all, and one crossing angle stands for every plane pair in a pair of inclination families. For the large single-shell archetype that single value is 8.57 km/s, in place of a true range of zero to 12.12 km/s across RAAN separation. Both choices are appropriate for what each

was built for: an expectation over the whole population for the global environment, a resolved lattice for a single constellation whose per-plane risk needs attributing.

Both errors ultimately matter only insofar as they move a collision across the catastrophic threshold, $\mathrm{EMR} > 40\,\mathrm{J/g}$. Table 6.5 gives the fraction of the rate the RAAN-resolved lattice misclassifies at this threshold, on the population's own masses. A speed error of nearly a fifth becomes a misclassification of only a twentieth, because most pairs sit well clear of the threshold, and the error is almost entirely false negatives: the lattice under-creates fragment clouds rather than over-creating them, which compounds forward through the simulation instead of merely being conservative. Refinement does not reduce this either, for the same geometric reason as the mean speed error itself.

Table 6.5: Misclassification of catastrophic collisions for PRISM's RAAN-resolved lattice and for the RAAN-collapsed estimator of Wang et al. (2025). The resolved lattice's false positives are below 0.8 % for every archetype.

	RAAN-resolved		RAAN-collapsed		
Archetype	**Misclassified**	**False neg.**	**Misclassified**	**False pos.**	**False neg.**
Large single-shell	3.4 %	3.25 %	78.4 %	75.3 %	3.1 %
Multi-shell	4.5 %	4.04 %	81.2 %	79.4 %	1.8 %
Small polar	2.9 %	2.30 %	29.4 %	22.4 %	7.0 %
High-altitude polar	5.2 %	4.49 %	86.3 %	85.6 %	0.7 %

On the opposite side, table 6.5 shows the uniform-RAAN estimator is an order of magnitude and a half worse in total, and fails almost entirely by false positives, because an estimator returning 8 km/s to 11 km/s for every plane pair of a shell calls catastrophic precisely the near-co-orbital pairs that the lattice merges and calls harmless. On the dangerous direction, false negatives, the two are comparable. The RAAN-resolved estimator's advantage is therefore in total misclassification and in its unique ability to resolve a plane at all in the first place.

In conclusion, the three lattice axes fail in different ways, and only one of them is the binding constraint on PRISM. The altitude axis, section 6.2.1, can diverge without bound whenever the bin width outruns the scale on which the population is actually resolved, moving the realized rate by a factor of several across the campaign's archetypes. The inclination and RAAN axes never diverge, because the geometric overlap between angular bins cancels the same density normalization that drives the altitude divergence; what they cost instead is a bounded discretization bias, dominated by inclination, that reaches the catastrophic-collision classification at only 2.9 % to 5.2 % of the rate, and that a practitioner can shrink further at no computational cost by aligning a bin centre with the population's operational inclination. Altitude resolution is therefore where PRISM's accuracy is won or lost; the angular lattice is a secondary, bounded, and controllable source of error by comparison.

6.3. Time Discretization

The population is frozen for the duration of a step, and the residence weight of eq. (4.33) averages over the orbit with the elements held at their values at the beginning of it. An object coasting under drag therefore does not cross the altitude lattice continuously: it occupies one bin for the whole of a step and moves to another at the boundary. How far it moves is measured by the number of bins it covers in one step,

$$C = \frac{|\dot{a}|\Delta t}{\Delta h} \tag{6.18}$$

which is the Courant number of the transport. Below unity the object visits every bin along its path, and the only error is that the time credited to each is quantized to a multiple of Δt instead of the transit time $\Delta h/|\dot{a}|$, which averages out over where the step grid happens to fall against the bin edges. Above unity the object is recorded in one bin at t_k and in a non-adjacent one

at t_{k+1}, the bins in between are never occupied at any sampled instant, and the rate between the object and whatever is resident there is zero by construction. When the step is too coarse, objects effectively teleport between non-adjacent bins, skipping the shells and their residents entirely. Satellites under thrust are exempt from it, since eq. (4.38) integrates their residence over the step itself and spreads them across every shell they genuinely cross.

Not every channel is exposed to the effect. Figure 6.13 resolves the time step's variance share by species pair, and almost all of it sits on the active-derelict channel, with every other channel at or below a few percent. What separates them is the geometry of the pair and not the species composing it. Where both members are pinned, as they are for two station-keeping satellites, the occupancy product is constant in time and no step size moves it. Where both members drift together, as they do within a derelict band that has reached steady state or within a fragment cloud, the population is translation invariant along the altitude axis, so advancing every member by one bin leaves the number of co-located pairs unchanged and the aliasing displaces the population without altering the count. The active-derelict channel is the only one in which one member is fixed at the shell while the other sweeps through it. Its overlap exists solely over the window a derelict takes to cross the bin holding the active satellites, and a window of finite duration is the only thing a coarse step can miss altogether.

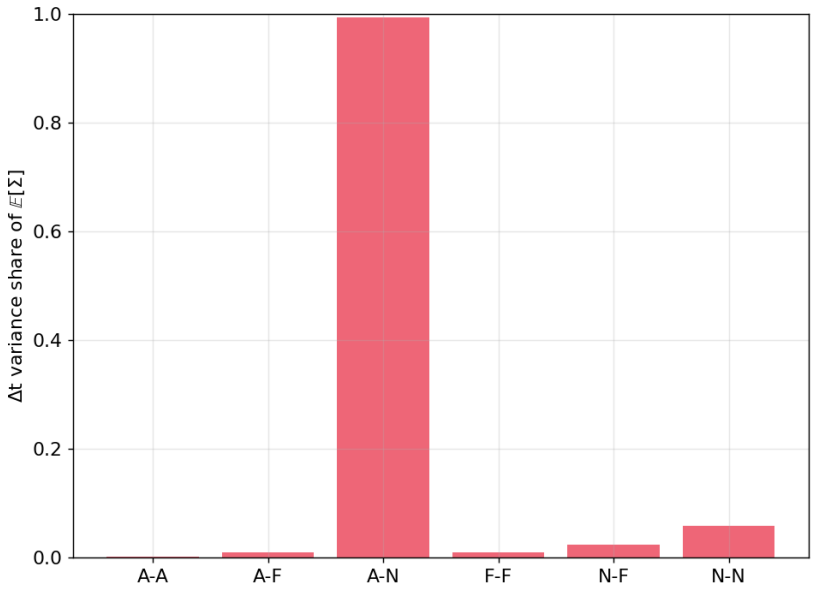

Figure 6.13: Share of the variance of $\mathbb{E}[\Sigma]$ attributable to the time step, resolved by species-pair channel.

Three conditions must therefore hold together before the time step registers at all, and only the multi-shell archetype carrying derelicts alone meets all three. Something has to be moving through the lattice, which excludes both polar archetypes: their decay times of 708 and 4900 years leave their bands deepening by under 3 km over a decade, the same immobility that makes them near-delta layers in section 6.2.1, and their time step share is indistinguishable from zero. A crossing channel has to carry a meaningful share of the rate, which excludes every composition containing a fragment cloud, since the cloud takes 96 % of the realized rate there and leaves the active-derelict channel too small to move the total. And the lattice has to be fine enough for the drift to cover a bin within a step: the multi-shell archetype is held to 8 km by its 10 km shell separation, the finest lattice of any archetype whose derelicts decay appreciably, while the large single-shell archetype sits at 36.6 km at a comparable decay rate. Only that combination brings C near unity anywhere on the ladder from 5 d to 60 d, and it is the combination behind the 36.1 % share reported in fig. 6.2.

The time step is consequently a second-order parameter and remains one. No case moves $\mathbb{E}[\Sigma]$ by more than a factor 1.7 across the whole $(\Delta i, \Delta\Omega, \Delta t)$ cube, against the factors of 5 to 30

the altitude ladder produces on its own, and the single case in which the step leaves a mark is one archetype in one of its two compositions. Since the cost of a run scales inversely with the step, a campaign run at 30 d instead of 5 d loses nothing in almost any scenario. What is worth carrying forward is the direction in which the error grows, which runs opposite to every other axis of this chapter: refining the altitude lattice raises C and makes the time step error worse, not better. A fine altitude lattice combined with a coarse step, on a population with something decaying through it, is the one configuration where the step deserves to be checked.

7

Network Theory

This chapter provides the theoretical foundation to analyze PRISM's population evolution model, built in chapter 4, by recasting it as a network where nodes are orbit bins and edges are interactions between them—collisions, transport, fragmentation. The network structure reveals which regions of orbital space are tightly coupled and which channels carry the risk, insights used in the analysis of chapter 8. Section 7.1 defines the edge weights for these interactions. Section 7.2 constructs adjacency matrices and node-strength centrality measures, which identify dominant pathways. Finally, section 7.3 analyzes long-term behavior through the eigenstructure of a linearized population model.

7.1. Definition of the Network Links

Within the network model, it is possible to define a set of links between nodes that represent the interactions between the different orbit bins and species. These can be *undirected*, such as the collision rate between two nodes, or *directed*, such as the flow of objects from one node to another. This makes the network a *mixed graph*. Furthermore, multiple links can run in parallel between the same two nodes. For instance, two F nodes in adjacent RAAN bins of the same shell can collide, can produce fragments which are deposited in either or both of the two nodes (as well as in other nodes), and can exchange objects through the secular drift of one RAAN bin into the next. This makes the network a *multigraph*. Because the fragments generated by one node can end up in the same node, the network has *self-loops*. Therefore, the network can be formally described as a mixed multigraph with self-loops and weights associated to each link, also called a *mixed weighted pseudograph*.

Two designs were considered to define and compute the links. The first accumulates every pairwise rate at every step of every Monte Carlo realization and averages the resulting adjacency across seeds. It was tried and abandoned. It defeats the FFT aggregate that drives the Poisson collision sampler, since that method works precisely by never materializing individual pair rates; it scales memory with the number of node pairs, almost none of which ever see a realized event; it suffers from the time aliasing issue described in section 6.3, where objects teleport between nodes if the time steps are too large, and it makes the fragment-flow links visible only on the rare steps where a collision is actually sampled, so the quantity being averaged is mostly zero. Furthermore, averaging a nonlinear quantity such as a collision rate or an adjacency matrix across seeds is also not the same as evaluating it on the averaged counts, the difference being the cross-seed covariance, and capturing that correctly requires exactly the per-seed per-step accumulation the design is trying to avoid.

The second design, adopted here, defines every link as a closed-form rate evaluated on a single recorded node state $x_{S_i}(t_k)$, computed only when a network is requested (at a chosen

epoch of a single run, or of an ensemble-averaged node table) and never accumulated during propagation. The simulation writes only the node table, and the edge set of any recorded epoch is rebuilt from it on demand. Two graphs are built per epoch: the undirected graph carrying the collision links, and the directed graph carrying everything else, with parallel typed links on one ordered pair collapsed into a summed weight, plus the per-channel breakdown retained alongside it so that any single channel can be isolated afterwards. The following definitions for the collision, natural transport, fragment redistribution and PMD are used. The transport links of maneuvering objects and the deposition of rising objects in their target orbit are excluded for simplicity.

Collision Links The rate of collision events between nodes S_i and S_j is exactly $\dot{\eta}_{S_iS_j}$ derived in section 2.3, scaled by the CAM failure factor α_{CAM} of eq. (4.3)

$$\chi^{\mathcal{C}}_{S_iS_j} = \alpha_{\text{CAM}}\dot{\eta}_{S_iS_j}$$

This is a function only of the lattice and the node counts, so it can be reconstructed exactly at no additional cost. Setting $\alpha_{\text{CAM}} = 1$ recovers the purely geometric conjunction rate, which is the quantity to use when the question concerns the lattice itself rather than the risk carried by it. The two readings must not be mixed within one comparison, since α_{CAM} suppresses an active against active pair by $(1 - s_{\text{CAM}})^2$ and a fragment against fragment pair not at all, which is enough to reorder which species pairs dominate. Note that the self-node diagonal is exactly zero since two objects in the same orbit bin have zero relative velocity under the bin-center approximation, so $\chi^{\mathcal{C}}_{S_iS_i} = 0$ for every node.

Natural Transport The semi-major axis decays under drag and the RAAN regresses under J_2. Evaluated at the bin centers, the circular-orbit limit of Martinusi et al. (2015) and the classical secular J_2 rate are

$$\dot{a}_\ell = -2C_0\sqrt{\mu_\oplus a^*_\ell} \qquad \dot{\Omega}_{\ell m} = -\frac{3}{2}\sqrt{\frac{\mu_\oplus}{(a^*_\ell)^3}}J_2\left(\frac{R_{eq}}{a^*_\ell}\right)^2\cos i^*_m$$

so the flow of objects out of one node $S_i = (\ell, m, n, S)$ into an adjacent node is

$$\chi^{\mathcal{D}}_{S_iS_j} = \begin{cases} x_{S_i}\dfrac{|\dot{a}_\ell|}{\Delta h_\ell} & \text{if } S_j = (\ell - 1, m, n, S) \text{ and } S \in \{N, F\} \\ x_{S_i}\dfrac{|\dot{\Omega}_{\ell m}|}{\Delta\Omega_n} & \text{if } S_j = (\ell, m, n', S) \text{ and } S \in \{N, F\} \\ 0 & \text{otherwise} \end{cases}$$

where $n' = \left(n + \operatorname{sgn}\dot{\Omega}_{\ell m}\right) \bmod N_\Omega$ is the next RAAN bin in the direction of motion. The drag term $C_0 = B\rho/2$ is the same as the one in eq. (4.32), with B the ballistic coefficient. However, since the individual object trajectories are unknown a posteriori, the ballistic coefficient must be substituted by a representative per-species value and the density by its value at the shell-center altitude.

Only the N and F species carry these links. An active satellite station-keeps, so neither its altitude nor its RAAN is free to evolve, while the rising and de-orbiting species do drift but are dominated by the thrust that defines them, which is the transport excluded above.

The two entries are independent links rather than one diagonal link into $(\ell - 1, m, n')$. A bin has an altitude face and a RAAN face, and an object leaving the bin crosses exactly one of the two: the outflux integrated over the whole boundary is

$$\chi^{\mathcal{D}}_{S_iS_j} = x_{S_i}\left(\frac{|\dot{a}_\ell|}{\Delta h_\ell} + \frac{|\dot{\Omega}_{\ell m}|}{\Delta\Omega_n}\right) \qquad \text{with } S_j = (\ell - 1, m, n', S)$$

which is exactly the sum of the two entries. The boundaries of the lattice are treated accordingly: the lowest shell has no link downwards, since an object leaving it has re-entered and left the network, and the unbounded top shell has no center altitude at which either rate could be evaluated. The RAAN axis has no boundary at all, being periodic, and its bins are uniform by construction so that a single width divides the drift rate everywhere.

For the same reason that the ballistic coefficient has to be a representative, the natural transport law is a bin-center closure that holds no information regarding the distribution of objects within a bin. The occupancy x_{S_i} is eccentricity-resolved but the flux between occupancies is not: an object with its perigee inside the shell and its apogee far above decays much faster than the shell-center circular rate says, and a residence-weighted count cannot express the difference. The transport is therefore a mean-field approximation. On the other hand, the rates above contain no time step, so unlike an accumulated flux they can neither teleport objects across non-adjacent bins nor suppress the instantaneous flow when the step is coarse, the failure described in section 6.3. How coarse the step is against the drift timescale remains a statement about the time resolution of the simulation and not about the link.

PMD and EOL Transitions When active objects reach EOL they transition to an N node or a D node in the same orbit site, in the proportions γ and $1-\gamma$ respectively, where γ is the PMD failure rate. The links from an active node A_i to its corresponding non-active and de-orbiting nodes are therefore

$$\chi^{\text{PMD}}_{A_i N_i} = \gamma \frac{x_{A_i}}{T_{\text{EOL}}} \qquad \chi^{\text{EOL}}_{A_i D_i} = (1-\gamma)\frac{x_{A_i}}{T_{\text{EOL}}}$$

the second of which exists only in those scenarios where a disposal maneuver is modeled at all. This formulation treats every active object as retiring at the fleet-average rate $1/T_{\text{EOL}}$, irrespective of its actual time on orbit, which is the mean-field approximation of the PMD stochastic process.

Fragment Flow and Non-Catastrophic Collisions The flow of fragments from a collision between two nodes S_i and S_k into a fragment node F_j is

$$\chi_{S_i F_j} = \sum_k \sum_{S_k \in \mathcal{S}} \alpha_{\text{CAM}} \dot{\eta}_{S_i S_k} \zeta_{S_i S_k} \xi^{F_j}_{S_i S_k}$$

where the inner sum runs over the species and the outer over the orbit bins, so together over every node of the network, although only the nodes sharing the shell of S_i contribute since $\dot{\eta}$ vanishes otherwise. The three factors are distinct: $\dot{\eta}_{S_i S_k}$ is how often the pair collides, $\zeta_{S_i S_k}$ is how many tracked fragments one such collision makes, taken from eq. (4.42), and $\xi^{F_j}_{S_i S_k}$ is the fraction of those fragments that are deposited in F_j. The link is therefore a fragment rate and not a collision rate, in fragments per unit time, and it is the product of a frequency and a per-event yield. The same production appears on the link out of S_k, so a breakup is credited in full to both of its parents rather than split between them, and when S_i is itself a fragment node the destination can be S_i, which is where the self-loops of the network come from.

The fraction ξ is approximated by a constant once the breakup site is fixed (Wang et al., 2025), and the reason lies in the NASA SBM formulation. The number of fragments above the tracked size threshold depends on the mass and the energy of the collision, whereas the distribution of their ejection velocities depends on each fragment's characteristic length and area-to-mass ratio and not on the masses of the parents. Scaling the parents therefore scales the count and leaves the normalized distribution of destinations unchanged, so ξ is a property of the site rather than of the event. Two properties of the population make it cheap to tabulate as well. The parents are near-circular, and for a circular orbit the radius and the speed are constant around the orbit, so the true anomaly at which the collision occurs does not enter the post-kick distribution of semi-major axes; and the ejecta are isotropic, so no direction of the kick is privileged either. What remains is a function of the breakup site alone, which can be sampled once offline and reused for every collision. The notation $\xi^{F_j}_{S_i S_k}$ retains the indices of both parents, but the quantity tabulated here depends on the species of neither: they enter only through the fragment count.

The table is built by the recipe of Wang et al. (2025), which is to initialize the breakup model at every source bin and histogram where the fragments land. A source sits at the shell-center altitude and the inclination-bin center, circular, with the true anomaly drawn uniformly, and the fragments it produces are binned through the same Keplerian residence weights of eq. (4.33) that define the node occupancy. The draw is repeated and averaged, so the stored fractions are expectations rather than one realization of a stochastic breakup, and in that sense the fragment flow is a mean-field approximation of the redistribution process in the same way the transport links are of the transport. Destinations receiving less than a fraction 10^{-4} of the fragments produced by a given source are discarded as single-fragment sampling noise. The surviving destinations keep their sampled fractions and are not rescaled to recover the discarded mass, since doing so would inflate every retained destination.

The table is keyed by a pair and returns a triple. A row is the shell and inclination bin of the heavier of the two parents, and a destination cell is a shell, an inclination bin, and a signed offset in RAAN bins relative to the parent's own, from which the absolute destination is recovered at use time as $n_j = (n_t + \Delta n) \bmod N_\Omega$. J_2, a spherically symmetric atmosphere and an isotropic ejecta velocity are all axisymmetric about the polar axis, so rotating the parent orbit about that axis by an arbitrary angle shifts every fragment's RAAN by the same angle and changes nothing else. The distribution of the offset therefore carries no residual dependence on the absolute RAAN of the breakup, and a RAAN row index would have nothing left to resolve. Collapsing the RAAN axis onto an offset removes a factor of N_Ω from both the storage and the sampling budget, and it improves the estimate since an absolute RAAN axis would spend the same number of samples on N_Ω independent noisy copies of one distribution instead of one well sampled copy.

Rows sum to at most one, the deficit being the fragments that re-enter, that escape, and those in the discarded noise cells above. There is no angular analogue of that deficit, because the inclination bins span $[0, \pi]$ and the RAAN bins span the full circle by construction, so a full row of the joint table sums to exactly the row sum of its altitude marginal. Keying the row on the heavier parent is an approximation in its own right, since two parents share a shell but not in general an inclination bin, and the cloud of the lighter one is attributed to the plane of the heavier. Otherwise this is identical to the fragment flow model of Wang et al. (2025) with one axis added, the RAAN axis.

Finally, in Wang et al. (2025), their NESSY model accounts for collisions with small fragments under 10 cm by asserting that collisions with small objects occur at a rate approximately 5.3 times the rate of collisions with large objects, and are carried by interactions with every F node, demoting an active satellite to a non-active one without producing any new fragments. Here, fragments below the 10 cm threshold are not modeled at all, so that channel is absent entirely. Non-catastrophic collisions between tracked objects are modeled, and they are a different mechanism producing the same outcome: when the energy-to-mass ratio of a collision is too low to fragment the heavier of the two objects, that object is demoted from active to non-active, with a rate

$$\chi^{\mathrm{SC}}_{S_i N_i} = \sum_k \sum_{S_k \in \mathcal{S}} \alpha_{\mathrm{CAM}} \dot{\eta}_{S_i S_k} \qquad \text{with } S_i \in \mathcal{A},\ m_{S_i} \geq m_{S_k},\ \mathrm{EMR}_{S_i S_k} < 40\,\mathrm{J/g}$$

where $\mathrm{EMR}_{S_i S_k} = m_{S_k} v_{rel}^2 / 2 m_{S_i}$ is the kinetic energy of the impactor per unit mass of the target, evaluated with the representative masses of the two species and the bin-center relative velocity of eq. (4.28). Nothing in the definition restricts the species of the impactor, but in practice only fragment nodes hold sufficiently small objects: between two intact objects the mass ratio is of order one, and at orbital relative velocities the EMR then exceeds the threshold by more than two orders of magnitude. Unlike the small-debris channel it stands in for, this one does produce fragments, from the Krisko scaling of the impactor's kinetic energy rather than from the total mass of the pair, so the demotion accompanies a fragment flow instead of replacing it.

7.2. Network Topology

The link definitions of the previous section assign a rate to every node pair from the node counts alone, so the whole edge set of any recorded epoch can be rebuilt a posteriori without the simulation having stored a single link. This section defines what that edge set carries: the adjacency matrices that hold the links, and the per-node strengths by which nodes are ranked against each other.

7.2.1. Adjacency Matrices

An adjacency matrix represents a set of links as a square array, one row and one column per node, whose entry at row S_i and column S_j holds the weight of the link from S_i to S_j, and zero wherever no such link exists. Each node here is the pair of a flattened orbit-bin index and a specie, so the network has $5J$ nodes and its adjacency array is $5J \times 5J$. Two arrays are kept rather than one. The collision adjacency matrix $\Xi^{\mathcal{C}}$ collects the collision links,

$$\Xi^{\mathcal{C}}_{S_i S_j} = \chi^{\mathcal{C}}_{S_i S_j} = \Xi^{\mathcal{C}}_{S_j S_i}$$

and is symmetric, because a collision links the unordered pair $\{S_i, S_j\}$ and has no direction. Its diagonal is formed but vanishes identically. The entry $\Xi^{\mathcal{C}}_{S_i S_i}$ would be the rate of collisions between two objects of the same specie sharing the same orbit bin, and the pair is enumerated, but the relative-velocity term $2\sin(\theta/2)$ is zero at equal inclination and zero RAAN separation. Two objects in the same bin are, at this resolution, co-orbital and never close on each other. Within-bin conjunctions are therefore not a channel of this network, and the diagonal is empty in every result reported here. The matrix is also sparse: two objects in different altitude shells never share a point of space in this formulation, so $\Xi^{\mathcal{C}}$ is block diagonal in the altitude index, with one block per shell. Within a block it is dense, since every inclination and RAAN bin of a shell can in principle intersect every other.

