

**High-Fidelity Simulation of Small-Body Lander/Rover
Spacecraft (Draft)**

by

Stefaan Van wal

B.Sc. Aerospace Engineering, Delft University of Technology

Delft, The Netherlands, 2012

M.Sc. Aerospace Engineering, Delft University of Technology

Delft, The Netherlands, 2015

A thesis submitted to the

Faculty of the Graduate School of the

University of Colorado in partial fulfillment

of the requirements for the degree of

Doctor of Philosophy

Smead Department of Aerospace Engineering Sciences

2018

This thesis entitled:
High-Fidelity Simulation of Small-Body Lander/Rover Spacecraft (Draft)
written by Stefaan Van wal
has been approved for the Smead Department of Aerospace Engineering Sciences

Prof. Daniel Scheeres

Prof. Hanspeter Schaub

Prof. Jay McMahon

Prof. Elizabeth Bradley

Dr. Simon Tardivel

Date _____

The final copy of this thesis has been examined by the signatories, and we find that both the content and the form meet acceptable presentation standards of scholarly work in the above mentioned discipline.

Van wal, Stefaan (Ph.D., Aerospace Engineering Sciences)

High-Fidelity Simulation of Small-Body Lander/Rover Spacecraft (Draft)

Thesis directed by Prof. Daniel Scheeres

The scientific return of spacecraft missions that explore solar system small bodies can be increased through the inclusion of surface exploration with deployed probes. In this dissertation, a methodology is presented that allows for fast, parallel simulation of bouncing trajectories of arbitrary-shaped ballistic probes in the small-body environment. This enables planning of probe deployment and operation, and supports their inclusion on future missions.

The coarse small-body shape is modeled using an implicit signed distance field (SDF) that allows for fast collision detection. Statistical features are included onto the SDF using procedural generation techniques. The small-body gravity field is captured using a voxelization of the classical constant-density polyhedron. Surface interactions between a probe and the surface are accounted for using a hard contact model that takes into account restitution and friction. These models are implemented in a GPU environment to allow for the parallel execution of multiple trajectories.

The developed simulation framework is applied to perform parametric investigations of probe deployment, which quantify the effects of relevant properties of a probe and its target small body. The probe shape and internal mass distribution are found to strongly affect its deployment dynamics, with near-spherical probes dispersing over greater regions than more distorted shapes. The effect of the surface interactions coefficients on the different shapes variants is quantified. The presence of statistical surface features is also shown to further influence probe dynamics.

Finally, the framework is applied to perform a pre-arrival deployment analysis of the MINERVA-II rovers onboard the Hayabusa-2 spacecraft. This analysis identified challenges in the rover deployment and was used to redesign aspects of the nominal rover release sequence. These models will be used to inform the target site selection and follow-on analysis for the Hayabusa-2 mission rover deployments.

Dedication

*To my parents, Dirk and Ilse,
my sister, Astrid,
and my grandparents,
for their love, support, and sacrifice.*

Acknowledgements

My sincerest gratitude goes to my advisor, Prof. Dan Scheeres, for his guidance and support throughout my graduate studies at CU. Not only did he go above and beyond in accommodating me as a foreign student in the CSML lab, but he also gave me the freedom to develop my research along my preferred direction, all the while helping me prepare for my career. My deepest thanks also go to Dr. Simon Tardivel, Dr. Robert Reid, and Prof. Yuichi Tsuda for their mentoring. Everybody: thanks, merci, cheers, arigato!

This was easily the largest project I have worked on so far, and I could not have completed it without the support and love from my family and friends. Katie, thanks for putting up with me for five whole years and for helping me make a home halfway across the world. Richard, thanks for being the best travel buddy I could ever imagine, may we continue to check out beautiful destinations together. Yusuf, thanks for all the great days at Blue Sky Basin; don't kill the momentum! Kristi, thanks for keeping me from breaking down on *that one hike* and for our (weekly?) dinner parties. Emili, thanks for getting me into bouldering and for various impulsive trips.

I doubt I would have survived the perils of grad school without the outdoor access provided by the Hiking Club at CU. The love for the outdoors that I discovered through the Club has changed my life forever; thanks to everyone for all the great trips and leadership opportunities. Special thanks goes out to Meredith and Claire for the great collaboration as first CUHC Trifecta.

Last, but not least, thanks to my fellow CSML members for your company over the years. I am running out of space, but want to give special shoutouts to Nicola, Samm, Jeroen, Alex, and Ben, for all your friendship. I hope to see you somewhere down the road.