The flow adjacency matrix $\Xi^{\mathcal{T}}$ collects everything else, that is, every directed transition a single object can undergo: natural drag decay across an altitude boundary and J_2-induced RAAN drift across a RAAN boundary, the $A \to N$ transition from failed PMD, the $A \to N$ demotion from a non-catastrophic collision, and the post-collision fragment redistribution $\chi_{S_i F_j}$. Several of these can act on the same ordered pair at once, so the entry

$$\Xi^{\mathcal{T}}_{S_i S_j} = \sum_{\tau} \chi^{(\tau)}_{S_i S_j}$$

is the sum over the channels τ, with the per-channel breakdown retained alongside it so that any single channel can be isolated afterwards. Keeping the two matrices apart is a modeling choice, not a bookkeeping one. Collision coupling and directed mass transport answer different questions: the first is a symmetric statement about which parts of the lattice are dangerous to each other, the second a conservation statement about where objects come from and where they end up. Collapsing them into one signed array would make both harder to read, and would put an undirected rate and a directed rate in the same sum.

7.2.2. Node Strength

On an unweighted graph, a node's degree is the number of links attached to it. Every link here carries a rate rather than a plain count, so the natural per-node measure is the strength instead: the sum of the weights of the links attached to a node, in place of their number. On the collision side,

$$s^{\mathcal{C}}_{S_i} = \sum_{j} \sum_{S_j \in \mathcal{S}} \Xi^{\mathcal{C}}_{S_i S_j} \tag{7.1}$$

is the total collision rate the node participates in, and is the primary frequency-based ranking of nodes as risk drivers. The relative velocity is already folded into each link through the overlap integral of section 4.2, so no further velocity weighting is applied on top: the strength is a rate,

in events per unit time, and not a score. One caution follows from the definition. Summing $s^{\mathcal{C}}$ over the nodes does not give the population's total collision rate, because every link is counted once at each of its two endpoints. The strength is therefore a ranking, and whenever a share of a total is wanted instead, the attribution has to be by link rate.

Two directed strengths follow from $\Xi^{\mathcal{T}}$,

$$s^{\text{out}}_{S_i} = \sum_{S_j} \Xi^{\mathcal{T}}_{S_i S_j} \,, \qquad s^{\text{in}}_{S_i} = \sum_{S_j} \Xi^{\mathcal{T}}_{S_j S_i}$$

whose difference $s^{\text{out}}_{S_i} - s^{\text{in}}_{S_i}$ separates sources from sinks: positive marks a node that emits objects faster than it receives them, as an operational shell does when its first generation reaches end of life, negative a node that accumulates them, as the low shells do where decaying objects pile up before re-entry. Restricting the out-strength to the fragment-flow channel alone gives

$$s^{F}_{S_i} = \sum_{j} \chi_{S_i F_j} \tag{7.2}$$

the expected fragment production rate of the node, in fragments per unit time, which ranks nodes by how much debris their collisions actually make rather than by how often they collide.

The two collision-side strengths are not independent quantities, and the relation between them is the one a risk assessment usually writes as likelihood times consequence. Every term of s^F carries the same factor $\dot{\eta}_{S_i S_k}$ that builds $s^{\mathcal{C}}$, multiplied by the per-event yield $\zeta\xi$, so s^F is a collision rate already weighted by the debris each of those collisions produces. The collision strength is the likelihood term alone, counting events in collisions per unit time and knowing nothing of their outcome. The ratio $s^F/s^{\mathcal{C}}$ is the consequence term alone, a mean fragment count per collision at that node, carrying no information about how often the collision occurs. Their product is s^F itself, which is therefore the combined risk and the only one of the three whose units are those of a hazard accumulating in time. Note also that α_{CAM} enters both strengths identically, so it scales the likelihood and the risk but cancels from the ratio, leaving the consequence term a statement about the breakup physics alone regardless of the avoidance assumption.

The severity so measured is a fragment count at the node where the collision happens, and says nothing about the mass or the lethality of those fragments nor about where they end up. A high yield is also not by itself a reason to act on a node, if the frequency multiplying it is negligible, which is why the product and not either factor alone is the quantity to rank on.

Centrality measures rank nodes by their position in the graph's structure rather than by the weight of their links, and the classical ones are unweighted: degree itself, betweenness, how often a node lies on the shortest path between two others, and closeness, the inverse of a node's average shortest-path distance to every other node. None of them are used here. Within one altitude shell the collision graph is close to complete by construction, since the lattice is a geometric partition and not an empirically sparse network, so the unweighted degree of a node there is essentially the number of occupied bins in its shell and says nothing about that node. Betweenness and closeness fail on $\Xi^{\mathcal{C}}$ for the same reason, since its blocks have a shortest-path diameter of one or two regardless of the physics. All of the information sits in the weights, which is why every measure defined here is rate-weighted instead.

7.3. Network Dynamics

Section 4.1 derived the network's evolutionary equations in their discrete-time, stochastic form, eq. (4.1), where $\boldsymbol{Y}(\boldsymbol{X}_k)$ collects every Poisson jump process sampled on the links of section 7.1. Taking the mean-field limit of that recursion, that is, replacing each sampled jump by its expected rate and letting the step Δt_k shrink, turns the sampled increment $\boldsymbol{Y}(\boldsymbol{X}_k) + \boldsymbol{g}(t_k)$ per step into a deterministic flux per unit time, and eq. (4.1) becomes the continuous-time ordinary differential equation

$$\dot{\boldsymbol{X}} = \boldsymbol{L}(\boldsymbol{X}) + \boldsymbol{B}(\boldsymbol{X}, \boldsymbol{X}) + \boldsymbol{g} \tag{7.3}$$

where $\boldsymbol{X}$ denotes the expected node populations, in place of the sampled $\boldsymbol{X}_k$. This is the same approximation already applied link by link in section 7.1, to the natural transport rates, the PMD transition, and the fragment redistribution table; here it is applied to the system as a whole, in place of the stochastic simulator, and it yields a deterministic system of the same dimension.

The reduction is needed because the questions this chapter asks of the network are not questions a single stochastic realization can answer. One realization is one trajectory, sampled once from the jump process of section 4.1: it has no equilibrium and no stability margin, only a sample path. Asking whether the network settles into a steady population, and whether that state is robust to a small perturbation, requires a deterministic system with well-defined fixed points and a Jacobian to linearize around, which eq. (7.3) supplies and eq. (4.1) does not.

The two terms of eq. (7.3) are the continuous-state evaluations of the two adjacency matrices of section 7.2.1, and their degree in $\boldsymbol{X}$ follows from the physical rate each matrix carries. $\boldsymbol{L}$ is linear because every link of the flow adjacency matrix $\Xi^{\mathcal{T}}$ is a rate proportional to a single node's occupancy, section 7.1: its S_i-th entry is

$$L_{S_i}(\boldsymbol{X}) = \sum_{S_j} \Xi^{\mathcal{T}}_{S_j S_i}(\boldsymbol{X}) - \sum_{S_j} \Xi^{\mathcal{T}}_{S_i S_j}(\boldsymbol{X})$$

the inflow into S_i minus its outflow, the same balance that separates sources from sinks in section 7.2.2, signed the other way and evaluated at a general population instead of a single recorded epoch. $\boldsymbol{B}$ is bilinear because every link of the collision adjacency matrix $\Xi^{\mathcal{C}}$ is a rate proportional to the product of both endpoints' occupancies, eq. (4.10), so it is a map $\boldsymbol{B}(\boldsymbol{X}_1, \boldsymbol{X}_2)$ linear in each argument separately: $\Xi^{\mathcal{C}}$ removes population from the colliding nodes and $\chi_{S_i F_j}$ deposits fragments elsewhere, and evaluating both arguments at the same $\boldsymbol{X}$ recovers the physical collision term. Keeping $\boldsymbol{L}$ and $\boldsymbol{B}$ as separate operators, rather than combining them into one, is therefore not a bookkeeping choice: a linear term and a bilinear term cannot be added into a single linear system, since they scale differently with $\boldsymbol{X}$.

An equilibrium point $\boldsymbol{X}^*$ of eq. (7.3) satisfies $\boldsymbol{L}(\boldsymbol{X}^*) + \boldsymbol{B}(\boldsymbol{X}^*, \boldsymbol{X}^*) + \boldsymbol{g} = \boldsymbol{0}$. To linearize around it, write $\boldsymbol{X} = \boldsymbol{X}^* + \delta\boldsymbol{X}$ and substitute into eq. (7.3):

$$\begin{aligned}\delta\dot{\boldsymbol{X}} &= \boldsymbol{L}(\boldsymbol{X}^* + \delta\boldsymbol{X}) + \boldsymbol{B}(\boldsymbol{X}^* + \delta\boldsymbol{X}, \boldsymbol{X}^* + \delta\boldsymbol{X}) + \boldsymbol{g} \\ &= \underbrace{\boldsymbol{L}(\boldsymbol{X}^*) + \boldsymbol{B}(\boldsymbol{X}^*, \boldsymbol{X}^*) + \boldsymbol{g}}_{=\boldsymbol{0}} + \boldsymbol{L}(\delta\boldsymbol{X}) + \boldsymbol{B}(\boldsymbol{X}^*, \delta\boldsymbol{X}) + \boldsymbol{B}(\delta\boldsymbol{X}, \boldsymbol{X}^*) + \boldsymbol{B}(\delta\boldsymbol{X}, \delta\boldsymbol{X})\end{aligned}$$

using the linearity of $\boldsymbol{L}$ and the bilinearity of $\boldsymbol{B}$ to expand every term exactly, with no approximation yet. The first three terms of the second line cancel by the equilibrium condition. Because $\Xi^{\mathcal{C}}$ is symmetric, $\boldsymbol{B}$ is symmetric in its two arguments, so $\boldsymbol{B}(\boldsymbol{X}^*, \delta\boldsymbol{X}) + \boldsymbol{B}(\delta\boldsymbol{X}, \boldsymbol{X}^*) = 2\boldsymbol{B}(\boldsymbol{X}^*, \delta\boldsymbol{X})$, leaving

$$\delta\dot{\boldsymbol{X}} = \underbrace{\boldsymbol{L}(\delta\boldsymbol{X}) + 2\boldsymbol{B}(\boldsymbol{X}^*, \delta\boldsymbol{X})}_{\text{first order in } \delta\boldsymbol{X}} + \underbrace{\boldsymbol{B}(\delta\boldsymbol{X}, \delta\boldsymbol{X})}_{\text{second order in } \delta\boldsymbol{X}}$$

which is the exact Taylor expansion of the right-hand side of eq. (7.3) about $\boldsymbol{X}^*$, since $\boldsymbol{L}$ and $\boldsymbol{B}$ are polynomial in $\boldsymbol{X}$ and the expansion terminates at second order. Dropping the second-order term for a small enough perturbation gives the linearized dynamics $\delta\dot{\boldsymbol{X}} \approx \boldsymbol{J}\,\delta\boldsymbol{X}$, with the Jacobian

$$\boldsymbol{J} = \boldsymbol{L} + 2\boldsymbol{B}(\boldsymbol{X}^*, \cdot)$$

acting on $\delta\boldsymbol{X}$, and the sign of its spectral abscissa, the largest real part among its eigenvalues, decides whether $\boldsymbol{X}^*$ is stable. This Jacobian-and-equilibrium framework is the general tool of which section 8.3.1's carrying-capacity construction is a reduced instance. It uses the same operators $\boldsymbol{L}$ and $\boldsymbol{B}$, and so weighs the same fragment production against the same removal, but on a coarser state and in a form where that balance can be evaluated directly, without ever solving for $\boldsymbol{X}^*$ or forming its Jacobian.

8

Results

The previous chapters discussed the formulation and implementation of a network model for the long-term evolution of a satellite constellation. Here, PRISM is applied to three real scenarios, namely Starlink, OneWeb and Iridium-NEXT, to assess their associated self-induced risks and their long-term stability. Section 8.1 traces the baseline evolution of each constellation. Section 8.2 examines how the orbital plane structure concentrates or disperses the collision rate. Section 8.3 introduces the carrying-capacity indicator ϱ_{crit}, which measures the margin before a shell stops clearing its own debris. Section 8.4 then sweeps three policy variables — avoidance performance, disposal reliability, and launch strategy — to assess each constellation's sensitivity to these levers.

8.1. Baseline Constellation Evolution

PRISM is applied to three real mega-constellations, built from their DISCOS data as of July 2026: Starlink, OneWeb, and Iridium-NEXT. Each is propagated on a time horizon that covers at least one full replacement cycle past the end of deployment. Table 8.1 reports the baseline parameters shared by all three scenarios. In all three scenarios, a new satellite is launched every time an active one reaches end-of-life, so as to maintain a constant population after the ramp-up phase.

Table 8.1: Baseline simulation parameters, shared by all three scenarios.

Parameter		**Value**
PMD failure	γ	5%
CAM success	s_{CAM}	99.99%
Life-time of active satellites	T_{EOL}	5 years
Time-step	Δt	5 d
Time horizon	t_f	25 years
Thrust acceleration	ϵ	5×10^{-8} km/s^2
Launch strategy		one-for-one

Table 8.2 summarizes the initial population for each: injection semi-major axis a_{inj}, target semi-major axis a_{tgt}, eccentricity e, inclination i, launch year and month, number of planes, and number of satellites.

Table 8.2: Injection and target orbit per shell for the three constellations, from DISCOS.

Shell	a_{inj} [km]	a_{tgt} [km]	e	i [deg]	**Launch date**	**N. planes**	**N. sats**
Starlink							
Gen2-group12	6666	6720	0.0001	53.0	2024-11	28	560
Gen2-group9	6664	6744	0.0001	53.0	2024-06	28	252
Gen2-group6	6664	6828	0.0001	43.0	2023-02	96	2496
Gen2-group10	6645	6830	0.0001	53.0	2024-06	10	1400
Gen2-group11	6641	6830	0.0001	53.0	2024-08	10	1400
Gen2-group8	6671	6848	0.0001	53.0	2024-04	28	198
Gen2-group7	6671	6860	0.0001	53.0	2023-08	72	360
Gen2-group15	6655	6878	0.0001	70.0	2025-02	10	1400
Gen1-group4	6663	6918	0.0013	53.2	2021-11	72	1512
Gen2-group13	6691	6928	0.0001	97.6	2025-01	10	1400
Gen2-group17	6651	6928	0.0001	97.6	2025-05	10	1400
Gen1-group1	6648	6928	0.0013	53.0	2019-05	72	1584
Gen2-group5	6695	6937	0.0010	53.0	2022-12	28	672
Gen1-group3	6695	6938	0.0013	97.6	2021-01	6	228
Gen1-group2	6651	6948	0.0009	70.0	2021-09	18	396
OneWeb							
–	6831	7578	0.0001	87.9	2019-02	12	588
Iridium-NEXT							
–	6994	7156	0.0001	86.4	2017-01	6	66

8.1.1. Plain Monte Carlo

Figures 8.1 to 8.3 show each constellation's fleet composition over its horizon, one row per populated species (A, R, D, N, F, as defined in section 4.1), the 200-seed plain Monte Carlo ensemble on the left and the exact conditional branch of section 5.2.3 on the right, both as the across-seed mean. Only the left column carries a standard deviation band: the conditioning writer stores a per-node variance, and summing that over the nodes of a species would drop the cross-node covariance, so no band is drawn on the right rather than a wrong one.

The right column is the conditional expectation $\mathbb{E}[\cdot \mid \Sigma \geq 1]$, the branch in which a collision has occurred, and it is not weighted by how likely that branch is. Its fragment row is therefore the $\mathbb{E}[x_F \mid \Sigma \geq 1]$ of table 8.3 resolved in time, and not the constellation's expected fragment count. The unconditional expectation follows from it by the single scalar factor $P(\Sigma \geq 1)$ of eq. (5.1), exactly and at every epoch, because the collision-free branch contributes no fragments at all. Plotting it alongside would therefore reproduce the same curve rescaled, with no difference in shape and none in timing, which is why only the conditional branch is drawn and the weight is reported as a number instead.

Starlink ramps up over roughly its first five years and holds its active fleet flat thereafter, around 10 500 satellites, the same one-for-one closed-loop replacement as OneWeb and Iridium-NEXT below. The derelict channel N opens around year three, once the first satellites fail PMD at end of life, and settles to roughly 150–175 objects for the remainder of the horizon. The fragment channel F is populated, but only barely and only late: it first appears around year seven, peaks under six objects, and decays back toward one by the horizon end, with a standard deviation band far wider than the mean itself. That is consistent with the constellation's own expected collision count over the horizon, $\Lambda = 1.647 \times 10^{-2}$ (table 8.3 below): across 200 seeds only three or four collisions are expected in total, so what fig. 8.1 shows in F is a handful of seeds that happened to collide, not a converged estimate of the channel.

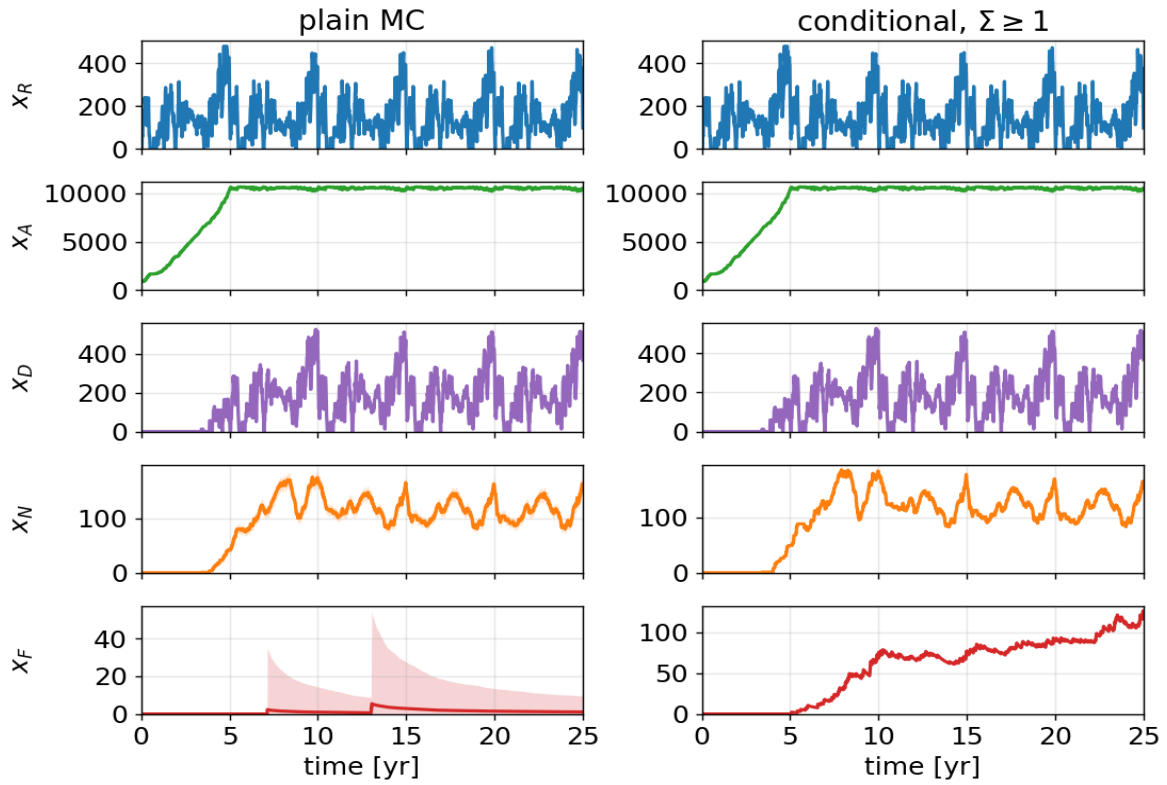

Figure 8.1: Starlink population evolution: plain Monte Carlo ensemble (left) versus conditional branch (right).

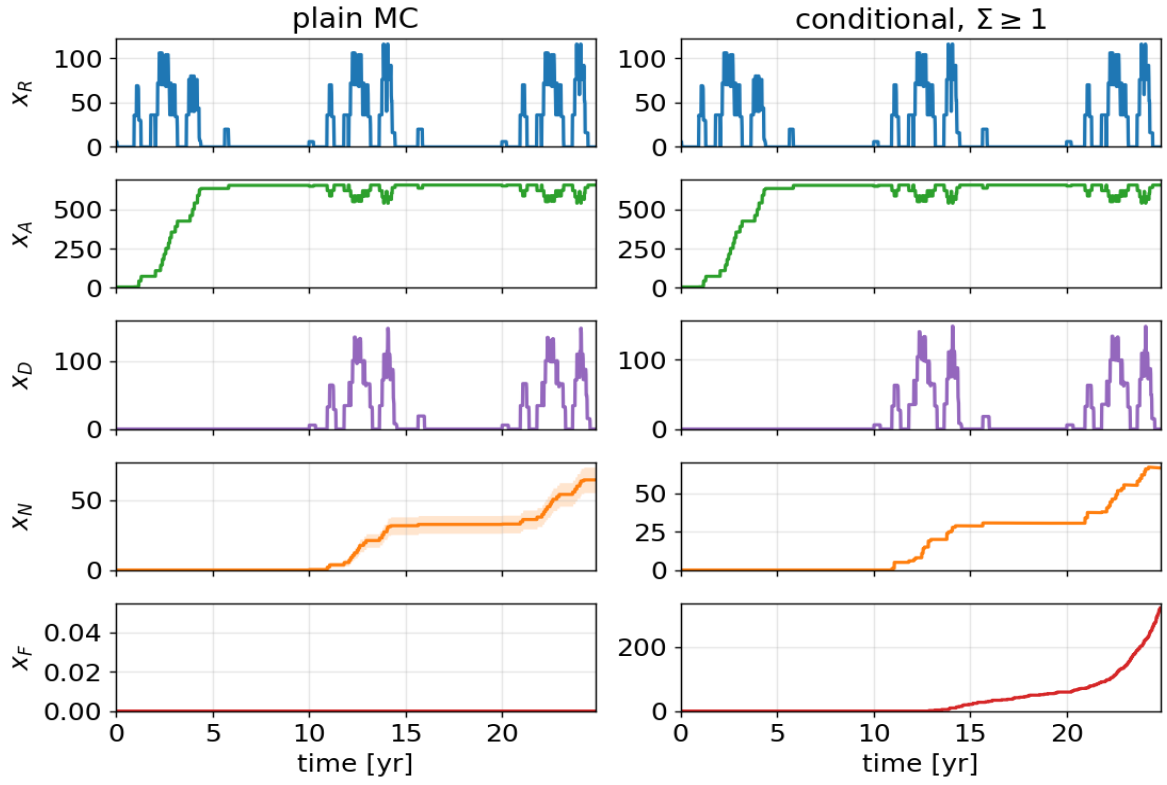

Figure 8.2: OneWeb population evolution: plain Monte Carlo ensemble (left) versus conditional branch (right).

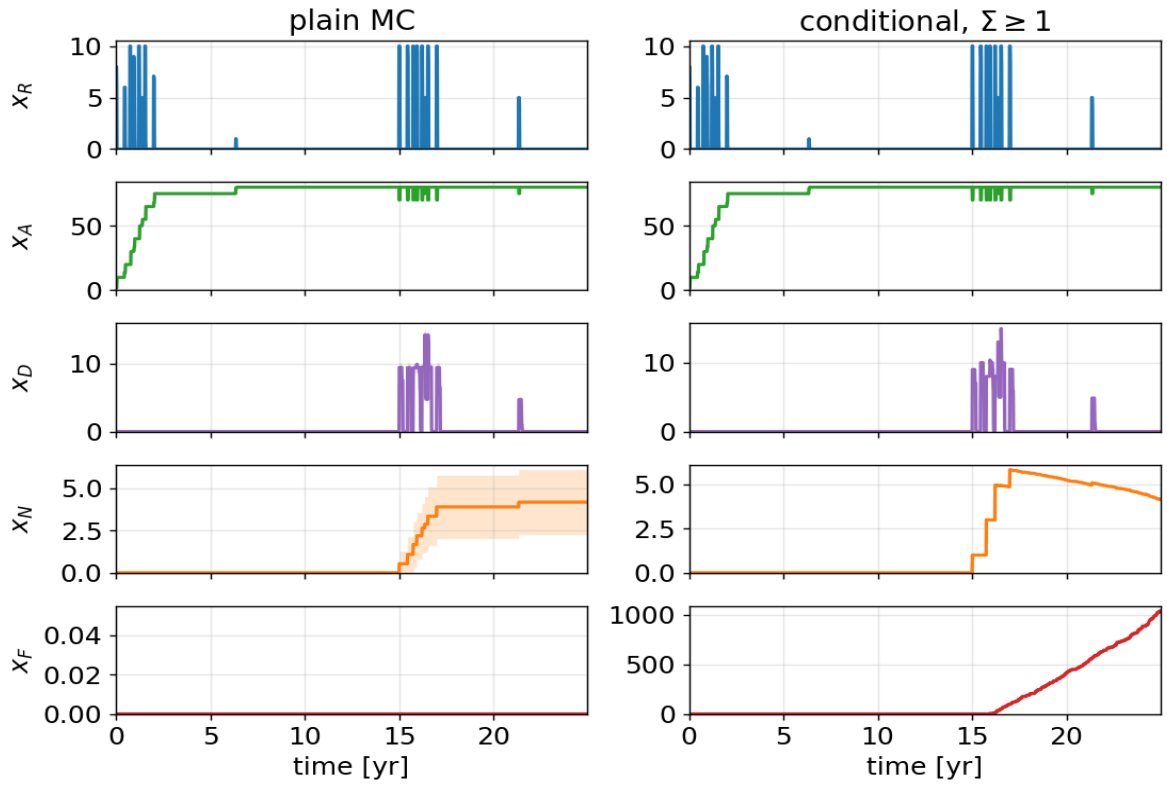

Figure 8.3: Iridium-NEXT population evolution: plain Monte Carlo ensemble (left) versus conditional branch (right).

OneWeb and Iridium-NEXT show the same behavior on a smaller scale. The active fleet reaches its full complement after the initial ramp-up phase and holds it flat thereafter, as designed. The derelict channel opens at the operational lifetime, and climbs as the first batch of satellites reach end-of-life before flattening. The plain Monte Carlo column's fragment row stays flat at zero: with $\Lambda = 2.524 \times 10^{-4}$ for OneWeb and $\Lambda = 9.946 \times 10^{-6}$ for Iridium-NEXT, a 200-seed ensemble expects on the order of 0.05 and 0.002 collisions in total, respectively, so the plain ensemble never once samples the channel. The conditional column's F row is by construction never empty, since it forces exactly one collision and simulates the fragment branch that follows.

Read across the two columns, the branches track each other closely in the intact species and diverge completely in F, and both halves of that statement are expected. They share the same launch schedule, the same end-of-life clock and the same PMD failure rate, and differ only by one forced breakup, which perturbs the deployment ramp and the disposal channel not at all: the active row is identical to the last object in all three populations at the final epoch, and Starlink's derelict and de-orbiting rows differ by well under a percent. The one intact channel that does move is the derelict count of the small populations, and the reason is that a collision consumes the objects that caused it. The parents are lost with the event, they are usually derelicts, and a constellation holding a handful of them feels the loss where one holding hundreds does not: the derelict row falls by a few percent for OneWeb and by roughly a third for Iridium-NEXT, whose derelict population is six objects.

The F row is where the two columns stop answering the same question. On the left it is a plain Monte Carlo estimate of a channel that opens only when a collision is sampled, and with Λ between 1.647×10^{-2} and 9.946×10^{-6} the 200-seed ensemble samples it three or four times for Starlink and never for the other two, so the left F row is either undersampled noise or an empty axis. On the right it is an exact conditional expectation, converged because every seed carries a collision by construction. The two disagreeing is the entire reason the conditioning scheme exists.

8.1.2. Forced Collision

Because the constellations are sparse, Λ ranges from 1.647×10^{-2} down to 9.946×10^{-6} over their respective horizons, and a plain Monte Carlo ensemble sees practically no collisions. What few it does see are noise, not converged statistics. The conditioning runner (section 5.2.3) replaces it with the exact decomposition of eq. (5.1), drawing the first collision from its true first-arrival time and simulating the fragment-carrying branch explicitly.

Table 8.3: Baseline conditioning summary. Λ is the expected number of collisions over the horizon, eq. (5.2), and $P(\Sigma \geq 1) = 1 - e^{-\Lambda}$ follows from it exactly. $\mathbb{E}[x_F]$ is the unconditional expected fragment count at the horizon, the $\mathbb{E}[X]$ of eq. (5.1) evaluated on $X = F$; $\mathbb{E}[x_F \mid \Sigma \geq 1]$ is the conditional branch alone, before it is weighted by how unlikely a collision is.

Population	Λ	$P(\Sigma \geq 1)$	$\mathbb{E}[x_F]$	$\mathbb{E}[x_F \mid \Sigma \geq 1]$
Starlink	1.647×10^{-2}	1.633×10^{-2}	2.257	124
OneWeb	2.524×10^{-4}	2.524×10^{-4}	0.085	321
Iridium-NEXT	9.946×10^{-6}	9.946×10^{-6}	0.011	1038

The two fragment columns of table 8.3 differ by the weight $P(\Sigma \geq 1)$ that eq. (5.1) attaches to the collision branch. The conditional count describes that branch: a few hundred fragments still on orbit at the horizon after a single forced breakup, which is what the constellation would face if a collision did occur. The unconditional count is that same cloud discounted by the probability that it occurs at all.

Of the three fragment estimates this section produces, it is the unconditional $\mathbb{E}[x_F]$ that should be read as the projection. The plain Monte Carlo estimate is undersampled past the point of usefulness, and is reported to demonstrate that rather than to be believed. The conditional $\mathbb{E}[x_F \mid \Sigma \geq 1]$ is not a forecast either, despite being the converged quantity and the one plotted in figs. 8.1 to 8.3: it describes a branch whose probability of occurring is between 1.6 % and one in a hundred thousand, and it exists as the intermediate that eq. (5.1) needs. Weighting it by $P(\Sigma \geq 1)$ closes the decomposition and gives $\mathbb{E}[x_F]$, the expected fragment contribution of the constellation over its horizon and the number an environmental budget should carry: 2.257 fragments for Starlink, 0.085 for OneWeb, 0.011 for Iridium-NEXT. All three falling below one is a statement about how unlikely the collision is, not about how mild it would be, and the severity if it does occur is the conditional column beside it. As an internal consistency check, integrating the baseline collision-rate over the horizon independently reproduces the conditioning runner's Λ to within a few percent for every constellation, a residual attributable to the trapezoidal quadrature on a strided time grid rather than to any disagreement between the two computations.

The headline result is the rate feedback: how much a single collision changes the constellation's own future collision rate, relative to the collision-free baseline. Table 8.4 reports the mean rate over the whole horizon and the rate at the final epoch, baseline and conditioned on one collision.

Table 8.4: Rate feedback factor: baseline (no collision) versus exact conditional ($\Sigma \geq 1$) collision rate.

Population	**Horizon mean [yr^{-1}]**			**Final epoch [yr^{-1}]**		
	baseline	$\Sigma \geq 1$	factor	baseline	$\Sigma \geq 1$	factor
Starlink	8.2935×10^{-4}	9.9690×10^{-4}	1.20	3.9239×10^{-3}	5.9949×10^{-3}	1.53
OneWeb	2.6853×10^{-6}	7.0698×10^{-6}	2.63	9.0782×10^{-6}	4.1407×10^{-5}	4.56
Iridium-NEXT	3.9824×10^{-7}	1.5546×10^{-5}	39.04	1.1687×10^{-6}	1.0119×10^{-4}	86.58

Table 8.5 breaks the same rate down by species pair, at the final epoch, showing which channel carries it before and after the collision.

Table 8.5: Leading species-pair channels at the final epoch, baseline vs. conditional branch, CAM-weighted rate, ranked by contribution. A dash marks a population with fewer than four channels above the reporting threshold.

Population	Branch	1st [yr^{-1}]	2nd [yr^{-1}]	3rd [yr^{-1}]	4th [yr^{-1}]
Starlink	baseline	N-N: 3.85×10^{-3}	A-N: 6.93×10^{-5}	N-D: 1.30×10^{-6}	N-R: 7.34×10^{-7}
	conditional	N-N: 4.33×10^{-3}	N-F: 1.57×10^{-3}	A-N: 7.34×10^{-5}	A-F: 1.40×10^{-5}
OneWeb	baseline	N-N: 9.01×10^{-6}	A-N: 6.99×10^{-8}	A-A: 9.17×10^{-11}	–
	conditional	N-F: 3.12×10^{-5}	N-N: 7.67×10^{-6}	F-F: 2.40×10^{-6}	A-F: 1.09×10^{-7}
Iridium-NEXT	baseline	N-N: 1.16×10^{-6}	A-N: 6.37×10^{-9}	A-A: 6.38×10^{-12}	–
	conditional	N-F: 6.22×10^{-5}	F-F: 3.83×10^{-5}	N-N: 5.06×10^{-7}	A-F: 1.72×10^{-7}

The feedback factor runs *inversely* with how likely the collision was to happen in the first place. Starlink is by far the most likely to suffer a breakup and the least perturbed by it (factor 1.20 on the horizon mean), because its baseline risk is already dominated by a large standing derelict population, so one fresh cloud is a marginal addition to N-N. Iridium-NEXT is the least likely and the most perturbed (factor 39.04, or 86.58 at the final epoch), because its baseline N-N term is tiny and a single cloud dwarfs it; OneWeb sits in between. A second reason is altitude: OneWeb and Iridium-NEXT sit much higher, where drag is far weaker, so any fragmentation sampled there dwells much longer, while a cloud at Starlink's altitude re-enters faster. Notably, in all cases A-A sits eight orders of magnitude below N-N everywhere, the direct consequence of the quadratic CAM suppression in eq. (4.3) at $s_{\text{CAM}} = 0.9999$: an operating constellation contributes essentially nothing to its own collision rate, which is entirely carried by what it has already failed to dispose of, plus what that collides with.

8.2. Network Topology and Plane Resolution

Everything so far has been an aggregate: a population count, a total rate, a channel ranking. None of it needs RAAN-resolution, which is in fact PRISM's distinguishing feature against the evolutionary models it descends from. This section uses that feature to study how the risk distributes over the constellation planes, how a fresh fragment cloud spreads across them, and how risk-drivers rank in the network topology.

Two post-processing readouts of the topology of section 7.2 are needed first. Both are evaluated on the collision adjacency matrix or on the node counts of a recorded epoch, so any epoch of any run can be interrogated after the fact. The first counts planes. Let f_n be a specie's RAAN marginal, its node counts summed over altitude and inclination and normalized, and let

$$n_{\text{eff}} = \frac{\left(\sum_n f_n\right)^2}{\sum_n f_n^2} = \frac{1}{\sum_n f_n^2} \tag{8.1}$$

be the participation ratio of that marginal, an effective number of occupied planes. If the specie sits in equal parts on k planes and nowhere else, then $f_n = 1/k$ on those planes and zero on the rest

$$\sum_n f_n^2 = k \cdot \frac{1}{k^2} = \frac{1}{k}$$

and $n_{\text{eff}} = k$. An uneven split returns less than the number of non-empty planes, since every plane enters weighted by its own share and one holding a negligible fraction of the specie contributes negligibly to the count: a specie present on ten planes but concentrated in a single one reads close to 1 rather than close to 10. The two limits follow as special cases, $n_{\text{eff}} = 1$ for a specie confined to one plane and $n_{\text{eff}} = N_\Omega$ for one spread evenly over all of them.

This is the same functional as the effective width of eq. (6.6), but the use here is the opposite. There the functional returns a length and is used prescriptively, to set the altitude bin width that resolves a population without diluting its rate. Here it returns a count and is used descriptively. The RAAN axis is periodic and naturally discrete for a constellation, since the planes are physical objects of the constellation's design, not artifacts created by binning. There is therefore no

bin width to choose: the bins are the planes themselves. What n_{eff} places instead is where a population sits between the two limits above. The upper one, a specie smeared uniformly over every plane, is precisely what a model without a RAAN dimension assumes from its first step, so the gap between a measured n_{eff} and N_Ω is the structure a resolved model sees and a blind one cannot. At the 5° resolution used throughout, $N_\Omega = 72$ planes total, so a specie fully smeared by a RAAN-blind model would show 1.39 % in each of them.

The second readout localizes the risk. The total collision rate is split over lattice groups, a group being one altitude shell together with one unordered pair of inclination bins. Cross-shell links are structurally zero, so the shell is unambiguous; the inclination pair is unordered because the links themselves are. A group g at shell ℓ and inclination pair $\{m_1, m_2\}$ is the node set $g = \{(\ell, m, n, S) : m \in \{m_1, m_2\}\}$, and the rate its plane n carries is

$$R_g(n) = \sum_{\substack{S_i \in g \\ n(S_i)=n}} \sum_{S_j \in g} \Xi^{\mathcal{C}}_{S_i S_j} \tag{8.2}$$

the same sum as the node strength of eq. (7.1), restricted to the group and collected by plane rather than by node. Two numbers are read off the resulting distribution $\{R_g(n)\}$: the busiest plane's share of the group, $\max_n R_g(n) / \sum_n R_g(n)$, normalized by the number of planes a group actually occupies n_Ω^{occ}, which is often smaller than n_Ω; and the ratio $\max_n R_g(n) / \min_n R_g(n)$ of the busiest to the quietest plane. A group of identical, uniformly phased planes should show the first metric at that uniform reference and the second at unity. Any departure from these values indicates the plane structure the model claims to see.

8.2.1. Population per Plane

Table 8.6 reports the participation ratio of each populated specie at each constellation's final epoch. The active totals are larger than the satellite counts of table 8.2 for that reason: those are the nominal shell sizes as DISCOS declares them, summing to 15 258 satellites for Starlink, whereas the table below counts what is actually on orbit ten years later. Several Starlink shells begin deploying after the simulation epoch and are still filling during the run, and the open-loop launch model keeps replacing satellites as they retire rather than holding the fleet at its nominal size, so the standing active count at the final epoch is 16 888. The two numbers describe different things and are not expected to agree.

The active row is the verification. OneWeb returns $n_{\text{eff}} = 12.00$ against its twelve launch planes and Iridium-NEXT 5.99 against its six. The lattice recovers the true Walker plane count from the population alone, which is a direct check that the 5° binning resolves planes rather than smearing them.

Table 8.6: Participation ratio and busiest-plane share of each populated specie at the final epoch, over $N_\Omega = 72$ planes. Uniform reference: $n_{\text{eff}} = 72$, busiest share 1.39 %.

Population	Specie	Total	n_{eff}	Busiest share
Starlink	A	16 888	25.47	6.44 %
	N	238.8	66.44	2.97 %
	F	282.4	66.55	3.00 %
OneWeb	A	656	12.00	8.38 %
	N	28.7	13.43	13.14 %
	F	320.5	42.17	5.22 %
Iridium-NEXT	A	80	5.99	17.50 %
	N	4.1	6.56	17.11 %
	F	1037.7	42.14	6.46 %

Starlink's pooled $n_{\text{eff}} = 25.47$ matches no single plane count, unsurprisingly. OneWeb and Iridium-NEXT are each one launch group in one shell at one inclination, whereas Starlink is

built from fourteen named launch groups spanning several genuinely different shells, so pooling its active count before taking the participation ratio blends together shells that do not share a plane count. Splitting the same final-epoch count by shell resolves the ambiguity: five distinct shells are occupied, at 300 km and 52.5° with $n_{\text{eff}} = 27.99$ over 797 satellites, and at 500 km with 64.00 at 37.5° (2592 satellites), 29.82 at 52.5° (7468), 13.20 at 67.5° (2098) and 10.81 at 97.5° (3933).

None of these five is a whole number, and the reason follows from the definition. n_{eff} coincides with a plane count only when the occupied planes hold equal shares, which is why the two single-group constellations return 12.00 and 5.99 and why the 300 km shell, whose 28 occupied planes hold 29 satellites each bar one, returns 27.99. A shell fed by more than one launch group is not equally occupied. The 500 km, 52.5° cell is the clearest case: 67 of the 72 planes hold satellites, but ten of them hold around 400 each and most of the rest hold a quarter of that or less, so the participation ratio reads 29.82 rather than 67. The number is smaller than the count of non-empty planes exactly to the degree that the population is concentrated in a few of them, and a fractional value simply reflects that the shares are unequal rather than that a fraction of a plane exists.

The derelict and fragment rows say something different and more consequential. Both sit far above the active row in n_{eff}, and for Starlink both are within 8 % of the uniform 72. A specie that starts confined to the operational planes ends up nearly uniform in RAAN, because J_2-induced nodal drift is altitude-dependent and no longer compensated once station-keeping stops, so an uncontrolled object leaves its plane and eventually visits all of them. The constellation's own payloads are therefore the one specie for which a RAAN-averaged treatment is least defensible, and its debris the one for which it is most, but only once the drift has had years to run its course; section 8.2.2 tracks how long that takes and shows the debris population spends most of the constellation's operating life far from uniform.

8.2.2. Risk per Plane

The link structure over RAAN separation is shown in fig. 8.4. Panel (b) shows the relative velocity between two objects sharing a shell but sitting in different orbit planes, plotted against the RAAN separation $\Delta\Omega$ between those planes and evaluated at the bin centers with eq. (4.28). There is no surprise in it: the closing speed grows monotonically from zero for co-planar pairs to 12 km/s to 15 km/s for counter-rotating ones at 180° of separation, and it is larger for the more inclined constellations, since inclination enters the same expression alongside the plane separation. Panel (a) shows the consequence, cleanly separating the two regimes. OneWeb and Iridium-NEXT occupy only a handful of RAAN separation classes, since their planes are evenly spaced by design: with P planes spaced $360/P$ degrees apart, the only separations any pair can have are multiples of that spacing, folded onto $[0, 180]$ degrees because the pair is unordered, which leaves $P/2$ non-co-planar classes. That is exactly the handful of peaks panel (a) shows: three for Iridium-NEXT's six planes, at 60°, 120° and 180° of separation, and six for OneWeb's twelve, at multiples of 30°. Their rate per link rises by four to five orders of magnitude from the co-planar class to the seam at 180°: the seam of a Walker-star pattern, where ascending and descending planes meet head-on, carries the population's worst links by a wide margin. Starlink occupies many more RAAN separation classes instead, since its shells carry different numbers of planes, and its rate per link varies by only about an order of magnitude across the whole range, because by the final epoch its links at every separation are dominated by its derelict population, which section 8.2.1 shows has drifted to near uniformity in RAAN, swamping most of the structure the design geometry would otherwise show.

What survives that swamping are the spikes riding on the continuum, and they are the constellation's own plane structure showing through. Starlink is not one Walker pattern but several superimposed: its launch groups carry 6, 10, 18, 28, 72 and 96 planes, and each admits only separations that are multiples of its own $360/P$ spacing. The finely divided groups, at 72 and 96 planes, place pairs every 5° and 3.75° and so fill in the near-continuous background; the five groups of ten planes can place pairs only at multiples of 36°, concentrating their rate onto that handful of classes, which is where the peaks sit. The coarsest pattern is the one that shows, and the finer ones are what it shows against. The co-planar class sits at the floor in all three

cases, the self-pair statement made earlier seen directly: objects sharing a plane have no closing speed and contribute essentially nothing to the collision risk.

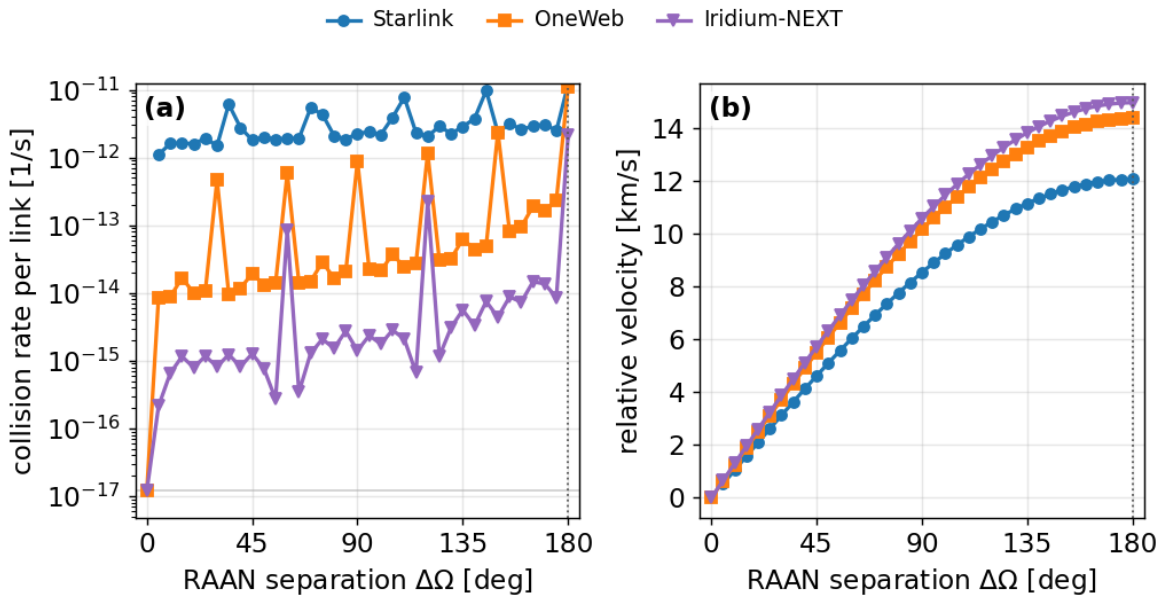

Figure 8.4: Collision link structure against RAAN separation at the final epoch. (a) Mean collision rate per link in each separation class. (b) Pair relative velocity. The dotted line marks the 180° seam.

Table 8.7 turns this into the concentration statement, for both branches of the conditioning run. The table measures concentration of collision rate over the planes. The ”Occupied planes” column counts the planes that appear at either end of at least one collision link, however small the rate that link carries. For the baseline branch the count exceeds the constellation’s operational plane count, since it includes every plane a derelict has drifted into and now collides in. For the conditional rows the forced breakup spreads fragments over the whole RAAN axis, so every one of the 72 planes qualifies. The ”Busiest share” and ”Uniform” columns are the two numbers defined at the start of this section: the busiest plane’s share of the group’s rate, and the uniform reference $1/n_\Omega^{\mathrm{occ}}$ it is measured against. The ”Max/min” column is the other of the two, the ratio between the busiest and quietest planes. In every row the busiest share sits well above the uniform reference and the max/min ratio well above unity, confirming that the plane structure these metrics are built to detect is actually there.

Table 8.7: Concentration of the collision rate over the RAAN planes at the final epoch. The uniform reference is the reciprocal of the number of occupied planes.

Population	Branch	Occupied planes	Busiest share	Uniform	Max/min
Starlink	baseline	72	5.34 %	1.39 %	2632
	conditional	72	2.55 %	1.39 %	3.80
OneWeb	baseline	26	15.41 %	3.85 %	1017
	conditional	72	10.76 %	1.39 %	76.9
Iridium-NEXT	baseline	12	21.32 %	8.33 %	1083
	conditional	72	9.15 %	1.39 %	35.5

The conditional row’s per-plane rate is estimated rather than exact: the true conditional rate would require reconstructing information the aggregation scheme discards for computational efficiency, so the row is built from mean-field links evaluated on the conditional population instead. The estimate gets the structure right, which plane is busiest and how concentrated it is,

but the absolute magnitude should be treated with caution. Every conclusion drawn from the table is qualitative about that structure, not quantitative.