Contents

Chapter

1	Introduction	1
1.1	Small bodies	1
1.2	Small-body surface exploration	6
1.3	Literature review	9
1.4	Thesis outline	11
2	Modeling	13
2.1	Equations of motion	14
2.2	Shape	16
2.2.1	Spherical harmonics	17
2.2.2	Polyhedron	18
2.2.3	Signed distance field	23
2.3	Statistical features	30
2.3.1	Polyhedral rocks	34
2.3.2	Fractal noise	36
2.3.3	Procedurally seeded rocks	48
2.4	Gravity	57
2.4.1	Constant-density polyhedron	57
2.4.2	Linearization	59

2.4.3	Model resolution	59
2.4.4	Voxelization	63
2.5	Contact	68
2.5.1	Geometry	71
2.5.2	Collisions	74
2.5.3	Contact motion	86
2.6	Regolith	91
2.7	Numerical methods	95
2.7.1	Trajectory propagation	95
2.7.2	Sequential simulation on the CPU	97
2.7.3	Parallel simulation on the GPU	98
3	The lift-off velocity	102
3.1	Deriving the lift-off velocity	103
3.1.1	Geometry	104
3.1.2	General expression	105
3.1.3	Lift-off from a curved surface	108
3.1.4	Lift-off from a ridge	112
3.1.5	Lift-off from a plane	116
3.1.6	Lift-off from a concave region	120
3.2	Surface curvature	123
3.2.1	General properties	123
3.2.2	Ellipsoid	126
3.2.3	Spherical harmonics	127
3.2.4	Polyhedron	130
3.3	Applications to simple shapes	135
3.3.1	Applications to the sphere	135

3.3.2	Applications to ellipsoids	139
3.4	Applications to Solar System bodies	147
3.4.1	1999 KW ₄ Alpha	147
3.4.2	Bennu	154
3.4.3	Eros	158
3.4.4	Phobos	162
4	Parametric studies	164
4.1	Spherical probes	165
4.1.1	Nominal scenarios	167
4.1.2	Effect of surface interaction coefficients	171
4.1.3	Effect of surface rocks	178
4.1.4	Summary: Deployment sensitivities in favorable environments	185
4.1.5	Deployment sensitivities in challenging environments	186
4.2	Complex landers	203
4.2.1	Nominal deployment scenario	204
4.2.2	Effect of probe shape	208
4.2.3	Effect of surface interaction coefficients	227
4.2.4	Effect of probe mass distribution	230
4.3	Surface topography	237
4.3.1	Effect of surface noise	238
4.3.2	Effect of procedural rocks	247
4.3.3	Effect of surface regolith	251
5	Case study: Hayabusa2 rovers	261
5.1	The Hayabusa2 mission	263
5.1.1	Asteroid Ryugu	263
5.1.2	Mothership and rovers	264

5.1.3	Rover release	266
5.2	Deployment to spherical Ryugu	269
5.2.1	Effect of surface interaction coefficients	272
5.2.2	Effect of release geometry	276
5.2.3	Pre-release maneuver	279
5.2.4	Discussion: Deployment to the reference sphere	282
5.3	Deployment to the Ryugu training model	283
5.3.1	Epoch 1 deployment	285
5.3.2	Epoch 2 deployment	288
5.3.3	Discussion: Deployment to the training model	290
5.4	Effect of the asteroid mass	293
5.5	Summary: The deployment of MINERVA-II	297
6	Conclusions	300
6.1	Simulation of small-body probes	300
6.2	Future work	304
	Bibliography	307

Tables

Table

3.1	Characteristic properties of asteroid 1999 KW ₄ Alpha [85].	147
3.2	Characteristic properties of asteroid (101955) Bennu [105].	155
3.3	Characteristic properties of asteroid (433) Eros [69, 70].	158
4.1	Characteristic properties of asteroid 25143 Itokawa [35, 36].	167
4.2	Nominal spacecraft and environmental parameters.	167
4.3	Deployment parameters of Itokawa scenarios.	168
4.4	Characteristic properties of asteroid 1999 KW ₄ Alpha [85, 87].	187
4.5	Deployment parameters of 1999 KW ₄ Alpha scenarios.	196
4.6	Nominal release parameters of complex lander deployment to Itokawa.	205
4.7	Relevant properties of tested symmetric lander shapes.	210
4.8	Critical impact angles θ_E for the tested spacecraft shapes and regolith sizes.	252
5.1	Physical properties of the MINERVA-II rovers and their release.	266
5.2	Release sequence parameters and 3σ errors.	269
5.3	Nominal parameters used in Ryugu reference sphere simulations.	270
5.4	Release azimuth values of the MINERVA-II rovers at both release epochs.	278