The baseline rows are the conditioning run's collision-free branch, the $\Sigma = 0$ term of eq. (5.1), in which collisions are disabled outright and no fragment is ever produced. They are not the plain Monte Carlo ensemble of section 8.1.1, which does sample a handful of collisions for Starlink across its 200 seeds. Read on those terms, the table shows the risk to be structurally non-uniform before any collision has occurred at all. The busiest plane carries 3.8 times the uniform share for Starlink, 4.0 for OneWeb and 2.6 for Iridium-NEXT, with busiest-to-quietest ratios of order 1×10^3 in all three. A RAAN-averaged model reports the mean of that distribution and can say nothing about which plane it belongs to.

Measured at the final epoch, the collision risk *homogenizes*. In every case the conditional busiest share is lower than the baseline share, and the busiest-to-quietest ratio collapses from order 1×10^3 to between 4 and 77. The mechanism is the drift of table 8.6: the cloud fills planes that carried little or no risk before, so the distribution becomes flatter and broader even as its total rises. The two effects are separate, and a plane-resolved model is needed to keep them apart at all, since the aggregate rate goes up while its concentration goes down. However that is not the picture along the full trajectory of the fragmentation event. Table 8.8 tracks the fragment participation ratio in the conditional branch from the first fragment epoch onwards.

The x_F column is the number of fragments of one breakup weighted by the probability that the breakup has already taken place, i.e., the unconditional fragment count. The conditioning fixes a single collision somewhere in the horizon and draws its epoch from the first-arrival distribution of section 5.2.3, so at an early epoch most realizations have not yet reached their collision and contribute no fragments to the average, while at a late one nearly all of them have. The column consequently rises toward the yield of one event, less what has decayed in the meantime: Iridium-NEXT reaches 1038 fragments against the 1302 that a collision between objects of that mass produces. Fragments so produced are themselves subject to collision, and those secondary events are sampled like any other, but at these rates they add nothing measurable to the totals.

Table 8.8: Fragment cloud spreading over the RAAN planes, conditional branch, from the first fragment epoch. Uniform reference $n_{\text{eff}} = 72$ at a busiest share of 1.39 %.

Population	Epoch [yr]	x_F	n_{eff}	Busiest share of x_F
Starlink	0.00	2.2	1.29	87.78 %
	1.03	13.1	41.70	7.11 %
	3.08	76.4	57.77	4.20 %
	5.15	282.4	66.55	3.00 %
OneWeb	0.00	1.6	1.21	90.82 %
	1.51	24.0	13.14	15.74 %
	4.53	149.0	28.22	6.53 %
	7.57	320.5	42.17	5.22 %
Iridium-NEXT	0.00	6.0	1.00	99.91 %
	1.81	175.0	20.08	11.63 %
	5.42	557.2	32.74	8.36 %
	9.05	1037.7	42.14	6.46 %

At onset the cloud is a single-plane object in every case: n_{eff} between 1.00 and 1.29, with 88 % to 99.9 % of the fragments in one plane. This is the physically expected result, since a fresh cloud is born co-planar with its parent and the ejection speeds are small against the orbital speed, but it is also the point at which a blind model is furthest from the truth. Spreading a cloud over all 72 planes on day one misstates the near-field risk by close to two orders of magnitude, at precisely the epoch when an operator would need it, and it does so in the optimistic direction

for the parent plane. The spreading is also slow. Iridium-NEXT reaches only $n_{\text{eff}} = 42$ out of 72 after nine years and OneWeb the same value after 7.6 years; neither reaches uniform inside its mission horizon. Starlink is the fastest, at 66.6 out of 72 in five years, because its clouds are born in low shells where the differential nodal drift is strongest. In no case does the RAAN-blind assumption become correct within the horizon over which the constellation is actually operated.

While the cloud remains concentrated, the plane holding it is not the only one exposed. Figure 8.4 gives the rate a link carries as a function of the RAAN separation between its two planes, and the counter-rotating pairing at 180° carries four to five orders of magnitude more than the co-planar one. The plane facing the cloud head-on is therefore the one most at risk from it, more so than the planes adjacent to it in RAAN, and it stays that way for as long as the cloud has not spread.

8.2.3. Node Strength

The preceding two subsections attribute rate to groups of nodes and to planes. Section 7.2.2's node strengths attribute it to single nodes instead, on three axes: the collision strength $s^{\mathcal{C}}$ of eq. (7.1) ranks a node by how often it collides, the fragment strength s^{F} of eq. (7.2) by how much debris those collisions make, and the net transport strength $s^{\text{out}} - s^{\text{in}}$ separates the nodes that emit objects from those that accumulate them. Only the last of the three is independent of the other two. As set out in section 7.2.2, $s^{\mathcal{C}}$ is the likelihood of a collision at the node and $s^{F}/s^{\mathcal{C}}$ its consequence in fragments per collision, so their product, the combined risk that a likelihood-times-consequence ranking would ask for, is s^{F} itself and is reported directly below. The three columns of table 8.9 are read on those terms.

The node strengths reported below carry the scenario collision-avoidance success $s_{\text{CAM}} = 0.9999$, so they rank risk rather than the geometric conjunction rate, the opposite convention from fig. 8.4 and table 8.7. Those two use the geometric rate to ask whether the lattice itself produces a plane dependence, a question CAM would mask; a ranking by geometry instead of by risk, on the other hand, would be dominated by the active-active pairs operators already avoid, and would say little about which nodes still matter once avoidance is accounted for. The flow graph built here carries all modeled transport interactions: failed PMD, natural transport in altitude and RAAN, post-collision fragment redistribution, and the $A \to N$ damage from non-catastrophic impacts. Re-entry is absent as a link, since the lowest shell decays out of the lattice and has no destination node, which makes the subgraph exactly conservative: summed over every node the net strength is zero to roundoff.

Table 8.9: The three nodes of each population carrying the most collision rate at the final epoch, one lattice cell per row. Collision and fragment strengths include the CAM factor of eq. (4.3); the net strength is over the flow channels listed in the text. Every leading node is a derelict.

Population	**Node** (h, i, Ω)	**Specie**	$s^{\mathcal{C}}$ [1/yr]	s^{F} [1/yr]	$s^{\text{out}} - s^{\text{in}}$ [1/yr]
Starlink	500 km, 67.5°, 307.5°	N	9.80×10^{-5}	8.49×10^{-2}	355
	500 km, 37.5°, 277.5°	N	9.75×10^{-5}	8.40×10^{-2}	958
	500 km, 52.5°, 37.5°	N	8.55×10^{-5}	7.40×10^{-2}	294
OneWeb	1300 km, 86.2°, 202.5°	N	3.80×10^{-6}	1.28×10^{-3}	44.9
	1300 km, 86.2°, 17.5°	N	3.20×10^{-6}	1.08×10^{-3}	89.0
	1300 km, 86.2°, 167.5°	N	3.10×10^{-6}	1.05×10^{-3}	46.0
Iridium-NEXT	700 km, 86.2°, 317.5°	N	1.01×10^{-7}	1.31×10^{-4}	5.31
	700 km, 86.2°, 172.5°	N	8.83×10^{-8}	1.15×10^{-4}	7.63
	700 km, 86.2°, 222.5°	N	6.82×10^{-8}	8.88×10^{-5}	0.168

Table 8.9 ranks the three busiest nodes of each population by collision strength $s^{\mathcal{C}}$, alongside their fragment and net transport strengths. The first thing it shows is a statement about specie, not geometry: every leading node is a derelict, in all three populations, at every rank. The active

payloads that outnumber the derelicts by one to two orders of magnitude in table 8.6 appear nowhere near the top, precisely because of CAM suppression. The node-level geometry follows the same pattern as the previous two subsections: Starlink's leading cells span its 500 km shell across three different inclinations, whereas OneWeb and Iridium-NEXT sit in single shells at a single inclination, so for them the node index is effectively the plane index alone.

The net strength column counts transits, not objects brought into existence, and for a derelict node it is almost entirely J_2 nodal drift. At 500 km and 37.5° the line of nodes regresses at 6.1 °/d, so a single object crosses a 5° RAAN bin about 440 times a year. The second Starlink node holds 2.4 derelicts and therefore sends some 1050 crossings per year out of itself, while the bin upstream of it holds 0.22 objects and returns about 100, leaving the 958 in the table. Every one of those departures is an arrival in the neighbouring bin, so no object is added anywhere: summed over the network the net strengths cancel to roundoff, as noted above.

Table 8.10: Flow budget by specie at the final epoch: how many nodes of the specie are sources, how many are sinks, and the specie's total net strength.

Population	Specie	Nodes	Sources	Sinks	$\sum(s^{\text{out}} - s^{\text{in}})$ [1/yr]
Starlink	A	204	204	0	3378
	N	581	217	364	−160.8
	R	131	131	0	6.40×10^{-4}
	D	316	112	204	−3209
	F	3083	506	2577	−8.14
OneWeb	A	12	12	0	65.6
	N	99	30	69	−3.27
	D	12	0	12	−62.3
	F	1218	0	1218	-9.17×10^{-3}
Iridium-NEXT	A	6	6	0	5.33
	N	78	24	54	−0.265
	D	6	0	6	−5.07
	F	880	0	880	-1.32×10^{-3}

Table 8.10 tallies the specie-level budget. Every active node is a source and its net strength is almost entirely the end-of-life transition, 95 % of Starlink's 3378 per year being $A \to D$ and the remainder failed PMD, matched within a fraction of a percent by the D row it feeds. The D row counts 112 sources against 204 sinks despite receiving that whole transition, because a de-orbiting object can still be hit: it is credited as a parent of the fragments its own collision produces, and it can be demoted to N by a non-catastrophic impact. Both channels are open to it and neither is significant, together carrying some 1×10^{-11} objects per year against the 3209 per year the specie receives, so they are enough to make a node a nominal source without moving the budget at all. The derelict and fragment species are net sinks. N receives the failed disposals faster than its collisions consume it, and F receives every fragment produced anywhere in the lattice: an F node only becomes a source where transport carries fragments out of it, down in altitude or around in RAAN, faster than new fragments arrive there, whether from a collision elsewhere or from transport in. Furthermore, every fragment leaving an F node lands on another F node, or leaves the lattice through re-entry at the lowest shell, and never crosses to another species, so every source's outflow becomes some other node's inflow, which leaves few fragment sources against many sinks in every population.

Table 8.11: Fragment yield s^F/s^C by specie at the final epoch, given as the lowest and highest value over the nodes where both strengths are non-zero. The ratio is a fragment count per collision, so it is unaffected by the CAM factor common to both strengths.

Population	Specie	Nodes	s^F/s^C	Spread
Starlink	A	204	[839.8, 866.5]	1.03
	N	425	[314.1, 866.5]	2.76
	R	131	[840.0, 866.1]	1.03
	D	218	[839.8, 866.2]	1.03
	F	776	[0.954, 506]	530
OneWeb	A	12	[336.9, 336.9]	1.00
	N	39	[336.9, 347.0]	1.03
Iridium-NEXT	A	6	[1302, 1302]	1.00
	N	33	[1302, 1304]	1.002

Each row of table 8.11 covers the nodes of that specie at which both strengths are non-zero, which is what the ratio requires: a node carrying no collision link has no s^C to divide by, and one whose collisions produce no fragments has no s^F. The bracketed pair is the lowest and the highest ratio among those nodes, and the Spread column is the highest divided by the lowest. Every intact specie's Spread sits within a few percent of unity, because a collision involving an intact object is catastrophic at essentially any closing speed the lattice produces and the fragment count then depends only on the involved mass. The fragment nodes carry a Spread of 530 instead, because a fragment-fragment event is close to the catastrophic threshold and its yield is dominated by which side of that threshold the pair falls on.

Within an intact specie, then, s^F is just s^C times a fixed number. The two therefore rank its nodes in the same order, and ranking by risk gives the same answer as ranking by frequency. Only among the F nodes, where the yield varies by a factor 530, can the two orderings differ.

8.2.4. The Case for Plane Resolution

The cleanest argument is calibration. Table 8.12 compares the horizon expected collision count at two RAAN resolutions. At 60° several real planes are merged into a single bin, while at 5° one bin holds one plane. Only OneWeb and Iridium-NEXT appear, since they are the two constellations that coarsening actually affects: their derelicts carry most of the collision rate and are still concentrated on a handful of planes at the final epoch. Table 8.6 gives $n_{\text{eff}} = 13.43$ for OneWeb and 6.56 for Iridium-NEXT against 12 and 6 design planes, so merging those planes at the coarser resolution discards real structure. Starlink is left out of the table for the opposite reason: its derelicts have drifted to $n_{\text{eff}} = 66.44$ out of 72, close enough to uniform that merging planes does not merge anything the population has not already smeared evenly across them. The two constellations shown here move in opposite directions.

Table 8.12: Horizon expected collision count Λ on a 60° and a 5° RAAN lattice, for the two constellations whose planes the coarse lattice merges.

Population	60° lattice	5° lattice	Change
OneWeb	4.940×10^{-5}	5.333×10^{-5}	7.9 %
Iridium-NEXT	1.1744×10^{-5}	9.946×10^{-6}	−15.3 %

The coarse lattice is not conservative. It is wrong in a direction that cannot be predicted without resolving the planes, which is a stronger argument for the resolution than any single-signed bias would be: a systematic overestimate could be absorbed into a safety factor, whereas an unsigned error of order 10 % cannot.

Beyond calibration, this chapter has established three things that no RAAN-averaged model

can produce. A breakup is a localized event that stays localized for years, so the near-field risk after a fragmentation is knowable only with the planes resolved. The risk is structurally non-uniform in the absence of any collision, by a factor of 2.6 to 4.0 on the busiest plane and by orders of magnitude across the lattice, which makes "which of my planes is the exposed one" an answerable question. And the lattice reproduces the true Walker plane count from the population, which verifies the resolution rather than assuming it.

The honest counterweight is that for the *aggregate* rate the resolution buys almost nothing. The ratio $r = \dot{\eta}^{\text{blind}}/\dot{\eta}$ between a RAAN-blind rate, with each pair's overlap integral replaced by its mean over a uniform RAAN distribution, and the resolved rate gives $r = 1.012$ on Starlink's collision-free branch at the final epoch and 1.004 on the conditional branch. That is not a disappointment so much as the expected consequence of the calibration identity: a constellation's payloads are uniform in RAAN by design, its derelicts and fragments become nearly so within a few years, and a uniform population is exactly the case in which the blind overlap is right. The case for the RAAN dimension is therefore not that it changes the total: it never claimed to, and if it did, every established evolutionary model that averages over RAAN would have to be wrong, which they are not. The agreement is a consistency check, not a failure. The actual case is that resolving RAAN is the only way to obtain the spatial distribution of the risk and the time-varying localization of a fresh cloud.

RAAN averaging is not wrong in general: it is well suited to the heterogeneous, broadly distributed populations most evolutionary models track, and this chapter's own aggregate numbers show it working exactly as it should on that kind of population. What it is not suited to is a rare, localized event in a constellation with a small number of discrete planes, which is the regime under study, and resolving RAAN is what makes that mismatch visible instead of silently absorbed into the aggregate.

8.3. Carrying Capacity

Both quantities reported so far, the horizon expected collision count Λ and the expected fragment count, are integrals over a particular mission clock. They answer how much risk a given scenario accumulates, and they cannot answer the question an operator or a regulator actually asks, which is how much more traffic a shell will take before it stops clearing its own debris. That question is about a balance of rates at an instant, not about an accumulation over a horizon, and it is answered here with a single number, the carrying-capacity indicator ϱ_{crit}, computed from a reduced form of the mean-field model.

8.3.1. The Capacity Indicator

The indicator is not read off the stochastic simulator of the earlier sections, which produces one realization at a time and has no notion of an equilibrium; it needs a deterministic model instead. The question it answers is about a single shell's own balance of production and removal, not about where in the lattice that balance is tightest or how a maneuvering fleet gets there, so the model can discard everything the question does not need rather than solving the general system of section 7.3 in full. Therefore, inclination and RAAN are summed out, so a node is one altitude shell and one specie, and the raising and de-orbiting species are dropped, since the capacity study runs without maneuvers and injects launches directly into the active population. The dynamical state is therefore $\boldsymbol{X}$, holding $\{A, N, F\}$ per shell, and restricting eq. (7.3) to that smaller state, with $\boldsymbol{L}$ and $\boldsymbol{B}$ acting only on the surviving components, gives

$$\dot{\boldsymbol{X}} = \boldsymbol{G}(\boldsymbol{X}) = \boldsymbol{L}\boldsymbol{X} + \boldsymbol{B}(\boldsymbol{X}, \boldsymbol{X}) + \boldsymbol{g} \tag{8.3}$$

The three terms carry the same meaning they carry in chapter 7. The linear operator $\boldsymbol{L}$ holds drag transport, which moves N and F one shell down, the $A \to N$ transition from failed PMD, and the removal of active satellites at end of life. The bilinear operator $\boldsymbol{B}$ holds the collisions: a catastrophic event removes both parents, a non-catastrophic one removes the projectile and

downgrades an active target to N, and both deposit fragments, spread over the shells by the altitude redistribution kernel. The source $\boldsymbol{g}$ holds the exogenous launch traffic.

A single coefficient of $\boldsymbol{B}$ has to stand for many links. A node of the reduced model is an entire shell, so one number must represent the collisions of every resolved pair inside it, across all of that shell's inclination and RAAN bins, and the rate of those collisions depends strongly on which planes the objects occupy. Assuming them spread evenly over the angles would therefore give the wrong number, and section 8.2 shows by how much. The coefficients are measured instead. For one shell and one pair of species, the rate engine of section 4.2 is run on a recorded population at full resolution, its rate is summed over every resolved pair in that shell, and the total is divided by the product of the two species counts there. The coefficient that comes out reproduces the resolved rate exactly at that population, and continues to do so as the population is scaled up or down, since the rate is bilinear in the counts. What it cannot follow is a change in the shape of the angular distribution itself, which stays as it was recorded.

The capacity itself follows from how the two operators scale. Let $\boldsymbol{X}_0$ denote the recorded population and scale it by a common dimensionless factor ϱ, so that $\varrho = 1$ recovers the recorded state. The scaling is uniform: every shell and every specie is multiplied by the same ϱ, so the composition of the population and its distribution over the lattice are held fixed and only its density changes. Where the added objects are placed is therefore not a free choice here but a fixed one, and eq. (8.4) describes that one way of growing the population and no other. The two operators then scale by different powers. The collision rate of section 2.3 counts pairs, so the rate between two populations is proportional to the product of their counts and scaling both by ϱ multiplies it by ϱ^2; this is a property of pair counting rather than an empirical scaling law, and it is exact for the operator as defined. The linear operator acts on objects one at a time, so scaling its argument multiplies it by ϱ:

$$\boldsymbol{B}(\varrho\boldsymbol{X}_0, \varrho\boldsymbol{X}_0) = \varrho^2\, \boldsymbol{B}(\boldsymbol{X}_0, \boldsymbol{X}_0), \qquad \boldsymbol{L}(\varrho\boldsymbol{X}_0) = \varrho\, \boldsymbol{L}\boldsymbol{X}_0. \tag{8.4}$$

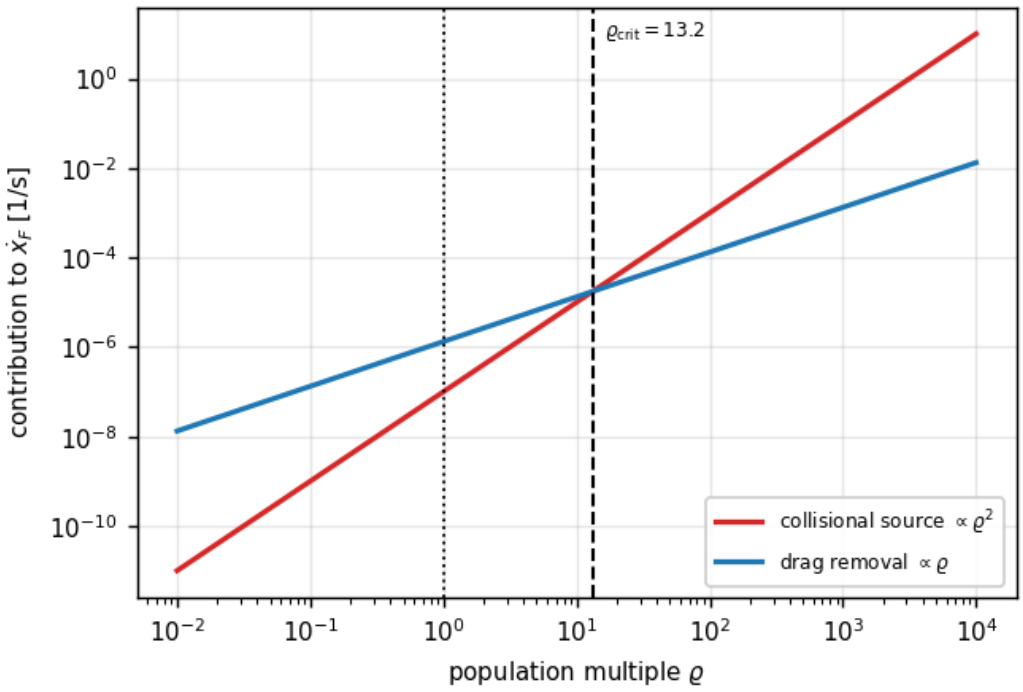

Figure 8.5: Carrying-capacity construction at Starlink's 500 km operating shell. The collisional source grows as ϱ^2 and drag removal as ϱ, and ϱ_{crit} is their single crossing. The dotted line marks the recorded population, $\varrho = 1$.

Write $G_F(\varrho\boldsymbol{X}_0)$ for the fragment component of $\boldsymbol{G}(\varrho\boldsymbol{X}_0)$ with the launch source set to zero,

that is, the net fragment production rate at ϱ times the recorded population. It is a quadratic in ϱ with a positive leading coefficient, so it changes sign at most once for $\varrho > 0$. That single crossing defines the carrying-capacity indicator ϱ_{crit}, through $G_F(\varrho_{\text{crit}}\boldsymbol{X}_0) = 0$: below it drag removes fragments faster than collisions make them and the shell clears its own debris, above it the balance has reversed. Because ϱ is a multiple of the recorded population, the reading is a margin rather than a rate: $\varrho_{\text{crit}} = 13$ means the shell would need thirteen times its recorded population before collisions make fragments faster than drag clears them, and $\varrho_{\text{crit}} < 1$ means it is already past that point at its present density. Figure 8.5 is the construction drawn for Starlink's operating shell: on logarithmic axes the collisional source is a straight line of slope two and the drag removal a line of slope one, and two straight lines of different slope meet exactly once.

Locating ϱ_{crit} directly, by scanning ϱ for the sign change, is a deliberate choice over the general Jacobian-and-equilibrium framework of section 7.3. That framework would solve for the equilibrium $\boldsymbol{X}^*(\varrho)$ at each value of ϱ and take ϱ_{crit} to be the value at which the Jacobian's spectral abscissa changes sign, and it is the more general tool: it applies to any direction of change in the state, not just a uniform rescaling, and it certifies full stability rather than a single crossing. It is also the more expensive one. It requires a converged equilibrium at every sampled ϱ, and the fragment-active regime is numerically stiff, since collisions and fragmentation act on a much faster timescale than drag, so a bounded equilibrium solver does not reliably follow the equilibrium as ϱ passes ϱ_{crit}, and settles instead on a spurious fragment-free solution. Restricting the populations considered to multiples of the recorded one, $\varrho\boldsymbol{X}_0$ for $\varrho > 0$, removes the need to solve for an equilibrium at all: among those populations the condition $G_F = 0$ returns the same ϱ_{crit}, and it can be evaluated directly on $\boldsymbol{G}$ without a solver. The cost is that the answer describes only that one way of growing the population, and not the stability of $\boldsymbol{X}^*$ against an arbitrary perturbation.

The indicator carries five limitations that should be kept in mind before it is used.

- The uniform scaling makes the quantity a counterfactual about density rather than a realistic growth path, along which the active, derelict and fragment populations would not grow together, nor spread over the shells in the proportions they hold today.
- The composition is frozen at the recorded epoch, so the indicator is a property of a state and not a forecast.
- Most relevant to this chapter, the reduced state has no plane index at all: everything section 8.2 established about the spatial distribution of the risk is deliberately outside the scope of this section, which is a statement about a shell-averaged balance.
- It is a ratio, so a shell can be labelled supercritical while its absolute fragment production is negligible, and the net rate column is reported alongside ϱ_{crit} throughout for that reason.
- The fragment population holds nothing below 10 cm. That threshold is set in section 4.5 and applies to every result in this chapter, but it bears on the capacity most directly, because both sides of the balance are built from fragments. The SBM count of eq. (4.41) grows as $L_c^{-1.71}$ toward small sizes, so a breakup puts roughly fifty times as many objects above 1 cm as above 10 cm, none of which this model carries. Objects in that range are not harmless: at orbital closing speeds a centimetre-class fragment carries enough energy to disable a satellite outright, and against a sufficiently light target to fragment it.

Every ϱ_{crit} reported here is therefore an upper bound on the true capacity. The omitted population would add collisions without adding any removal, so it can only move the crossing to the left, and a shell reported here as having room may have less of it than the number suggests. The indicator locates where the margin is thin; it does not predict a collisional cascade.

A verification check has been performed on a century-scale full-environment population, the same one used for verification in chapter 5. There, the indicator recovers the altitude-gated structure that is independently known from the literature: the band around 900 km to 1100 km already supercritical, and the low shells drag-cleaned and far below their capacity.

8.3.2. Baseline Capacity per Constellation

Table 8.13 reports the indicator across the shells of each baseline scenario, with the operating shell marked. Entries of $< 10^{-2}$ and $> 10^{4}$ are the two ends of the scanned range: the first means no crossing was found even at a hundredth of the recorded population, so the shell is supercritical throughout, and the second that none was found up to ten thousand times it.

The three constellations occupy at most two shells each. Starlink flies at 300 and 500 km, Iridium-NEXT at 700 km and OneWeb at 1300 km; every other row of the table holds fragments and nothing else. A breakup ejects debris over a wide span of altitudes on both sides of the operational shell, so each scenario leaves a thinning tail of fragments across the whole lattice. Starlink's 1700 and 1900 km rows carry 3.6 and 2.6 fragments respectively and no satellite at all, and their supercritical label follows from that: the debris deposited there by a breakup below has essentially no drag to remove it, so what little production those fragments sustain is never balanced.

Table 8.13: Per-shell carrying capacity for each baseline scenario. The operating shell is marked $\star$.

Shell	Starlink		OneWeb		Iridium-NEXT	
	ϱ_{crit}	$\dot{x}_F$ [1/yr]	ϱ_{crit}	$\dot{x}_F$ [1/yr]	ϱ_{crit}	$\dot{x}_F$ [1/yr]
300 km	760.94	-3.76×10^{2}	$> 10^{4}$	-1.66×10^{1}	$> 10^{4}$	-6.58×10^{1}
500 km	13.07*	-3.88×10^{1}	$> 10^{4}$	-9.25×10^{-1}	1587.61	-5.48
700 km	1.04	-2.05×10^{-2}	471.58	-8.90×10^{-2}	377.56*	-8.09
900 km	0.15	1.67×10^{-1}	$< 10^{-2}$	1.93×10^{-2}	113.47	-4.40×10^{-1}
1100 km	0.05	9.64×10^{-2}	10.84	-1.02×10^{-2}	19.86	-2.65×10^{-2}
1300 km	0.02	5.92×10^{-2}	6.89*	-3.19×10^{-2}	7.42	-4.62×10^{-3}
1500 km	0.02	3.86×10^{-2}	3.10	-2.19×10^{-3}	5.81	-2.05×10^{-3}
1700 km	$< 10^{-2}$	2.73×10^{-2}	1.25	-9.75×10^{-5}	1.74	-2.08×10^{-4}
1900 km	$< 10^{-2}$	1.94×10^{-2}	2.07	-2.17×10^{-4}	3.06	-3.99×10^{-4}

The headline for each is the marked row. Starlink's 500 km shell sits a factor 13.07 below capacity, OneWeb's 1300 km shell a factor 6.89, roughly half the margin, and Iridium-NEXT's 700 km shell a factor 377.56, about two and a half orders of magnitude clear. The ordering follows population size: Starlink flies 17 000 satellites where Iridium-NEXT flies eighty, and the bilinear collision term turns that difference in population into a much larger difference in margin.

The altitude trend is the more general result. Capacity falls steeply with altitude in every scenario, because drag removal is exponential in the scale height while the collisional source is not, so the sink term collapses over a few hundred kilometres while the source term barely notices. For Starlink the crossing of unity happens between 700 km ($\varrho_{\mathrm{crit}} = 1.04$) and 900 km ($0.15$), and the two shells above 1700 km are supercritical at their recorded population. However, the net fragment rates there are 2×10^{-2} to 3×10^{-2} fragments per year, so the supercritical label describes a balance between two negligibly small numbers rather than an ongoing cascade. The sign is nonetheless correct, and it is the quantitative version of the well-known statement that the region above roughly 800 km does not clean itself on operational timescales.

A few rows break the monotonic trend. OneWeb's 900 km shell is supercritical while the 1100 km and 1300 km shells above it are not, and Iridium-NEXT's 1900 km row sits above its 1700 km one. These are shells the constellation does not occupy, and their balance is set by their neighbours rather than by themselves: the fragment component of G at such a shell is dominated by what the redistribution kernel deposits into it and by the drag flux arriving from above, both of which are inherited, while its own bilinear production is essentially zero. The net rate column confirms the reading, at order 1×10^{-4} fragments per year in those rows. The indicator is a statement about a shell a population actually occupies; for one it merely borrows debris from, the sign is real but the ordering carries no information.

8.4. Case Studies

This section asks what the operator can actually change. Each of the three studies that follow takes the Starlink baseline over a 10 year horizon, picks one policy parameter, and re-runs the whole scenario at each of a range of values for it while holding everything else fixed. Such a series of runs is referred to as a sweep, and the three axes swept are CAM performance, PMD compliance and launch strategy, one per subsection.

Every table reports the same five quantities for each swept value. The first two are the active and derelict populations at the end of the horizon, which are the state the capacity indicator of section 8.3.1 is then evaluated on. The next two are outcomes accumulated over the run: the exact expected collision count Λ from the collision-free pass, and the expected fragment count at the horizon. The last is the operating shell's ϱ_{crit}. The first four say what the scenario does; the fifth says how much room it leaves before the shell stops clearing its own debris. The fragment count, and the fragment panel of every figure in this section, is a mean over the 200-seed ensemble, which is why both are reported as fractional rather than integer counts.

8.4.1. Impact of Different CAM Policies

Figure 8.6 and table 8.14 sweep the CAM failure probability α over four decades from the baseline 1×10^{-4} to complete failure.

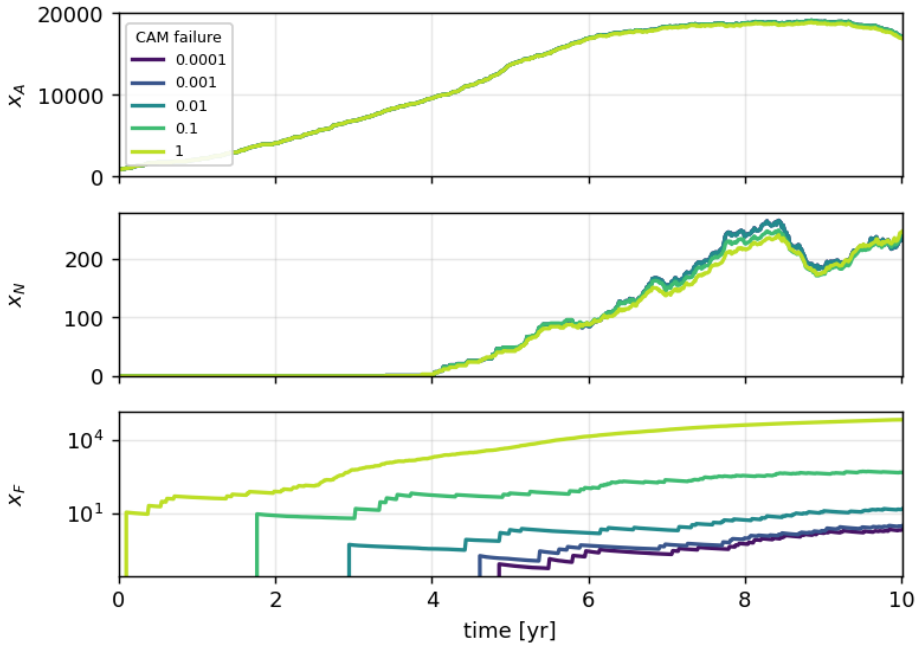

Figure 8.6: Starlink CAM failure sweep: active, derelict and fragment populations over the horizon.

Table 8.14: Starlink CAM failure sweep. † one forced collision per step misses coincident ones, so $\mathrm{E}[F]$ is a lower bound

CAM failure	Ref. A	Ref. N	Λ	$\mathbb{E}[x_F]$ at horizon	ϱ_{crit}
0.0001	17 148	239	8.32×10^{-3}	2.4	13.07
0.001	17 148	238	1.05×10^{-2}	3.4	13.16
0.01	17 147	241	4.72×10^{-2}	16.3	5.23
0.1	17 146	242	1.91	496.2	0.36
1	16 898	246	1.70×10^{2}	67 190.7†	0.36

The first two panels of fig. 8.6 are almost invariant: collision avoidance does not change how many satellites are launched or how many fail disposal, so the active and derelict curves are unaffected. On the other hand, the fragment panel spans four orders of magnitude across the same sweep. The reference columns of table 8.14 move for the same reason and by as little. They are counts at the end of the horizon, and the only route by which avoidance reaches them is the collisions themselves: a catastrophic event removes both of the objects involved, and a non-catastrophic one demotes an active satellite to a derelict. Neither is visible until collisions become numerous, which is why the active count holds near 17 150 across four of the five cases and falls by about 250 only at complete avoidance failure, the mark left there by 170 collisions.

The scaling of Λ across the sweep is the sum of two channels, one avoidance cannot touch and one it can. A derelict-derelict collision has no active endpoint, so nothing in eq. (4.3) suppresses it: no avoidance policy, however good, stops two objects that cannot maneuver from colliding. Table 8.5 identified that channel as already dominant at the baseline, and it sets a floor of roughly $\Lambda = 8.3 \times 10^{-3}$ that a ten-fold rise in the failure probability barely moves, from $\Lambda = 8.32 \times 10^{-3}$ at a failure probability of 1×10^{-4} to 1.05×10^{-2} at 1×10^{-3}. The second channel is the one CAM actually governs: active-derelict and active-active collisions, entering eq. (4.3) linearly and quadratically in the failure probability respectively, so avoidance suppresses active-active pairs far more effectively than it suppresses active-derelict ones. That protection erodes as failure grows, and once the active-endpoint channel has grown enough to catch up with the derelict-derelict floor it takes the total over: Λ reaches 4.72×10^{-2} at 1 % failure, 1.91 at 10 %, and 170 at complete failure, where both channels are unsuppressed and the active-active term, no longer quadratically damped, is what drives the last jump.

Read through the capacity indicator, the sweep has a threshold. The operating shell's ϱ_{crit} falls from 13.07 at baseline to 5.23 at 1 % failure and crosses unity between 1 % and 10 %, saturating near 0.36 thereafter. In other words a constellation of this size can absorb a roughly hundred-fold degradation of its avoidance performance and remain subcritical at its operating shell. That margin is bounded by the same size threshold as the capacity itself. Every object the sweep avoids, or fails to avoid, is at least 10 cm, and debris below that size is absent from the fragment population modeled here and from the catalogues on which avoidance depends, since an object that is not tracked cannot be screened against and no failure probability describes the encounter. The two omissions fall on the same population and compound, so the residual risk that survives good CAM performance is understated, and the hundred-fold margin is a statement about catalogued objects rather than about everything a satellite flies through.

At complete CAM failure the reported $\mathbb{E}[x_F]$ is a lower bound. The conditioning scheme of section 5.2.3 forces exactly one collision into the seeded step, drawn from the categorical distribution over pairs at that step; everywhere else in the run, collisions are drawn with their full Poisson multiplicity, so a step can and does produce more than one. At a five-day step and CAM failure of unity, the seeded step's own intensity Λ_k reaches 0.56, high enough that, had it not been forced, a second collision in that same step would have been a real possibility; because it is forced instead, that possibility is discarded, and the row is daggered in table 8.14 for it. This scenario also carries $\Lambda = 170$ over the whole horizon, which is a second, distinct problem: the conditioning scheme exists to make a rare event, one collision out of an otherwise near-empty run, tractable to sample, and 170 expected collisions is not a rare event. Hence, the scenario has moved outside the regime the scheme was built for and only a finer time step would recover the multiplicity the seeded step is missing.

8.4.2. Impact of Different PMD Success Rates

Figure 8.7 and table 8.15 sweep the PMD failure probability from perfect disposal to one satellite in two.

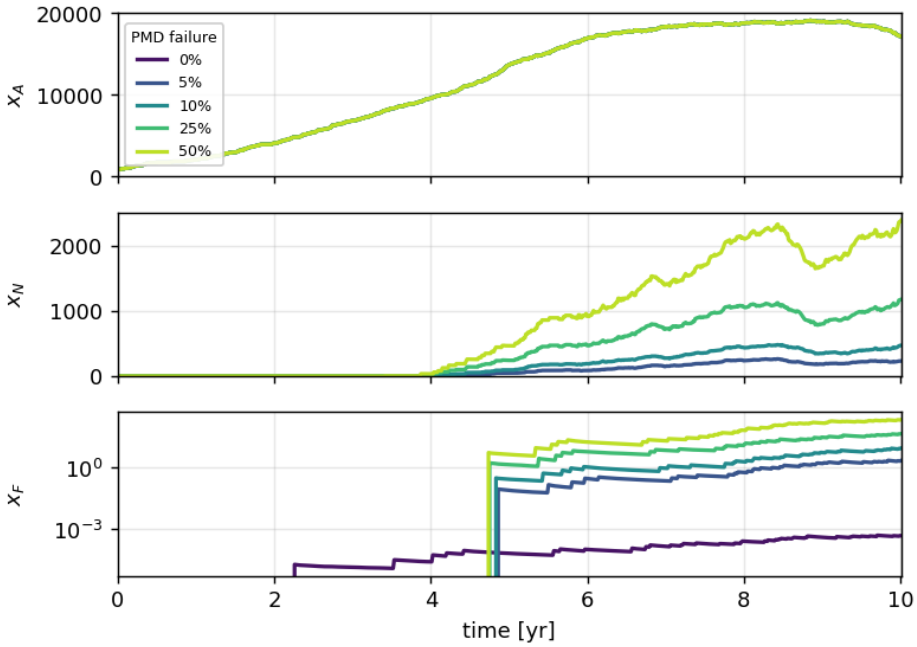

Figure 8.7: Starlink PMD failure sweep: active, derelict and fragment populations over the horizon.

Table 8.15: Starlink pmd failure sweep.

PMD failure	Ref. A	Ref. N	Λ	$\mathbb{E}[x_F]$ at horizon	ϱ_{crit}
0	17 147	0	1.68×10^{-6}	0.0	9781.29
0.05	17 148	239	8.32×10^{-3}	2.4	13.07
0.1	17 148	471	2.90×10^{-2}	9.5	4.21
0.25	17 148	1197	1.69×10^{-1}	48.3	0.67
0.5	17 148	2411	6.98×10^{-1}	233.3	0.25

The active-population panel of fig. 8.7 is unaffected by the sweep: PMD failure changes what happens to a satellite once it reaches end of life, not how many are launched, so every curve tracks the same launch ramp. The derelict panel is where the sweep acts directly. Every curve is flat until around year four, when the first satellites reach end of life, and separates from there in the failure ordering, exactly as table 8.15 reports at the horizon. The reference derelict population scales exactly linearly with the failure probability, 239 objects at 5 % against 471 at 10 %, 1197 at 25 % and 2411 at 50 %, which is the expected result for a fixed number of retirements. The expected collision count grows much faster than that: a ten-fold increase in derelicts, from 5 % to 50 % failure, raises Λ from 8.32×10^{-3} to 0.698, a factor 84. That superlinear growth is consistent with the channel structure of table 8.5: if the risk were carried by anything other than derelict-derelict pairs, which scale with the square of that population, it would not accelerate like this as the failure probability rises.

The fragment panel of fig. 8.7 has the same switch-on structure. For every failure probability above zero the curve is flat until roughly year five, then steps up sharply and settles into a slow climb; the failure levels differ mainly in how high the step lands, not in when it happens. The step is where the ensemble mean crosses from almost no seed having produced a collision yet to a growing fraction having done so, which tracks the derelict-derelict channel's rate leaving zero once the first satellites reach end of life, a timing set by satellite lifetime rather than by the failure probability. At zero failure no derelict is ever created, so this channel never exists, and the only one left is active-active, the pairing CAM suppresses to a baseline success of 0.9999 but

does not eliminate. That channel is present from the first launches rather than waiting on a derelict population to build up, which is why its curve starts climbing around year two instead of five, and being far more strongly suppressed than an unresisted derelict-derelict pairing, it accumulates roughly three orders of magnitude fewer fragments over the same horizon.

With no failed disposals at all, Λ falls to 1.68×10^{-6}, five thousand times below the baseline, the fragment count to zero, and the operating shell's capacity margin rises to 9781. Remarkably, a constellation that disposes of every satellite is, on this model's terms, essentially free of endogenous collision risk. The threshold is correspondingly close. The operating shell's ϱ_{crit} falls from 13.07 at the baseline 5 % to 4.21 at 10 % and crosses unity between 10 % and 25 %. Set against the previous subsection, driving the shell supercritical takes a roughly thousand-fold degradation in avoidance performance, or a threefold degradation in disposal reliability. Measured by distance from current practice to the threshold, disposal is by a wide margin the more fragile of the two levers.

8.4.3. Impact of Different Launch Strategies

Figure 8.8 and table 8.16 replace the baseline one-to-one launch model with three alternatives: no future launches at all, a ramp up to the target population followed by replenishment, and unbounded growth.

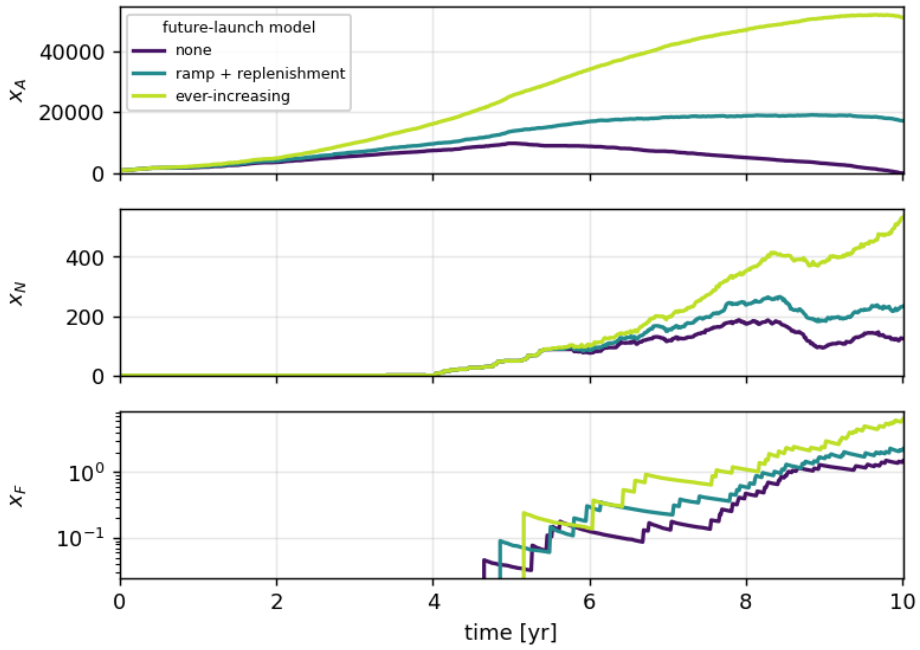

Figure 8.8: Starlink launch-model sweep: active, derelict and fragment populations over the horizon.

Table 8.16: Starlink launch model sweep.

Launch model	Ref. A	Ref. N	Λ	$\mathbb{E}[x_F]$ at horizon	ϱ_{crit}
none	0	127	4.21×10^{-3}	1.6	63.87
ramp + replenishment	17 148	239	8.32×10^{-3}	2.4	13.07
ever-increasing	51 012	542	2.15×10^{-2}	6.7	3.68

Here the active panel is the discriminator: with no future launches the fleet decays to nothing by the horizon, the baseline ramp-plus-replenishment policy plateaus at 17 148, and unbounded growth reaches 51 012, three times the plateau. The derelict panel opens around year four in all

three curves, the same switch-on seen in the two sweeps above, and separates from there in the same ordering as the active fleet, with one exception: the no-launch derelict count peaks around year eight and then falls back, because with no new satellites left to fail disposal, natural decay removes existing derelicts faster than the empty fleet can add them. The fragment panel steps up at nearly the same time in every curve, around year five, for the same reason it did in the disposal sweep: it is tracking when the derelict-derelict channel switches on, not when each scenario's active fleet grows. The expected collision count instead moves from 4.21×10^{-3} to 2.15×10^{-2} across the whole sweep, a factor five, and the unbounded case, with three times the baseline's fleet, carries only 2.6 times its risk.

The reason is in the derelict column. The unbounded-growth case flies three times the satellites but holds only 542 derelicts against the baseline's 239, a factor 2.3, because a fleet that is still growing is a young fleet and comparatively few of its satellites have yet reached end of life. Since the risk is carried by the derelict population, it tracks the number of satellites that have already retired and failed disposal, not the number currently flying. Within a mission horizon, aggressive deployment is therefore much less dangerous than it looks.

The no-launch case makes the same point from the other direction. At the horizon that scenario has no operational satellites whatsoever, and it still carries $\Lambda = 4.21 \times 10^{-3}$, half the baseline, from the 127 derelicts it produced before stopping. Its capacity margin recovers from 13.07 to 63.87, so the shell is substantially better off, but it is not cleared: ceasing operations does not retire the risk, because the risk is in what was already left behind. This study is also the only one of the three whose entire sweep range stays subcritical, with $\varrho_{\mathrm{crit}} = 3.68$ even under unbounded growth.

Taken together the three sweeps rank the levers for this scenario: disposal reliability dominates, avoidance performance matters only once it degrades by orders of magnitude, and launch volume is the weakest of the three within a mission horizon. Disposal comes first because it sets the size of the one population avoidance cannot touch. A satellite that fails disposal becomes a derelict, and a collision between two derelicts has no active endpoint for a maneuver to remove. The primacy of disposal is not particular to PRISM. J.-C. Liou et al. (2018) put the same question to large constellations specifically, and reach the same answer from century-scale LEGEND projections of three of them against a background population, recommending a disposal success rate of at least 99 %. All three statements are properties of a constellation whose risk is carried by its own derelicts, which is what section 8.2 and table 8.5 established independently, so the ranking should be expected to hold for any population with the same channel structure. They are also statements about that constellation in isolation: as section 4.1 noted, the model excludes rocket bodies and the background population, so every ϱ_{crit} and every threshold in this chapter measures a constellation's stability against its own debris alone. Adding collisions with objects outside the constellation would only add risk, never remove it, so a shell that appears comfortably subcritical here would cross into supercritical sooner once that background is accounted for. The debris the constellation makes itself is truncated in the same direction, since a breakup produces far more fragments below 10 cm than above it and only the tracked part is carried here. Both omissions push the same way, which makes every threshold in this chapter an optimistic one.