Figures

Figure

1.1	Halley's comet, as (<i>left</i>) shown on the 11th-century Bayeux Tapestry and (<i>right</i>) imaged by the Giotto spacecraft in 1986 [52].	2
1.2	Highlights of the Rosetta mission (image courtesy ESA, see rosetta.esa.int).	3
1.3	(<i>left</i>) The impact of Deep Impact on comet Tempel 1 [1], (<i>middle</i>) the surface of Eros as seen during the descent of NEAR-Shoemaker [136], and (<i>right</i>) Itokawa as seen by Hayabusa [33].	5
2.1	The applied inertial (\mathcal{N}), target (\mathcal{T}), and spacecraft (\mathcal{S}) frames.	14
2.2	(<i>left</i>) Spherical harmonic and (<i>right</i>) polyhedron shape model of Phobos	18
2.3	(<i>left</i>) High- and (<i>right</i>) low-resolution polyhedron shape model of asteroid Itokawa.	20
2.4	Illustration of the grids that comprise (<i>left</i>) the Z-atlas, (<i>middle</i>) the X-atlas, and (<i>right</i>) the features included in a sample local world.	21
2.5	Illustration of local world switching during collision convergence.	22
2.6	Illustration of local world sizes on the polyhedron model of asteroid Itokawa.	22
2.7	Illustration of bounding sphere usage in collision detection with a polyhedron model.	23
2.8	The signed distance function of a closed two-dimensional object.	24
2.9	Sections of the signed distance function $d_s(\mathbf{x})$ of a sphere.	26
2.10	Illustration of the tri-linear SDF interpolation.	27
2.11	Sections of the SDF of comet 67P/C-G.	29

2.12	Surface shapes captured by a tri-linear interpolator. Taken from [61].	31
2.13	Stochastic rock models bias lander settling towards low-slope regions.	33
2.14	Stochastic rock models may fail to capture grazing impacts.	33
2.15	Illustration of the rock creation algorithm.	35
2.16	(left) Itokawa's rocky surface [74] and (right) active local world with procedurally generated rocks.	35
2.17	Interpolation of a one-dimensional noise field.	38
2.18	Interpolation of a one-dimensional noise field for varying A , with fixed $f = 1 \text{ m}^{-1}$	39
2.19	Interpolation of a one-dimensional noise field for varying f , with fixed $A = 1 \text{ m}$	39
2.20	Interpolation of a one-dimensional noise field for varying (A, f) pairs.	40
2.21	Interpolation of a one-dimensional noise field for increasing number of octaves.	41
2.22	Interpolation of a two-dimensional noise field for increasing number of octaves.	41
2.23	Interpolation of a three-dimensional noise field for increasing number of octaves.	42
2.24	Facet resolutions on Itokawa.	43
2.25	(left) Low- and (right) high-resolution probe models resolve surface impacts differently.	45
2.26	Progressive tessellation of an octahedron.	46
2.27	Vertex resolution as a function of the tessellation level for the platonic solids.	46
2.28	Modification of the impact normal.	47
2.29	Height field of a spherical rock field.	49
2.30	Periodic tiling of the spherical rock field.	49
2.31	(left) Texture index and (right) corresponding aperiodic tiling of spherical rocks.	51
2.32	Three-dimensional view of the aperiodic tiling of spherical rocks.	51
2.33	Realization of generalized distance fields of various shapes with $p = 100$	52
2.34	Distance field of an icosahedron ($a = 4$ and $b = 13$), for different values of the norm p	52
2.35	Three-dimensional view of a distance field used to generate procedurally seeded rocks.	53
2.36	Two-dimensional view of a distance field used to generate procedurally seeded rocks.	54