9

Conclusion

This chapter closes the thesis by consolidating the contributions of the preceding chapters and answering the research questions posed in chapter 1. Section 9.1 summarizes PRISM, its sensitivities and results, and answers the four sub-questions and the main research question in turn. Section 9.2 identifies limitations and open problems, organized by their nature and the effort they would require.

9.1. Conclusions

The contribution of this thesis is PRISM, a stochastic model that evolves the population of a satellite constellation and assesses the endogenous collision risk that population carries. Chapter 4 builds the formulation on the network architecture of NESSY (Wang et al., 2025). The major difference from their work, is that here the nodes are the bins of a lattice defined jointly by altitude, inclination, and RAAN, so that the risk can be attributed to a specific plane rather than averaged over the shell. Individual objects are advanced through the lattice and the population evolution is recovered by averaging over realizations in a Monte Carlo fashion, so the model is not a source-sink model in the strict sense. The mean-field description enters only as a reduced deterministic companion, used in chapter 7 to define equilibrium and a carrying-capacity margin. Chapter 5 implements the formulation as a simulator and verifies it against MOCAT-MC (Jang et al., 2024) on the full heterogeneous LEO catalogue. Finally, chapter 6 performs a sensitivity analysis, and chapter 8 then applies the model to Starlink, OneWeb, and Iridium-NEXT. There, the plane resolution shows its operational value: over the mission-length horizons an operator plans against, neither a derelict population nor a freshly generated fragment cloud has had time to spread uniformly in RAAN, so the risk stays concentrated on a small number of planes, which the model succesfully identifies.

The four sub-questions posed in section 1.3 are answered in turn below, followed by the main research question.

Q1 How can traditional evolutionary models be adapted for constellation-only populations?

Traditional evolutionary models operate on century-scale scenarios of the whole heterogeneous LEO population, where collisions are comparatively common and are driven mostly by the vast amount of debris. That setting licenses two of their standard devices: averaging the angular orbital elements within a bin, which the secular precession of RAAN under J_2 justifies over a century, and brute-force sampling of the stochastic events. Neither survives the move to a constellation-only population, which is small, station-kept, deliberately phased, and observed over a mission-length horizon. Collisions are rare, at an

expected count of 8.3×10^{-3} over ten years for the Starlink baseline of chapter 8, so a plain Monte Carlo would need tens to hundreds of thousands of runs before returning reliable estimates; and a decade is not long enough for nodal precession to fill a shell, so averaging the planes hides the source of risk. As quantified in section 6.2.3 the RAAN-blind approach of Wang et al. (2025) returns one averaged relative velocity of 8.57 km/s where the true value runs from zero to 12.12 km/s with RAAN separation, and it misclassifies 78 % to 86 % of collisions at the catastrophic threshold, against 2.9 % to 5.2 % for the resolved lattice. On the real constellations of section 8.2, coarsening the RAAN intervals from 5° to 60° moves the expected collision count in an unpredictable direction by roughly 10 %, upward for OneWeb and downward for Iridium-NEXT.

The two obstacles are addressed by a corresponding pair of modifications. The rarity is handled by the conditioning scheme of section 5.2.3, which forces a collision at a step drawn from the exact first-arrival distribution of the underlying Poisson process and recombines that branch with the collision-free one by the law of total expectation. The geometry is handled by the collision rate of section 4.2, where the rate between two nodes is the geometric overlap of the two finite orbital bands they represent, evaluated in the (h, i, Ω) phase space with only the argument of latitude averaged out. This method, which is developed from the traditional gas-kinetic approach, introduces two novelties: the RAAN dimension, and the fact that the quantity computed is a rate between two populations, whereas the flux-based methods of Klinkrad (1993) and its descendants reviewed in section 2.3.2 compute the flux of a field population onto one designated target orbit. The species taxonomy of section 4.1 additionally carries a rising and a de-orbiting specie, so that shell-crossing traffic from orbit raising and disposal enters the model explicitly.

Q2 Which stochastic processes need to be modeled to accurately capture the constellation's evolution?

Two stochastic processes are modeled in this work, PMD failures and collisions, and the policy sweeps of section 8.4 show that they are also the two that carry the risk. A satellite that fails disposal becomes a derelict that can no longer maneuver, and a collision between two derelicts is the one channel no avoidance policy can suppress. The expected collision count consequently grows superlinearly with the disposal failure probability, and falls five thousand times below the baseline at perfect disposal. Conversely, avoidance performance has to degrade by orders of magnitude before it registers at all, and launch volume is weaker still, because the risk tracks the satellites that have already retired and failed disposal rather than those currently flying. This ordering agrees with Lewis et al. (2017), who identify conjunctions involving failed satellites as the dominant long-term self-induced risk driver in a constellation. Where those derelicts sit matters as much as how many there are. Derelicts are continuously supplied at their parent shell and drift down from there, forming a band that widens over time. At low altitudes drag spreads the band and eventually clears it, while at high ones it does neither.

Collisions are rare enough that a baseline run essentially never produces one, but they carry the largest consequence of any process modeled. A fresh cloud is born as a clustered object, with 88 % to 99.9 % of its fragments near the parent plane, and it spreads slowly, so it does not become uniform within the horizon over which the constellation is operated. The hazard is correspondingly front-loaded, losing its rate contribution over a decade as the cloud spreads and thins. What the model has to carry is therefore the time-resolved evolution of the cloud in altitude and RAAN.

Two classes of process are deliberately excluded. Collisions with objects outside the constellation, rocket bodies and the background population, fall outside an endogenous risk assessment, though nothing in the formulation prevents adding them to the initial population and evolving them alongside the constellation. Their omission means the thresholds reported here measure a constellation's stability against its own debris alone, and accounting for the background can only move a shell toward instability, never away from it. Random in-orbit failures and explosions are likewise left out for simplicity, and no statement about their contribution is made here, since none was measured.

Q3 How does the network formalism compare to traditional models in accuracy and computational efficiency, and what unique constellation-level insights does the network approach provide?

On the full heterogeneous LEO catalogue of section 5.3, PRISM reproduces MOCAT-MC 's century-long trajectory within one standard deviation, at a fraction of the computational cost, approximately 4 times faster, though the implementation is not yet optimized as far as it could be. The two reach that agreement by structurally unrelated routes, MOCAT-MC through the CUBE method and PRISM through the rate derived in chapter 4.

The largest computational gain is seen on the constellation-only scenarios this thesis targets. The CUBE method can register a collision between two planes whose orbits never intersect, a failure mode Lewis et al. (2019) document for constellation geometries, and avoiding it forces the cube size, and with it the time step, down far enough to resolve the true separation, raising the cost of a run. PRISM needs no such correction: once the altitude bin is narrower than the shell separation, two objects that cannot occupy the same orbital band are never assigned a collision rate, and the time step carries only the negligible aliasing error quantified in chapter 6. More generally, PRISM never compares every object against every other, as Orbit-Trace does, nor every object against every grid cell it could occupy, as CUBE and Klinkrad's (1993) method do. Instead, it evaluates a rate between a small, fixed number of node pairs, and samples a discrete event only when one occurs, which for a constellation-only population is rare. Consequently, simulating the full LEO environment evolution over 100-years takes approximately 20 minutes across 8 CPUs, compared to the 2 hours taken by MOCAT-MC , but the cost decreases considerably for shorter simulations and smaller populations, such as the constellations of chapter 8. By comparison Reiland et al. (2021) report 13 h and 40 h on 200 CPUs for a single pairwise assessment of OneWeb and Starlink, and NEODEEM (Sugita et al., 2021) consumes over 560 000 core-hours per study. On the other hand, the gain narrows for a heterogeneous population, as collisions occur often enough that the running costs of the two families converge. Furthermore, PRISM still averages out the in-plane phasing that the other methods could in principle retain, though they randomize the mean anomaly anyway, so recovering accurate phase information would need a different approach on either side.

The main contribution of the network formalism of chapter 7 is the same as the one identified by Wang et al. (2025), i.e., building the evolutionary model from the network definition itself rather than layering a network representation on a pre-existing model as post-processing. However, it should not be oversold as a source of topological insight, since the network is a geometric lattice whose links are set by known orbital dynamics, not a social network: collision links never cross between shells and are close to complete within one, so classical unweighted centrality measures would only recover a shell's own occupancy pattern. Regardless, rate-weighted metrics allow to rank nodes' risk based on their collision or transport strengths. Separating out-strength from in-strength turns the same graph into a flow budget where it is possible to identify sources and sinks, while the ratio of a node's fragment to collision strength separates how often it collides from how much debris it produces. Additionally, aggregating the same links by RAAN separation isolates the 180° seam of a Walker-star pattern. The network also admits a mean-field limit, in the manner of a source-sink model, from which the carrying-capacity margin of a shell is defined. This remains a proof of concept, since the links are rebuilt on demand from a run's recorded populations rather than accumulated during the run, which rests on simplifying assumptions of its own.

Q4 How well does the stochastic dynamic model generalize to different constellation geometries and mission profiles?

PRISM has a known degeneracy for a population concentrated in a single, thin altitude band, and refining the lattice around a delta-distributed population makes the rate diverge, as quantified in section 6.2.1. A business-as-usual scenario, in which every satellite operates as designed and stays confined to its assigned shell, is exactly such a population, so its results would be unreliable. However, those scenarios carry virtually no risk to begin

with, since a healthy constellation performs collision avoidance, suppressing the collision rate, and the degeneracy lifts as soon as the population becomes heterogeneous in altitude, through staggered failures or a fragmentation event, which is where the model is accurate. Rarity is handled cleanly by the conditioning scheme regardless of geometry, but the model is sensitive to noise for a sparsely populated constellation, and behaves best for a denser mega-constellation.

PRISM was exercised on four idealized Walker-constellation archetypes chosen to span the relevant parameter space: a large single-shell constellation, a multi-shell one with shells only kilometers apart, a small polar constellation, and a high-altitude polar one (chapter 6), as well as on the real Starlink, OneWeb, and Iridium-NEXT constellations built from DISCOS data (chapter 8). The sensitivity analysis exposed how the model's parameters are not arbitrary and need to be tuned to the specific constellation's geometry. Altitude matters because it sets how fast drag pulls a dead satellite down and out of its shell: below roughly 600 km the resulting derelict band forms within ten years no matter how draggy the satellite is, while above that altitude it barely moves, deepening by under 3 km a decade, which is precisely the thin, undeveloped layer that triggers the degeneracy above. Inclination matters differently: two orbital planes cross each other at an angle that depends on how far each sits from the pole, and that dependence is weakest exactly at the pole, which is why a coarse angular lattice costs the two mid-inclination archetypes 11.6 % and 10.4 % in relative-velocity error against only 3.2 % and 5.4 % for the two polar ones, at the same bin count. A multi-shell constellation adds one more requirement on top: the altitude bins have to be narrower than the spacing between its shells, or the model starts merging shells that are meant to stay apart.

The model's scope is otherwise bounded by the perturbations it implements, drag and J_2, chosen because they dominate in LEO. Higher-order gravity harmonics, third-body perturbations, and solar radiation pressure, relevant at higher altitudes or other orbital regimes, are not modeled. Although the modular propagator interface would allow a different one to be substituted, doing so is expected to take considerable effort, and no conclusion about that regime can be drawn here.

Taken together, these four answers settle the main research question posed in section 1.3:

To what extent can evolutionary space environment models assess long-term endogenous collision risk in satellite constellations?

They can, on the condition of resolving the individual orbital planes of a constellation and not merely the altitude shells those planes occupy. Evolutionary models were built for a different question, the fate of the whole LEO population over a century, and the two devices that make that question tractable are averaging over the angular elements and sampling collisions by brute force. Both remove precisely what a constellation study needs, namely which planes interact and what happens in the tail of a distribution whose mean is close to zero. This thesis replaced them with a RAAN-resolved collision rate and a conditioning scheme, and showed that the resulting model still evolves the full LEO population in agreement with MOCAT-MC over a century. With those two substitutions, an evolutionary model assesses endogenous risk to the extent of reporting an expected collision count together with the plane, shell, and specie interactions that carry it, and a margin against the maximum sustainable capacity of the shells the constellation occupies. Chapter 8 shows those quantities to be strongly non-uniform, with risk concentrated on a small number of planes, so the resolution is what makes the assessment informative.

The cost of that resolution is that the lattice has to be tuned for each constellation, with more care than an evolutionary model normally demands. Refining it does not necessarily buy accuracy, and the bin widths have to be derived from the constellation itself, from its shell spacing, its inclinations and the drag-driven spread of the derelict population it will eventually produce. The resulting error is nonetheless bounded, and bounded before any run is performed, since chapter 6 sizes it from those same mission parameters. What this thesis demonstrates

is that a model built on that basis is tractable, that it agrees with established evolutionary models while costing no more to run, and that on a single constellation it returns the plane-level attribution those models cannot produce at all.

9.2. Recommendations for Future Work

The recommendations that follow are ordered by what each would demand of the model. The first five are processes the present formulation could already carry: they need no new theory, only a larger population, further bookkeeping, or a different propagator. The next two are assumptions built into the lattice and the collision rate themselves, which cannot be relaxed without deriving eq. (4.22) again or solving it numerically, and are correspondingly expensive. The last three concern the implementation and its testing rather than the model: closing a known discrepancy against MOCAT-MC, testing one of the predictions against observation, and removing the single quantitative caveat the network analysis carries.

Inclusion of the background population The risk assessed here is endogenous, so collisions with rocket bodies, background debris, and satellites belonging to other operators are excluded. Nothing in the formulation prevents them: the species taxonomy would need a class for objects that are neither constellation payloads nor their fragments, and those objects would then enter the initial population and evolve on the same lattice alongside the constellation. This is the extension that would most change the numbers reported here, which are a floor on what a constellation faces in a shared environment, and it would also allow the carrying-capacity margin of a shell to be evaluated against its actual occupancy rather than against one operator's share of it.

Random failures and explosions A satellite becomes non-maneuverable in this model only by failing PMD at end of life. In reality it can fail at any point during its operational lifetime, and stored energy in batteries or residual propellant can cause it to explode. Both are Poisson processes acting on node populations the simulator already samples, so they are inexpensive to add. What they contribute is not known in advance, since a random failure releases a derelict at an operational altitude earlier than a disposal failure does.

More realistic maneuvers Rising and de-orbiting satellites are modeled, but their contribution to risk is not studied in detail here, on the grounds that they retain the ability to maneuver and therefore take part in virtually no collisions. Polli et al. (2023) nonetheless identify shell-crossing traffic as a meaningful source of risk, which makes it worth examining on its own terms. Beyond that, satellites are currently injected directly at their target RAAN and inclination and only change altitude, which already improves on traditional models that inject and remove them at the target orbit instantaneously. A more realistic profile would add a drifting phase at an intermediate parking orbit, which is how operators normally reach their assigned slot after release from a rideshare launch.

Extension of the environment The immediate step is to study the effect of different atmospheric models, since the static, layered-exponential atmosphere used here is the one structural difference that could not be eliminated in the verification of section 5.3. In the longer term, modeling the perturbations that dominate outside LEO would extend the model to other orbital regimes, which requires substituting the propagator.

Adding a mass budget Mass currently affects only drag and breakup. In reality every maneuver consumes propellant and changes a satellite's mass, and the budget is finite: a Starlink satellite carries enough for approximately 350 CAMs (SpaceX, 2022). Tracking it would let the

model represent a satellite that exhausts its propellant and becomes non-maneuverable before its nominal end of life, which is presently outside its scope.

Inclusion of in-plane phasing The model averages out the mean anomaly and is consequently blind to phasing. Including it would require either reformulating the collision rate or solving eq. (4.22) numerically, both at considerable cost. The approximation rests on the assumption that a mission designer spaces active satellites sensibly to maximize their separation, so a phase-aware model could detect a poor design that the present one cannot. More significantly, it would allow the chain reaction described in chapter 1 to be modeled explicitly: an individual CAM changes the phasing between neighbouring satellites and propagates through the constellation, whereas here CAM enters only as a suppressing factor on the collision intensity.

Inclusion of eccentricity Eccentricity currently enters only through the occupancy model, which weights a node's population by the time each object spends at the corresponding altitude over one orbital period. This keeps the definition of a node as an altitude shell clean, but introduces inaccuracies in both the relative velocity and the transport links, which assume circular velocities. Adding an eccentricity direction to the lattice would remove those assumptions and is theoretically possible on the basis of Giudici et al. (2023), at the cost of reformulating the collision rate.

Reconciling the breakup mass budget The comparison of section 5.3 leaves one discrepancy attributable to the implementation rather than to a modeling choice. Both models draw a fragment count from the same size distribution and assign each fragment a mass, but MOCAT-MC then closes the budget explicitly, adding remnant objects to absorb a shortfall or removing the largest fragments on an overshoot (Jang et al., 2024), whereas this implementation takes the fragments as drawn. Closing the budget the same way would remove one of the two known differences between the runs and allow the remaining one, the atmosphere, to be isolated.

Validation of the fragment cloud The model is verified against another model rather than against observation. A full validation against a real constellation's history is of limited value, for the reasons given in section 5.3, but one component admits an observational test on its own. The fragment cloud spreading of chapter 8, which drives the front-loaded hazard the results rest on, predicts a growth in altitude and RAAN extent that can be compared directly against the tracked debris from the Iridium 33 and Cosmos 2251 collision of February 2009 (Anz-Meador & Liou, 2010), the only observed catastrophic satellite-to-satellite fragmentation in LEO.

Exact network links on a conditioned run The network of chapter 7 is rebuilt on demand from a recorded epoch's node populations rather than accumulated during the run, a choice made because accumulating every pair rate would defeat the aggregation that makes the collision sampler efficient. The cost of that choice is visible in chapter 8, where the per-plane rate of the conditional branch had to be estimated from mean-field links evaluated on the conditional population, leaving its structure reliable but its absolute magnitude not. Recovering exact links would remove the one quantitative caveat carried by the network results.

References

Acciarini, G., & Vasile, M. (2020). A Multi-layer Temporal Network Model of the Space Environment. *Proceedings of the 71st International Astronautical Congress*.

Acciarini, G., & Vasile, M. (2021). A Network-based Evolutionary Model of the Space Environment. *Proceedings of the 8th European Conference on Space Debris*.

Alfano, S. (2005). A Numerical Implementation of Spherical Object Coliision Probability. *The Journal of the Astronautical Sciences*, *53*(1), 103–109.

Anz-Meador, P., & Liou, J.-C. (2010). Analysis and Consequences of the Iridium 33-Cosmos 2251 Collision. *Proceedings of the 38th COSPAR Scientific Assembly*, 3.

Arroyo-Parejo, A., Sánchez-Ortiz, N., & Ul Domínguez-González, R. (2021). *Effect of Mega-constellations on Collision Risk in Space* (tech. rep.). http://conference.sdo.esoc.esa.int,

Bendisch, J., Bunte, K., Klinkrad, H., Krag, H., Martin, C., Sdunnus, H., Walker, R., Wegener, P., & Wiedemann, C. (2004). The MASTER-2001 model. *Advances in Space Research*, *34*(5), 959–968. https://doi.org/10.1016/j.asr.2003.02.026

Blue Origin LLC. (2026, January). *Application File No. SAT-LOA-20260120-00033* (tech. rep.). Blue Origin LLC. https://cdn.geekwire.com/wp-content/uploads/2026/01/Terawave-Narrative.pdf

Brouwer, D. (1959). Solution of the problem of artificial satellite theory without drag. *The Astronomical Journal*, *64*, 378–396. https://doi.org/10.1086/107958

Chen, Y., Li, Y., Li, H., Liu, L., Ouyang, L., Yang, J., Li, J., Wu, J., Wu, Q., Liu, J., & Lai, Z. (2024). Instability of Self-Driving Satellite Mega-Constellation: From Theory to Practical Impacts on Network Lifetime and Capacity. http://arxiv.org/abs/2406.06068

Colombo, C., Letizia, F., & Lewis, H. G. (2016). Spatial Density Approach for Modeling of the Space Debris Population. *Proceedings of the 26th AAS/AIAA Space Flight Mechanics Meeting*.

D'Ambrosio, A., Servadio, S., Siew, P. M., Linares, R., & McGrath, C. N. (2023). Novel Source–Sink Model for Space Environment Evolution with Orbit Capacity Assessment. *Journal of Spacecraft and Rockets*, *60*(4), 1112–1126. https://doi.org/10.2514/1.A35579

Di Costanzo, R., Dolado-Perez, J., & Revelin, B. (2013). Introducing MEDEE - A New Orbital Debris Evolutionary Model. *Proceedings of the 6th International Conference on Astrodynamics Tools and Techniques*, 22–25.

ESA Space Debris Office. (2025, October). *ESA's Annual Space Environment Report* (tech. rep.). European Space Agency. https://www.esa.int/Space_Safety/Space_Debris/ESA_Space_Environment_Report_2025

Fanto, P., Liang, C., & Signoracci, A. (2023). A Stochastic Source-Sink Framework for Orbital Debris. *Proceedings of the 2nd Orbital Debris Conf. Papers*.

FCC. (2023, April). *DA 23-362 Order and Declaratory Ruling* (tech. rep.). Federal Communications Commision. https://docs.fcc.gov/public/attachments/DA-23-362A1.pdf

FCC. (2026a, January). *DA 26-36 Authorization and Order* (tech. rep.). Federal Communications Commission. https://docs.fcc.gov/public/attachments/DA-26-36A1.pdf

FCC. (2026b, February). *GRANT Kuiper Systems LLC* (tech. rep.). Federal Communications Commission. https://cdn.geekwire.com/wp-content/uploads/2026/02/AmazonLeo-Gen2.pdf

Frey, S., Colombo, C., & Lemmens, S. (2019). Application of density-based propagation to fragment clouds using the Starling suite. *First International Orbital Debris Conference*.

Frey, S., & Colombo, C. (2021). Transformation of satellite breakup distribution for probabilistic orbital collision hazard analysis. *Journal of Guidance, Control, and Dynamics*, *44*(1), 88–105. https://doi.org/10.2514/1.G004939

Giudici, L. (2021). *Continuum Formulation for Modelling the Impact of Fragmentation Events in Low Earth Orbit* [Doctoral dissertation, Politecnico di Milano].

Giudici, L., Colombo, C., Horstmann, A., Letizia, F., & Lemmens, S. (2024). Density-based evolutionary model of the space debris environment in low-Earth orbit. *Acta Astronautica*, *219*, 115–127. https://doi.org/10.1016/j.actaastro.2024.03.008

Giudici, L., Gonzalo, J. L., & Colombo, C. (2023). Density-based in-orbit collision risk model extension to any impact geometry. http://arxiv.org/abs/2309.03562

Gusmini, D., D'ambrosio, A., Servadio, S., Siew, P. M., Di Lizia, P., & Linares, R. (2024). Effects of Orbit Raising and Deorbiting in Source-Sink Evolutionary Models. *Journal of Spacecraft and Rockets*, *61*(3), 784–797. https://doi.org/10.2514/1.A35849

Iridium Satellite LLC. (2013, December). *Iridium NEXT Engineering Statement* (tech. rep.).