2.37 PSRs generated using (left) an unmodified and (right) a strictly negative distance field.	54
2.38 Illustration of the distortion of rock shapes by a strictly negative distance field.	55
2.39 Procedurally seeded rocks with varying minimum rock size d_{min}	56
2.40 Detailed view of procedurally seeded rocks with $d_{min} = 0.10$ cm.	56
2.41 Vectors used in the evaluation of the constant-density polyhedron model.	58
2.42 Itokawa shape models.	60
2.43 67P/C-G shape models.	60
2.44 Phobos shape models.	60
2.45 Gravity error of various polyhedron model resolutions.	61
2.46 Maximum relative gravity error for various polyhedron model resolutions.	63
2.47 Relative gravity error of voxelized acceleration and perturbation models of Phobos.	65
2.48 Itokawa relative surface gravity error for varying voxel grid size.	67
2.49 67P/C-G relative surface gravity error for varying voxel grid size.	67
2.50 Phobos relative surface gravity error for varying voxel grid size.	67
2.51 Alignment of the collision vector in central and eccentric collisions.	68
2.52 Snapshots of MASCOT impacting a granular surface, as simulated by Thuillet et al. [122]	70
2.53 Illustration of the contact geometry.	72
2.54 Illustration of velocity hodographs for different slip/stick behaviors.	81
2.55 Illustration of the different phases in the rover deployment to a small body.	86
2.56 Surface morphology changes at Philae's first impact site [11].	92
2.57 Definition of the impact angle for (left) a pointmass and (right) a complex probe.	93
2.58 Modulation of the coefficient of restitution e as a function of impact angle θ	94
2.59 Flowchart of event convergence algorithm.	97
2.60 Flowchart of serial simulation algorithm.	98
2.61 Flowchart of parallel simulation algorithm.	101

3.1	(left) Full and (right) departure plane geometry applied in the derivation of the lift-off velocity.	104
3.2	Geometry viewed along the $\hat{\mathbf{b}}_\theta$ axis.	108
3.3	(left) Variation of the two lift-off velocities at an arbitrary position on a rotating ellipsoid, and (right) illustration of the azimuthal angle η	111
3.4	Geometry of spherical object rolling on the edge between two facets.	113
3.5	Comparison of full and reduced ridge lift-off velocity expressions on Itokawa.	115
3.6	Geometry for the lift-off from a rotating plane.	116
3.7	Illustration of a particle lifting off from a rotating plane in (left) 2D and (right) 3D, where the green arrow indicates low lift-off velocity, the orange arrows indicate intermediate lift-off velocity, and the red arrows indicate infinite lift-off velocity.	118
3.8	Ratio of planar-to-curved lift-off velocities at an arbitrary position on a rotating ellipsoid.	119
3.9	Simulated lift-off of particles on a rotating plane.	120
3.10	Radius of curvature and lift-off velocity variation for a (left) convex and (right) convex-concave region.	122
3.11	Illustration of (<i>left</i>) general properties of curvature and (<i>right</i>) Meusnier's theorem, on a warped cylinder.	123
3.12	Geometry applied in the computation of polyhedron curvatures of vertex P_i	130
3.13	Scaled differences in Gaussian curvature between the spherical harmonic and various polyhedron models of Phobos.	133
3.14	Lift-off plane geometry on the surface of a sphere for variations in (left) departure point latitude ϕ with fixed $\eta = 0^\circ$ and (right) departure direction η with fixed $\phi = 30^\circ$	136
3.15	Simulated lift-off of a particle from the surface of a rotating sphere.	137
3.16	Lift-off velocity curves for variations in ϕ and η on the surface of a rotating sphere.	138

3.17 (left) Minimum lift-off velocity, (right) maximum lift-off velocity, and (bottom) extremum difference distribution across the surface of a rotating sphere. The arrows indicate the directions in which the respective extrema occur.	139
3.18 Simulated lift-off for particles on the surface of a general ellipsoid.	141
3.19 The tested oblate, prolate, and general ellipsoids.	142
3.20 (left) Minimum and (right) maximum ρ across the surface of the oblate ellipsoid. . .	142
3.21 (left) Minimum and (right) maximum ρ across the surface of the prolate ellipsoid. .	143
3.22 (left) Minimum and (right) maximum ρ across the surface of the general ellipsoid. .	143
3.23 (left) Minimum and (right) maximum V_θ across the surface of the oblate ellipsoid. .	144
3.24 (left) Minimum and (right) maximum V_θ across the surface of the prolate ellipsoid. .	144
3.25 (left) Minimum and (right) maximum V_θ across the surface of the general ellipsoid. .	144
3.26 (top) V_θ and (bottom) ρ vs. η on the surface of the oblate ellipsoid.	145
3.27 (top) V_θ and (bottom) ρ vs. η on the surface of the prolate ellipsoid.	146
3.28 (top) V_θ and (bottom) ρ vs. η on the surface of the general ellipsoid.	146
3.29 (left) Minimum and (right) maximum V_θ across the polyhedron model of 1999 KW ₄ Alpha.	148
3.30 Ridge lift-off velocity across the polyhedron model of 1999 KW ₄ Alpha.	148
3.31 RMS error of spherical harmonics fit of 1999 KW ₄ Alpha.	149
3.32 (left) Polyhedron and (right) $J = 32$ spherical harmonic model of 1999 KW ₄ Alpha.	150
3.33 XZ-plane intersection of the polyhedron and spherical harmonic model of 1999 KW ₄ Alpha.	151
3.34 (left) Minimum and (right) maximum lift-off velocity across the spherical harmonic model of 1999 KW ₄ Alpha.	151
3.35 Lines of minimum latitude of release on 1999 KW ₄ Alpha.	154
3.36 (left) Minimum and (right) maximum V_θ across the polyhedron model of Bennu. . .	156
3.37 Ridge lift-off velocity across the polyhedron model of Bennu.	156
3.38 RMS error of spherical harmonics fit of Bennu.	156