Jang, D., Gusmini, D., Siew, P. M., D'Ambrosio, A., Servadio, S., Machuca, P., & Linares, R. (2024). A New Monte-Carlo Model for the Space Environment. http://arxiv.org/abs/2405.10430

JeongAhn, Y., & Malhotra, R. (2017). Simplified derivation of the collision probability of two objects in independent Keplerian orbits. https://doi.org/10.3847/1538-3881/aa6aa7

Johnson, N. L., Krisko, P. H., Liou, J.-C., & Anz-Meador, P. D. (2001). Nasa's new breakup model of evolve 4.0. *Advancements in Space Research*, *28*, 1377–1384. https://doi.org/10.1016/S0273-1177(01)00423-9

Jones, A. (2026, January). China files ITU paperwork for megaconstellations totaling nearly 200,000 satellites. https://spacenews.com/china-files-itu-paperwork-for-megaconstellations-totaling-nearly-200000-satellites/

Kawamoto, S., Hirai, T., Kitajima, S., Abe, S., & Hanada, T. (2018). Evaluation of Space Debris Mitigation Measures Using a Debris Evolutionary Model. *Transactions of the Japan Society for Aeronautical and Space Sciences*, *16*(7), 599–603. https://doi.org/10.2322/tastj.16.599

Kawamoto, S., Nagaoka, N., Kitagawa, Y., & Hanada, T. (2021). Evaluation of Impacts of Large Constellations Using Debris Evolutionary Model for Considering Environment Capacity. *Proceedings of the 8th European Conference on Space Debris*.

Kawamoto, S., Nagaoka, N., Sato, T., & Hanada, T. (2020). Impact on collision probability by post mission disposal and active debris removal. *Journal of Space Safety Engineering*, *7*(3), 178–191. https://doi.org/10.1016/j.jsse.2020.07.012

Kerr, E., & Sánchez Ortiz, N. (2021). State of the Art and Future Needs in Conjunction Analysis Methods, Processes and Software. In T. Flohrer, S. Lemmens, & F. Schmitz (Eds.), *Proceedings of the 8th european conference on space debris*. ESA Space Debris Office.

Kessler, D. J. (1981). *Derivation of the Collision Probability between Orbiting Objects: The Lifetimes of Jupiter's Outer Moons* (tech. rep.).

Kessler, D. J., & Cour-Palais, B. G. (1978). Collision Frequency of Artificial Satellites: the Creation of a Debris Belt. *J Geophys Res*, *83*(A6), 2637–2646. https://doi.org/10.1029/JA083iA06p02637

Klinkrad, H. (1993). Collision Risk Analysis for Low Earth Orbits. *Advances in Space Research*, *13*(8), 177. https://doi.org/10.1016/0273-1177(93)90588-3

Klinkrad, H., Bendisch, J., Bunte, K. D., Krag, H., Sdunnus, H., & Wegene, P. (2001). The MASTER-99 Space Debris and Meteoroid Environment Model. *Advances in Space Research*, *28*(9), 1355–1366. https://doi.org/10.1016/S0273-1177(01)00416-1

Klinkrad, H. (2006). *Space Debris - Models and Risk Analysis*. Springer.

Krisko, P. (2011). Proper implementation of the 1998 nasa breakup model. *NASA Orbital Debris Quarterly News*, *15*, 4–5. www.orbitaldebris.jsc.nasa.gov

Letizia, F., Colombo, C., & Lewis, H. G. (2015). Analytical model for the propagation of small-debris-object clouds after fragmentations. *Journal of Guidance, Control, and Dynamics*, *38*(8), 1478–1491. https://doi.org/10.2514/1.G000695

Letizia, F., Colombo, C., & Lewis, H. G. (2016). Multidimensional extension of the continuity equation method for debris clouds evolution. *Advances in Space Research*, *57*(8), 1624–1640. https://doi.org/10.1016/j.asr.2015.11.035

Letizia, F., Colombo, C., Lewis, H. G., & Krag, H. (2016). Assessment of breakup severity on operational satellites. *Advances in Space Research*, *58*(7), 1255–1274. https://doi.org/10.1016/j.asr.2016.05.036

Lewis, H. G., Diserens, S., Maclay, T., & Sheehan, J. P. (2019). Limitations of the cube method for assessing large constellations. *First International Debris Conference*.

Lewis, H. G., Newland, R. J., Swinerd, G. G., & Saunders, A. (2010). A new analysis of debris mitigation and removal using networks. *Acta Astronautica*, *66*(1-2), 257–268. https://doi.org/10.1016/j.actaastro.2009.05.010

Lewis, H. G., Radtke, J., Beck, J., Bastida Virgili, B., & Krag, H. (2017). Self-induced Collision Risk Analysis for Large Constellations. *Proceedings of the 7th European Conference on Space Debris*.

Lewis, H. G., Swinerd, G., Williams, N., & Gittings, G. (2001). DAMAGE: A Dedicated GEO Debris Model Framework. *Proceedings of the 3rd European Conference on Space Debris*.

Lin, Q., Li, T., Li, Y., & Luo, Z. F. (2024). Risk Assessment of Satellite Constellations Using Network Theory. *2024 International Annual Conference on Complex Systems and Intelligent Science, CSIS-IAC 2024*, 522–529. https://doi.org/10.1109/CSIS-IAC63491.2024.10919439

Liou, J. C. (2006). Collision activities in the future orbital debris environment. *Advances in Space Research*, *38*(9), 2102–2106. https://doi.org/10.1016/j.asr.2005.06.021

Liou, J. C., Anilkumar, A. K., Bastida Virgili, B., Hanada, T., Krag, H., Lewis, H., Raj, M. X. J., Rao, M. M., Rossi, A., & Sharma, R. K. (2013). Stability of the Future LEO Environment - An IADC Comparison Study. *Proceedings of the 6th International Conference on Astrodynamics Tools and Techniques*.

Liou, J. C., Hall, D. T., Krisko, P. H., & Opiela, J. N. (2004). LEGEND - A three-dimensional LEO-to-GEO debris evolutionary model. *Advances in Space Research*, *34*(5), 981–986. https://doi.org/10.1016/j.asr.2003.02.027

Liou, J. C., Kessler, D. J., Matney, M., & Stansbery, G. (2003). A New Approach to Evaluate Collision Probabilities among Asteroids, Comets, and Kuiper Belt Objects. *Proceedings on the Lunar Planetary Science Conference*.

Liou, J.-C., Matney, M., Vavrin, A., Manis, A., & Gates, D. (2018). NASA ODPO's Large Constellation Study. *Orbital Debris Quarterly News*, *22*(3), 4–7. https://orbitaldebris.jsc.nasa.gov/quarterly-news/pdfs/odqnv22i3.pdf

Lyddane, R. H. (1963). Small eccentricities or inclinations in the Brouwer theory of the artificial satellite. *The Astronomical Journal*, *68*, 555–558. https://doi.org/10.1086/109179

Martinusi, V., Dell'Elce, L., & Kerschen, G. (2015). Analytic propagation of near-circular satellite orbits in the atmosphere of an oblate planet. *Celestial Mechanics and Dynamical Astronomy*, *123*(1), 85–103. https://doi.org/10.1007/s10569-015-9630-7

Matthiesen, B., Razmi, N., Leyva-Mayorga, I., Dekorsy, A., & Popovski, P. (2023). Federated Learning in Satellite Constellations. *IEEE Network*, *38*(2), 232–239. https://doi.org/10.1109/MNET.132.2200504

McInnes, C. R. (1993). An analytical model for the catastrophic production of orbital debris. *ESA Journal*, *17*(4), 293–305.

McInnes, C. R. (2000). Simple analytic model of the long-term evolution of nanosatellite constellations. *Journal of Guidance, Control, and Dynamics*, *23*(2), 332–338. https://doi.org/10.2514/2.4527

Mortari, D., Wilkins, M. P., & Bruccoleri, C. (2004). The Flower Constellations. *The Journal of the Astronautical Sciences*, *52*(2), 107–127.

Öpik, E. J. (1951). Collision Probabilities with the Planets and the Distribution of Interplanetary Matter. *Proceedings of the Royal Irish Academy*.

Owens-Fahrner, N., Wysack, J., & Kim, J. (2025). *Graph-Based Optimization for High-Density LEO Constellation Design* (tech. rep.). www.amostech.com

Patera, R. P. (2001). General method for calculating satellite collision probability. *Journal of Guidance, Control, and Dynamics*, *24*(4), 716–722. https://doi.org/10.2514/2.4771

Polli, E. M., Gonzalo, J. L., & Colombo, C. (2023). Analytical model for collision probability assessments with large satellite constellations. *Advances in Space Research*, *72*(7), 2515–2534. https://doi.org/10.1016/j.asr.2022.07.055

Radtke, J., Flegel, S., Gelhaus, J., Möckel, M., Braun, V., Kebschull, C., Wiedemann, C., Krag, H., Merz, K., & Vörsmann, P. (2013). Revision of Statistical Collision Analysis for Objects Inside of Satellite Constellations. *Proceedings of the 64th International Astronautical Congress*.

Radtke, J., Mueller, S., Schaus, V., & Stoll, E. (2017). LUCA2 - An Enhanced Long-term Utility for Collision Analysis. *Proceedings of the 7th European Conference on Space Debris*.

Reiland, N., Rosengren, A. J., Malhotra, R., & Bombardelli, C. (2021). Assessing and minimizing collisions in satellite mega-constellations. *Advances in Space Research*, *67*(11), 3755–3774. https://doi.org/10.1016/j.asr.2021.01.010

Romano, M., Carletti, T., & Daquin, J. (2024). The Resident Space Objects Network: A Complex System Approach for Shaping Space Sustainability. *Journal of the Astronautical Sciences*, *71*(4). https://doi.org/10.1007/s40295-024-00449-4

Rossi, A., Anselmo, L., Pardini, C., Jehn, R., & Valsecchi, G. B. (2009). The New Space Debris Mitigation (SDM 4.0) Long Term Evolution Code. *Proceedings of the 5th European Conference on Space Debris*.

Rossi, A., Cordelli, A., Farinella, P., & Anselmo, L. (1994). Collisional evolution of the Earth's orbital debris cloud. *Journal of Geophysical Research*, *99*(E11). https://doi.org/10.1029/94je02320

Rossi, A., Cordelli, A., Farinella, P., Anselmo, L., & Pardini, C. (1997). Long Term Evolution of the Space Debris Population. *Advances in Space Research*, *19*(2), 331–340.

Rossi, A., Valsecchi, G. B., & Alessi, E. M. (2015). The criticality of spacecraft index. *Advances in Space Research*, *56*(3), 449–460. https://doi.org/10.1016/j.asr.2015.02.027

Ruan, Y., Hu, M., Yun, C., & Xue, W. (2024). Preliminary Safety Analysis of Megaconstellations in Low Earth Orbit: Assessing Short-Term and Long-Term Collision Risks. *Applied Sciences (Switzerland)*, *14*(7). https://doi.org/10.3390/app14072953

SpaceX. (2022, October). *IBFS File Nos. SAT-LOA-20200526-00055 and SAT-AMD-20210818-00105* (tech. rep.). SpaceX.

Sugita, H., Kawamoto, S., Kitagawa, Y., & Nagaoka, N. (2021). *Long-term Orbital Environment Prediction by Orbital Debris Evolutionary Model* (tech. rep.). JAXA. https://www.jss.jaxa.jp/en/ar/e2020/14349/

Vallado, D. A., & McClain, W. D. (2013). *Fundamentals of Astrodynamics and Applications* (Fourth). Microcosm Press.

Virgili, B. B. (2016). DELTA (Debris Environment Long Term Analysis). *Proceedings of the 6th International Conference on Astrodynamics Tools and Techniques*.

Virgili-Llop, J. (2014, November). *Spacecraft flight in the atmosphere* [Doctoral dissertation]. https://doi.org/10.13140/RG.2.1.5192.8480

Vokrouhlický, D., Pokorný, P., & Nesvorný, D. (2012). Öpik-type collision probability for high-inclination orbits. *Icarus*, *219*(1), 150–160. https://doi.org/10.1016/j.icarus.2012.02.021

Wakker, K. (2015, January). *Fundamentals of astrodynamics*.

Walker, J. G. (1984). Satellite Constellations. *Journal of the British Interplanetary Society*, *37*.

Wang, Y., Marchi, P. D., & Vasile, M. (2025). A stochastic dynamic network model of the space environment. *Advances in Space Research*, *76*(9), 5168–5198. https://doi.org/10.1016/j.asr.2025.08.051

Wertz, J. R., & Larson, W. J. (1999). *Space Mission Analysis and Design* (The Space Technology Library, Ed.; 3rd ed.). Microcosm Press; Kluwer Academic Publishers.

Wetherill, G. W. (1967). Collisions in the asteroid belt. *Journal of Geophysical Research*, *72*(9), 2429–2444. https://doi.org/10.1029/jz072i009p02429

Yates, K. W., & Jonas, F. M. (1995, February). *Assessment of the NASA EVOLVE Long-term Orbital Debris Evolution Model* (tech. rep.). ORION International Technologies, Inc.

Yoon, Y. T., Ghezzo, P., Hervieu, C., & Palomo, I. D. A. (2023, December). Navigating a large satellite constellation in the new space era: An operational perspective. https://doi.org/10.1016/j.jsse.2023.10.006

A

Research Planning

The project ran from February to September 2026, organized into a literature review, five work packages and four fixed milestones. The breakdown below follows the logical order in which the model was built: a software skeleton and a population source first, the physics that advances that population second, the verification that it is implemented correctly third, and the simulation campaigns that produce the results last. Unit testing and report writing are the two exceptions to that sequence, since both run in parallel with everything else.

A.1. Work Package Breakdown

- **Literature review** (2 February to 12 March). Survey the state of the art and establish the gap the thesis addresses, ahead of any implementation. Two bodies of work were reviewed in parallel: the conjunction screening and pairwise probability methods used to assess a constellation over days or weeks, and the evolutionary, source-sink and continuum models used to project a whole population over decades. The objective was to show that neither family answers the question posed here, and to select the framework to be adapted. It delivers chapter 1 and chapter 2, and was scheduled to close before WP1 begins so that the species taxonomy and the lattice the codebase is built around were settled beforehand.
- **WP1: Foundation.** Set up the software infrastructure that every later package builds on, so that the physics could be added one module at a time without the surrounding framework having to change.
  - **WP1.1 Codebase skeleton** (16 March to 6 April). Define the architecture of PRISM described in section 5.1: the structure-of-arrays registry that holds the population, the narrow interfaces behind which the orbital dynamics, the atmosphere and the collision handling sit, the three-phase simulation step, the configuration and result-recording layers, and the worker pool that runs a Monte Carlo ensemble across the cores of a single node.
  - **WP1.2 Constellation data ingestion** (16 to 27 March). Implement the two population sources of section 5.1.2: the idealized Walker-pattern generator, and the retrieval of a real constellation from the DISCOS database, converted into the shared object schema.
- **WP2: Core Physics Model.** Derive and implement the evolutionary model of chapter 4, that is, every mechanism that moves an object between species or between nodes of the orbit-bin lattice. This was the longest package, and its sub-packages were sequenced so that each one could be tested against the model it extended.
  - **WP2.1 Collision rate kernel** (16 March to 22 May). Derive the closed-form collision rate between two nodes of the lattice from the phase-space density, reduce the overlap

integral to a form cheap enough to evaluate over the full lattice, and implement it: the quadrature of the overlap tensor, the exploitation of its circulance in RAAN to generate the whole table of bin pairs from a single slice, the Fourier contraction that makes the aggregate rate affordable at run time, and the relative velocity and average cross-section that weight it. Being both the mathematical core of the model and its dominant computational cost, this sub-package was allocated the largest share of the schedule.

  - **WP2.2 Natural decay and J_2 drift** (13 to 24 April). Implement the analytic secular propagator: the drag-driven decay of the semi-major axis over the layered exponential atmosphere, the secular precession rates induced by Earth's oblateness, and the re-entry criterion that removes an object from the population.
  - **WP2.3 EOL and PMD transitions** (20 to 24 April). Implement the species transitions that do not involve a change of orbit: random in-orbit failures, the end of the operational life, and the attempted post-mission disposal with its prescribed success rate, which turns a compliant satellite into a de-orbiting one and a non-compliant one into a derelict.
  - **WP2.4 Launches** (27 April to 8 May). Inject the scheduled launches of the static manifest into the mutable population as the simulated clock passes their epoch, and implement the replenishment strategies that hold a fleet at its target size over the horizon. This includes the offline fits of the historical launch cadence used to extend a catalog into a manifest that spans the whole simulated horizon.
  - **WP2.5 Orbit maneuvers** (11 to 22 May). Implement the low-thrust transfers that carry a satellite from its injection orbit up to its operational shell and, at disposal, back down through the shells below it, together with the bookkeeping that assigns a maneuvering object to the nodes it occupies during a step. This is the mechanism that produces the shell-crossing traffic.
  - **WP2.6 NASA SBM** (25 May to 10 June). Implement the Standard Breakup Model: the fragment count drawn from the characteristic-length distribution, the sampled area-to-mass ratios and velocity increments, the distinction between catastrophic and non-catastrophic impacts, and the placement of the resulting cloud on the lattice.

- **WP3: Verification and Validation.** Establish that the implementation reproduces both its own theoretical predictions and the output of an independently developed model, and leave the codebase in a state fit for the campaigns of WP4.

  - **WP3.1 Unit tests and sanity checks** (7 April to 19 June). Run alongside WP2 rather than after it, so that each mechanism was tested as soon as it was written: unit tests per module, convergence checks on the quadrature, limiting cases in which the rate has a known analytic value, and comparisons of the sampled event counts against the intensities they were drawn from.
  - **WP3.2 Cross-verification against MOCAT-MC** (11 to 16 June). Reproduce the published scenario of section 5.3, an ensemble of 200 seeds evolving the real LEO population over a century, and compare the cumulative collision counts and the population by species against it. MOCAT-MC reaches its collision counts by a route that shares nothing with the formulation of this model, which is what makes the agreement meaningful.
  - **WP3.3 Code clean-up** (16 to 29 June). Refactor, document and profile the codebase, and freeze the interfaces before the production runs, so that every campaign in WP4 could be attributed to a single version of the model.

- **WP4: Simulation Campaigns and Analyses.** Use the verified model to answer the research questions. The three sub-packages are sequenced by dependency: the sensitivity analysis fixes the settings the later runs use, and the network measures are needed to read the case studies.

  - **WP4.1 Sensitivity analysis** (17 June to 10 July). Carry out the global sensitivity analysis and the discretization study of chapter 6: identify which model parameters drive the outcome, quantify how finely the lattice must be resolved in altitude, inclination and RAAN and how short the time step must be before the answers can be trusted, and measure what a coarse choice costs.
  - **WP4.2 Network topology** (13 to 24 July). Recast the population as the network of chapter 7 and implement its topological and dynamical measures, adjacency matrices, node strength and the derived indicators, which identify the regions that carry the risk and the channels that transport it.
  - **WP4.3 Reference case studies** (27 July to 11 August). Run the campaigns of chapter 8 on Starlink, OneWeb and Iridium-NEXT: the baseline evolution of each constellation over its own horizon, the plane-resolved risk, the carrying capacity of each shell, and the sweeps over the collision-avoidance, disposal and launch policies an operator controls.

- **WP5: Report writing** (2 February to 25 August). Write this report, together with the intermediate deliverables the milestones require. The package spans the entire project rather than occupying a block at the end, so that each chapter was drafted while the work it describes was still fresh, and its last three weeks, after WP4 closes, were reserved for the final revision.

A.2. Milestones and Planning Overview

Four fixed dates anchor the schedule, and the two multi-week packages were placed so that a complete and verified model was available for the first and a complete set of results for the second. The midterm review (23 June) presents the model formulation, the software framework and the cross-verification, that is, the whole of WP1 to WP3, and obtains approval on the modeling choices before committing to the production campaigns; it falls one week after WP3.2 closes so that the verification result was available to present. The greenlight (12 August), one day after WP4.3 closes, demonstrates that the results are complete and the report far enough along to be finished within the remaining two weeks. The final submission (26 August) follows the day after WP5 closes, and the defense (9 September) two weeks later.

Figure A.1 places all of the above on a single timeline, together with the milestones, the non-working periods and the dependencies between sub-packages.

Figure A.1: Gantt chart of the research planning. Bars in blue are work packages and sub-packages, bars in red are the non-working periods, and diamonds are milestones. Arrows indicate the dependencies between packages.

B

Probability Foundations and Stochastic Processes

B.1. Common Definitions

Def B.1.1: Bernoulli Distribution

A *Bernoulli distribution*

$$\text{Bernoulli}(p)$$

describes the outcome of a single trial of a binary random variable that takes the value 1 (success) with probability p and 0 (failure) with probability $1-p$. Its probability mass function is

$$f(x \mid p) = p^x(1-p)^{1-x} \qquad x \in \{0,1\}$$

Its mean and variance are $\mathbb{E}[X] = p$ and $\text{Var}(X) = p(1-p)$.

A Bernoulli distribution can be used to describe the indicator variable of whether two objects fall in the same bin of an occupancy model, or, more generally, any single yes/no trial with a known probability of success, such as tossing a (possibly biased) coin.

Def B.1.2: Discrete Uniform Distribution

A *discrete uniform distribution*

$$\text{Uniform}\{1,\ldots,n\}$$

describes a discrete random variable that can take on one of n equally likely values. Its probability mass function is

$$f(x \mid n) = \frac{1}{n} \qquad x \in \{1,\ldots,n\}$$

A discrete uniform distribution describes the outcome of throwing a fair die, or the selection of one object among the n objects occupying a node, when every occupant is equally likely to be drawn, such as when identifying the two colliding objects once a colliding pair of nodes has been selected.

Def B.1.3: Categorical Distribution

A *categorical distribution*

$$\text{Categorical}(\boldsymbol{p})$$

describes the possible outcomes of a discrete random variable that can take on one of k possible categories, each one with its assigned probability p_i. The probability mass function of a categorical distribution is

$$f(x = i \mid \boldsymbol{p}) = p_i$$

with $\boldsymbol{p} = (p_1, \ldots, p_k)^T$. The individual probabilities are constrained only by the fact that they must be in the range 0 to 1, and they must sum to 1. A categorical distribution with $k = 2$ is a Bernoulli distribution, while one whose all probabilities are equal is a discrete uniform distribution.

A categorical distribution can be used to describe the outcome of throwing an unfair die, for which the probability of landing on any given face is known, or the pair of network nodes for which a collision occurs, given the node-pair probability of collision.

Def B.1.4: Multinomial Distribution

A *multinomial distribution*

$$\text{Multinomial}(n, \boldsymbol{p})$$

describes the possible outcomes of a discrete random variable that can take on one of K possible categories, each one with its assigned probability p_i, for n independent trials. In practice, it describes the outcomes of n independent draws from a categorical distribution. Its probability mass function is

$$f(x_1, \ldots, x_k \mid n, \boldsymbol{p}) = P(X_1 = x_1 \text{ and } \ldots \text{ and } X_k = x_k) = \frac{n!}{x_1! \cdots x_k!} p_1^{x_1} \cdots p_k^{x_k}$$

where x_i are non-negative integers and their sum is n. A multinomial distribution with $k \geq 2$ and $n = 1$ is a categorical distribution.

A multinomial distribution can be used to describe the outcome of throwing an unfair die n times, or all the pair of network nodes for which a collision occurs, given the node-pair probability of collision and the total number of collisions across all pairs.

Def B.1.5: Poisson Distribution

A *Poisson distribution*

$$\text{Poisson}(\Lambda)$$

describes the probability of a given number of events occurring in a fixed interval of time if these events occur with a known constant mean rate Λ and independently of the time since the last event. Its probability mass function is

$$f(k, \Lambda) = P(X = k) = \frac{\Lambda^k e^{-\Lambda}}{k!}$$

For a Poisson distribution the mean and variance are $\mathbb{E}[X] = \text{Var}(X) = \Lambda$.

Given an event whose only possible outcomes are to either occur or not, a Poisson distribution can be used to describe the number of occurrences when the number of trials is (infinitely) large and the probability of the event to occur is (infinitesimally) small, such as the number of collisions in LEO.

Def B.1.6: Continuous Uniform Distribution

A *continuous uniform distribution*

$$\text{Uniform}(a, b)$$

describes a continuous random variable all of whose possible values, within the interval $[a, b]$, are equally likely. Its probability density function is

$$f(x \mid a, b) = \begin{cases} \dfrac{1}{b-a} & a \leq x \leq b \\ 0 & \text{otherwise} \end{cases}$$

A continuous uniform distribution can be used to describe an angular orbital element, such as the RAAN, argument of periapsis, or mean anomaly of an object, when no value within its range is preferred over any other.

Def B.1.7: Normal Distribution

A *normal*, or *Gaussian*, distribution

$$\mathcal{N}(\mu, \sigma^2)$$

describes a continuous random variable clustered symmetrically around a mean μ with spread governed by the standard deviation σ. Its probability density function is

$$f(x \mid \mu, \sigma) = \frac{1}{\sigma\sqrt{2\pi}} \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right)$$

For a normal distribution the mean and variance are $\mathbb{E}[X] = \mu$ and $\text{Var}(X) = \sigma^2$.

A normal distribution can be used to describe a quantity that results from the additive combination of many independent effects.

Def B.1.8: Log-Normal Distribution

A random variable X is *log-normally distributed* if $\ln X$ is normally distributed, that is

$$\ln X \sim \mathcal{N}(\mu, \sigma^2)$$

Its probability density function is

$$f(x \mid \mu, \sigma) = \frac{1}{x\sigma\sqrt{2\pi}} \exp\left(-\frac{(\ln x - \mu)^2}{2\sigma^2}\right) \qquad x > 0$$

A log-normal distribution can be used to describe a strictly positive quantity whose order of magnitude, rather than its value, is normally distributed, such as a fragment's area-to-mass ratio or the magnitude of the velocity kick it receives in a fragmentation event.

B.2. Generalized Dirac-Delta Function

The Dirac delta is the tool that transfers a density from one phase-space to another. It is not a function but a *distribution*, defined by how it acts under an integral rather than by its values.

Def B.2.1: Dirac Delta

The *Dirac delta* $\delta(x)$ is the distribution that vanishes for $x \neq 0$, carries unit mass

$$\int_{-\infty}^{+\infty} \delta(x)\,\mathrm{d}x = 1$$

and reproduces the value of a test function at a point, a property known as *sifting*,

$$\int_{-\infty}^{+\infty} f(x)\delta(x-a)\,\mathrm{d}x = f(a) \tag{B.1}$$

It can be viewed as the limit of a sequence of normalized pulses of vanishing width, for instance a normal distribution as $\sigma \to 0$, and it inherits the following properties

$$\delta(-x) = \delta(x) \qquad \delta(ax) = \frac{\delta(x)}{|a|} \qquad x\,\delta(x) = 0 \tag{B.2}$$

The delta over $\mathbb{R}^n$ is the product of n one-dimensional deltas, one per component

$$\delta(\boldsymbol{x}) = \delta(x_1)\cdots\delta(x_n)$$

The scaling property makes the units explicit: a delta of an m-component argument carries the reciprocal units of that argument raised to the m-th power, which is what converts a phase-space density into a density over a different space rather than merely relabelling it.

The argument of the delta is often not a coordinate but a function of one. The delta is then evaluated on the zero set of that function, and the change of variables introduces a Jacobian.

Def B.2.2: Dirac Delta of a Function

Let $\boldsymbol{g} : \mathbb{R}^n \to \mathbb{R}^m$ with $m \leq n$ be continuously differentiable with simple zeros. Split the coordinates as $\boldsymbol{x} = (\boldsymbol{y}, \boldsymbol{z})$, with $\boldsymbol{y} \in \mathbb{R}^m$ the variables solved for and $\boldsymbol{z} \in \mathbb{R}^{n-m}$ left free, and for each $\boldsymbol{z}$ let $\boldsymbol{y}_\sigma(\boldsymbol{z})$ be the roots of $\boldsymbol{g}(\boldsymbol{y}, \boldsymbol{z}) = \boldsymbol{0}$. Then

$$\int_{\mathbb{R}^n} f(\boldsymbol{x})\,\delta[\boldsymbol{g}(\boldsymbol{x})]\,\mathrm{d}^n\boldsymbol{x} = \int_{\mathbb{R}^{n-m}} \sum_\sigma \frac{f(\boldsymbol{y}_\sigma(\boldsymbol{z}), \boldsymbol{z})}{|\det \boldsymbol{J}(\boldsymbol{y}_\sigma(\boldsymbol{z}), \boldsymbol{z})|}\,\mathrm{d}^{n-m}\boldsymbol{z} \tag{B.3}$$

where the sum runs over every root and the integral is taken over those $\boldsymbol{z}$ for which at least one root exists.

The delta therefore collapses exactly m of the n dimensions, leaving an integral over the $(n-m)$-dimensional solution manifold, and weights each root by the reciprocal of the determinant of the $m \times m$ Jacobian in the solved variables. Two special cases are worth naming. For $m = n$ the manifold is a set of isolated points and the result is the sum

$$\sum_\sigma \frac{f(\boldsymbol{x}_\sigma)}{|\det \boldsymbol{J}(\boldsymbol{x}_\sigma)|}$$

with no integral left. For $m = n = 1$ this reduces to

$$\delta[g(x)] = \sum_\sigma \frac{\delta(x - x_\sigma)}{|g'(x_\sigma)|}$$

If no root exists the integral is zero, which is how a geometric impossibility is expressed algebraically.

In section 4.2 the constraint is $\boldsymbol{g}(\boldsymbol{u}_1) = \boldsymbol{r}(\boldsymbol{u}_1) - \boldsymbol{r}(\boldsymbol{u}_2)$, a map $\mathbb{R}^4 \to \mathbb{R}^3$, so $m = 3$ of the four phase-space dimensions are collapsed and one is left free. The three variables solved for are (h_1, i_1, Ω_1), the free one is the argument of latitude u_1, and the weight is the determinant of eq. (4.21). The vanishing of the integral for non-overlapping altitude shells and the two inclination branches $\sigma = \pm$ of eq. (4.18) are both instances of the root structure above.

B.3. Poisson Point Processes

A stochastic process is a model used to describe systems that evolve in a random manner. Formally, it is a collection of random variables indexed by time or space. A Poisson point process is a specific type of stochastic process used to count the occurrence of random events over time, given the instantaneous rate at which events occur $\lambda(t)$ (or the average number of events per unit time). In a Poisson process, the random variables being collected are the number of events $N(t)$ at time t. To qualify as such, a Poisson point process must satisfy the following conditions:

1. The number of events $N(t_0, t_1)$ occurring in an interval $[t_0, t_1]$ follows a Poisson distribution whose mean is the time integral of the *intensity function* $\lambda(t)$

$$N(t_0, t_1) \sim \text{Poisson}(\Lambda(t_0, t_1)) \qquad \text{with} \quad \Lambda(t_0, t_1) = \int_{t_0}^{t_1} \lambda(t)\,\mathrm{d}t = \mathbb{E}\left[N(t_0, t_1)\right]$$

   Indeed, $\Lambda(t_0, t_1)$ is the expected number of events in that interval. It follows that the probability of the number of events occurred between t_0 and t_1 to be exactly k is

$$P(N(t_0, t_1) = k) = \frac{\Lambda(t_0, t_1)^k e^{-\Lambda(t_0, t_1)}}{k!}$$

2. The number of events that occur in completely disjoint (non-overlapping) time intervals are independent of each other. For example, the number of events between time t_0 and t_1 has no bearing on the number of events that occur in a later interval t_2 and t_3.

The process of randomly sampling the number of events from a Poisson distribution is often indicated with the following notation

$$N(t_0, t_1) = \int_{t_0}^{t_1} \mathrm{d}\mathcal{N}(\Lambda)$$

This corresponds to counting the number of vertical jumps occurred in the Poisson process between times t_0 and t_1.

Def B.3.1: Homogeneous Poisson Process

A Poisson process is said to be *homogeneous* if the intensity function is constant $\lambda(t) = \lambda$. Then, the expectation becomes $\mathbb{E}\left[N(t_0, t_1)\right] = \Lambda(t_0, t_1) = \lambda(t_1 - t_0)$.

First-Arrival Time

Consider the case in which $t_0 = 0$ and $N(0) = 0$. Then, the number of events occurred in a Poisson process up to time t follows a Poisson distribution

$$N(t) \sim \text{Poisson}(\Lambda(t)) \qquad \text{with} \quad \Lambda(t) = \int_0^t \lambda(\tau)\,\mathrm{d}\tau$$

and

$$P(N(t) = k) = \frac{\Lambda(t)^k e^{-\Lambda(t)}}{k!}$$

The probability that the first event fires strictly after time t is the same as the probability of having exactly zero events up to time t

$$P(T_1 > t) = P(N(t) = 0) = e^{-\Lambda(t)}$$

Therefore, the probability of the first event firing strictly before time t is

$$P(T_1 \leq t) = 1 - e^{-\Lambda(t)}$$

Finally, the probability density function of the first-arrival time is obtained by differentiating the above expression with respect to time

$$f_{T_1}(t) = \frac{\mathrm{d}}{\mathrm{d}t}\left(1 - e^{-\Lambda(t)}\right) = \lambda(t) e^{-\Lambda(t)}$$

C

Numerical Methods

C.1. Gauss-Legendre Quadrature

The overlap hyper-volume of eq. (4.22) has no closed form and is evaluated numerically. The scheme used is Gauss-Legendre quadrature, which approximates a definite integral by a weighted sum of the integrand at a set of points chosen so that both the points and the weights are optimal for a polynomial integrand.

Def C.1.1: Gauss-Legendre Quadrature

The Q-*point Gauss-Legendre rule* approximates an integral over the reference interval $[-1, 1]$ as

$$\int_{-1}^{1} f(\xi)\,\mathrm{d}\xi \simeq \sum_{q=1}^{Q} w_q f(\xi_q)$$

where the *abscissae* ξ_q are the Q roots of the Legendre polynomial P_Q and the *weights* are

$$w_q = \frac{2}{\left(1-\xi_q^2\right)\left[P'_Q(\xi_q)\right]^2}$$

The rule is exact for every polynomial of degree at most $2Q-1$. All abscissae lie strictly inside $(-1, 1)$: the Gauss-Legendre rule is *open*, and never samples an interval endpoint.

An integral over an arbitrary interval is mapped affinely onto the reference one,

$$\int_a^b g(x)\,\mathrm{d}x = \frac{b-a}{2}\sum_{q=1}^{Q} w_q\, g\left(\frac{b-a}{2}\xi_q + \frac{a+b}{2}\right) \tag{C.1}$$

and an integral over several variables is handled by a tensor-product rule, applying eq. (C.1) in each dimension independently, at a cost of $Q_1 Q_2 \cdots Q_d$ evaluations of the integrand for a d-dimensional integral. This is the form used in section 4.2, with Q_i, Q_Ω and Q_u points over the outer variables (i_2, Ω_2, u_2), each of which ranges over one lattice bin or one full revolution.

C.2. Fourier Methods for Circulant Sums

Two separate steps of the model rest on the same structural fact, established in eq. (4.25): at fixed inclination bins, the angular overlap of two nodes depends on their RAAN bins only through the

signed difference between them, not on the two bins individually. The first step is offline. Because the overlap is circulant in the RAAN indices, the full table of ordered bin pairs is generated by the third-order tensor $\Phi_{m_1 m_2 \Delta n}$ of eq. (4.26), and the quadrature is run N_Ω times fewer. The second step is at run time. The system-wide collision intensity Λ_k of eq. (5.2) is a sum of the rate over every ordered pair of populated nodes, and a circulant kernel can be contracted and evaluated efficiently via the discrete Fourier transform as is derived in this section.

Throughout, the imaginary unit is written $\mathrm{j} = \sqrt{-1}$, since i is reserved for inclination, and κ denotes a frequency index, since k is reserved for the time step. All sequences are indexed by RAAN bin and are of length N_Ω, with indices understood modulo N_Ω, so that they are functions on the cyclic group $\mathbb{Z}_{N_\Omega}$ rather than on a finite interval. This is not a modelling convenience but a property of the lattice: RAAN is periodic, the bins tile the full 360°, and bin $N_\Omega - 1$ is genuinely adjacent to bin 0.

C.2.1. The Discrete Fourier Transform

Def C.2.1: Discrete Fourier Transform

The *discrete Fourier transform* of a sequence $(a_0, \dots, a_{N_\Omega - 1})$ is the sequence $\hat{a}$ defined by

$$\hat{a}_\kappa = \sum_{n=0}^{N_\Omega - 1} a_n e^{-2\pi \mathrm{j} \kappa n / N_\Omega}$$

with $\kappa = 0, \dots, N_\Omega - 1$ and the transform is inverted by

$$a_n = \frac{1}{N_\Omega} \sum_{\kappa=0}^{N_\Omega - 1} \hat{a}_\kappa e^{+2\pi \mathrm{j} \kappa n / N_\Omega}$$

Evaluated directly the transform costs $O(N_\Omega^2)$ operations, which would defeat the purpose; the fast Fourier transform recursively factors the sum over n by the divisors of N_Ω and brings the cost to $O(N_\Omega \log N_\Omega)$. Every complexity statement in this section assumes the fast algorithm.

Def C.2.2: Circular Cross-Correlation

The *circular cross-correlation* of two sequences a and b of length N_Ω is the sequence

$$(a \star b)_{\Delta n} = \sum_{n=0}^{N_\Omega - 1} a_n b_{(n + \Delta n) \bmod N_\Omega}$$

The cross-correlation at lag Δn collects every ordered pair of entries separated by exactly Δn positions on the ring. This is precisely the bookkeeping a circulant kernel requires: if a quantity attached to a pair of RAAN bins depends on the pair only through its separation, then the entire double sum over bins collapses to a single sum over separations, with $a \star b$ supplying the population weight at each one.

The *correlation theorem* turns a cross-correlation into a pointwise product of transforms, and the *Plancherel identity* equates a sum of products of two sequences with the corresponding sum over their transforms. Together they place the contraction of a circulant kernel against a cross-correlation entirely in the frequency domain: for real sequences Φ, a and b of length N_Ω,

$$\sum_{\Delta n} \Phi_{\Delta n} \, (a \star b)_{\Delta n} = \frac{1}{N_\Omega} \sum_{\kappa=0}^{N_\Omega - 1} \hat{\Phi}_\kappa \, \hat{a}_\kappa \, \hat{b}^*_\kappa \tag{C.2}$$

where the asterisk denotes complex conjugation. No inverse transform appears: the contraction is read off from three forward transforms and a pointwise product, the cross-correlation itself never being wanted for its own sake. The left-hand side is manifestly real, so only the real part of the right-hand side need be formed.

The sum halves again. Since all sequences involved are real, their transforms are *Hermitian*, $\hat{a}_{N_\Omega-\kappa} = \hat{a}_\kappa^*$, so the summand of eq. (C.2) at κ and at $N_\Omega - \kappa$ are complex conjugates of one another and contribute twice their common real part. The terms that are their own conjugate partners, $\kappa = 0$ and, for even N_Ω, the Nyquist term $\kappa = N_\Omega/2$, are real and counted once. Grouping the spectrum into conjugate pairs therefore reduces eq. (C.2) to a sum over half the frequencies,

$$\sum_{\Delta n} \Phi_{\Delta n}(a \star b)_{\Delta n} = \frac{1}{N_\Omega} \operatorname{Re} \sum_{\kappa=0}^{\lfloor N_\Omega/2 \rfloor} w_\kappa \, \hat{\Phi}_\kappa \, \hat{a}_\kappa \, \hat{b}_\kappa^* \qquad w_\kappa = \begin{cases} 1 & \kappa = 0 \\ 1 & \kappa = N_\Omega/2, \ N_\Omega \text{ even} \\ 2 & \text{otherwise} \end{cases} \tag{C.3}$$

The reduction is exact, not an approximation: only half the spectrum is ever computed, at half the storage and half the arithmetic, and the discarded half carries no information that the *Parseval weights* w_κ do not restore.

C.2.2. Application to the Aggregate Collision Rate

It remains to show that the pair rate has the required form, namely a circulant kernel in the RAAN separation multiplied by a bilinear form in quantities carried by each node separately. Collect first the two ingredients that are circulant. The overlap tensor $\Phi^{\text{sym}}_{m_1 m_2 \Delta n}$ of eq. (4.26) is circulant by construction. The relative velocity of eq. (4.30) is circulant as well, since at bin centres

$$\cos\theta^*_{12} = \sin i^*_{m_1} \sin i^*_{m_2} \cos(\Delta n \, \Delta\Omega) + \cos i^*_{m_1} \cos i^*_{m_2}$$

depends on the two RAAN bins only through their separation, the bins being uniform of width $\Delta\Omega$, and the cosine being even in Δn. The two are therefore folded into a single *velocity-weighted overlap tensor*

$$\Phi^{v}_{m_1 m_2 \Delta n} = 2 \sin\left(\frac{\theta^*_{m_1 m_2 \Delta n}}{2}\right) \Phi^{\text{sym}}_{m_1 m_2 \Delta n} \tag{C.4}$$

which is real, circulant, and inherits from Φ^{sym} and the evenness of the kinematic factor the swap symmetry $\Phi^{v}_{m_1 m_2 \Delta n} = \Phi^{v}_{m_2 m_1 [(-\Delta n) \bmod N_\Omega]}$. It also vanishes on the diagonal, $\Phi^{v}_{mm0} = 0$, which is the algebraic statement of the suppression of within-node collisions noted in section 4.2: two objects in the same bin are co-orbital at this resolution and close on each other at zero relative velocity. Self-pairs therefore drop out of every sum below without needing to be excluded by hand.

Next, the population dependence. Writing the first and second *diameter moments* of a node as

$$d^{(1)}_{S_i} = \sum_q^{S_i} d^q_{S_i} \qquad\qquad d^{(2)}_{S_i} = \sum_q^{S_i} \left(d^q_{S_i}\right)^2 \tag{C.5}$$

the cross-node impact cross-section of eq. (4.31) reads

$$A_c = \frac{\pi}{4 x_{S_i} x_{S_j}} \underbrace{\left[x_{S_j} d^{(2)}_{S_i} + x_{S_i} d^{(2)}_{S_j} + 2 d^{(1)}_{S_i} d^{(1)}_{S_j} \right]}_{T_{S_i S_j}}$$

and the phase-space densities are $n^{(1)}_{\boldsymbol{u}} = x_{S_i}/|\mathcal{U}_{\ell m_1}|$ and $n^{(2)}_{\boldsymbol{u}} = x_{S_j}/|\mathcal{U}_{\ell m_2}|$, where $|\mathcal{U}_{\ell m}|$ is the measure of the bin, independent of the RAAN index for uniform RAAN bins. The crucial cancellation is that the object counts in the denominator of A_c are exactly the counts in the numerators of the two densities, so

$$A_c \, n^{(1)}_{\boldsymbol{u}} n^{(2)}_{\boldsymbol{u}} = \frac{\pi}{4 \, |\mathcal{U}_{\ell m_1}| \, |\mathcal{U}_{\ell m_2}|} T_{S_i S_j}$$

carries no explicit $x_{S_i} x_{S_j}$ denominator. Had it survived, the pair rate would not have been bilinear and none of what follows would apply. Absorbing the remaining altitude-only factors of the pair rate into

$$\Gamma_\ell = \frac{\pi v^*_\ell}{4} \left(\frac{1}{R_\oplus + h_\ell} - \frac{1}{R_\oplus + h_{\ell+1}} \right) \tag{C.6}$$

the rate between nodes $S_i = (\ell, m_1, n_1, S)$ and $S_j = (\ell, m_2, n_2, S')$ becomes

$$\dot{\eta}_{S_i S_j} = \frac{\Gamma_\ell}{|\mathcal{U}_{\ell m_1}|\,|\mathcal{U}_{\ell m_2}|}\,\Phi^{v}_{m_1 m_2 [(n_2 - n_1) \bmod N_\Omega]}\,T_{S_i S_j} \tag{C.7}$$

which is a circulant kernel in Δn times a form that is bilinear in the triples $(x, d^{(1)}, d^{(2)})$ of the two nodes, one factor drawn from each. This is the structure the previous subsection requires.

Fix now an altitude shell ℓ, a pair of inclination bins (m_1, m_2) and a pair of species (S, S'), and view the node quantities along the remaining RAAN axis as sequences of length N_Ω. Aggregating eq. (C.7) over every ordered pair of RAAN bins, and re-indexing the double sum by the separation Δn with $n_2 = (n_1 + \Delta n) \bmod N_\Omega$,

$$\dot{\eta}_{\ell m_1 S m_2 S'} = \sum_{n_1, n_2} \dot{\eta}_{S_i S_j} = \frac{\Gamma_\ell}{|\mathcal{U}_{\ell m_1}|\,|\mathcal{U}_{\ell m_2}|} \sum_{\Delta n} \Phi^{v}_{m_1 m_2 \Delta n}\,G_{\Delta n} \tag{C.8}$$

where the population weight at separation Δn is the sum of three circular cross-correlations, one per term of the trinomial,

$$G_{\Delta n} = \left(d^{(2)} \star x\right)_{\Delta n} + \left(x \star d^{(2)}\right)_{\Delta n} + 2\left(d^{(1)} \star d^{(1)}\right)_{\Delta n} \tag{C.9}$$

with the left operand of each product taken on the group (ℓ, m_1, S) and the right operand on (ℓ, m_2, S'). Applying eq. (C.3) term by term, the entire double sum over RAAN bins is evaluated as a single pass over the half-spectrum,

$$\dot{\eta}_{\ell m_1 S m_2 S'} = \frac{\Gamma_\ell}{N_\Omega\,|\mathcal{U}_{\ell m_1}|\,|\mathcal{U}_{\ell m_2}|}\,\mathrm{Re}\sum_{\kappa=0}^{\lfloor N_\Omega/2 \rfloor} w_\kappa\,\hat{\Phi}^{v}_{m_1 m_2 \kappa}\Big[\hat{d}^{(2)}_{m_1 S \kappa}\,\hat{x}^{*}_{m_2 S' \kappa} + \hat{x}_{m_1 S \kappa}\,\hat{d}^{(2)*}_{m_2 S' \kappa} + 2\,\hat{d}^{(1)}_{m_1 S \kappa}\,\hat{d}^{(1)*}_{m_2 S' \kappa}\Big] \tag{C.10}$$

The transform of the kernel is population-independent and is taken once at start-up, with the weights w_κ and the factor $1/N_\Omega$ folded into the stored array, so that at run time the kernel enters as a precomputed complex tensor of shape $N_i \times N_i \times (N_\Omega/2 + 1)$. What remains per step is three forward transforms of the population sequences, batched over all shells, inclination bins and species at once, followed by a pointwise product summed over the frequency axis. The system-wide intensity Λ_k of eq. (5.2) then follows by summing eq. (C.10) over shells, inclination pairs and species pairs and multiplying by $\alpha \Delta t_k$.

Two bookkeeping corrections attach to this. First, eq. (C.10) is an *ordered* sum: letting m_1 and m_2 run independently over all inclination bins counts each physical pair of distinct species once in each orientation, which is what the unordered species-pair accumulator wants, but for $S = S'$ it counts every unordered pair twice, and a factor of one half is applied after the inclination sum. Second, self-pairs need no exclusion, since $\Phi^{v}_{mm0} = 0$ as established above.