3.39	(<i>left</i>) Minimum and (<i>right</i>) lift-off velocity across Bennu’s spherical harmonic model.	157
3.40	Lines of minimum latitude of release on Bennu.	157
3.41	(<i>left</i>) Minimum and (<i>right</i>) maximum lift-off velocity across Eros’s polyhedron model.	159
3.42	Ridge lift-off velocity across Eros’s polyhedron model.	159
3.43	(<i>left</i>) Polyhedron, (<i>middle</i>) $J = 18$, and (<i>right</i>) $J = 40$ spherical harmonic model of Eros.	160
3.44	Sectors of the Eros polyhedron shape that are separately fitted.	160
3.45	RMS error of five-sector spherical harmonics fit of Eros.	160
3.46	(<i>left</i>) Polyhedron and (<i>right</i>) $J = 24$ <i>sectored</i> spherical harmonic shape model of Eros.	161
3.47	(<i>left</i>) Minimum and (<i>right</i>) maximum lift-off velocity across the $J = 24$, <i>sectored</i> spherical harmonic model of Eros.	161
3.48	Minimum lift-off velocity on Phobos’s surface.	163
3.49	Maximum lift-off velocity on Phobos’s surface.	163
3.50	Normalized ridge lift-off velocity on Phobos’s surface.	163
4.1	Illustration of mass configurations of a spherical probe.	166
4.2	Overview of the nominal Itokawa deployments.	168
4.3	Heat map and position scatter plot of nominal Itokawa deployments.	169
4.4	Statistics of nominal Itokawa deployments.	169
4.5	Settling time statistics for varying surface interaction coefficients.	172
4.6	Settling slope statistics for varying surface interaction coefficients.	172
4.7	Settling time statistics for varying surface interaction coefficients.	174
4.8	Settling slope statistics for varying surface interaction coefficients.	174
4.9	Heat map of settling locations for varying e .	174
4.10	Heat map of settling locations for varying C_{RR} .	175
4.11	Heat map of settling locations for varying f .	175
4.12	Shaded Muses Sea and Sagami-hara regions on Itokawa.	175

4.13 (top) Settling time and (bottom) settling angle on nominal and heterogeneous Itokawa surfaces.	176
4.14 Settling position heat maps of nominal and heterogeneous Itokawa surfaces.	176
4.15 Small-body surface for varying number of rocks.	179
4.16 Settling time and slope statistics for varying K	179
4.17 Heat map of settling locations for varying K in high-altitude Itokawa scenario.	180
4.18 Heat map of settling locations for varying K in low-altitude Itokawa scenario.	180
4.19 Settling time statistics for global and rock-only variation of e	183
4.20 Surface slopes on 1999 KW ₄ Alpha.	188
4.21 Velocity hodograph for impacts on a (<i>left</i>) flat and (<i>right</i>) rocky surface.	189
4.22 Geometry of first impact of a particle onto a rotating, spherical target.	190
4.23 Outgoing surface-relative particle velocity for varying surface inclinations.	193
4.24 Outgoing surface-relative particle velocity for varying periaipse and surface inclination.	193
4.25 Outgoing surface-relative particle velocity for varying coefficient of restitution.	194
4.26 Sample deployment orbits to 1999 KW ₄ Alpha with varying periaipse radius.	198
4.27 Surface dispersion of deployments to 1999 KW ₄ Alpha with varying periaipse radius.	198
4.28 Energy and impact angle for varying r_P deployments to 1999 KW ₄ Alpha.	199
4.29 Surface dispersion of deployments to 1999 KW ₄ Alpha with varying rock densities.	201
4.30 Energy and impact angle for varying rock density deployments to 1999 KW ₄ Alpha.	201
4.31 Settling time versus impact angle for varying rock density deployments to 1999 KW ₄ Alpha.	201
4.32 Typical deployment of a cube in the low-altitude Itokawa scenario.	205
4.33 Settling time statistics of the nominal Itokawa deployment scenarios.	206
4.34 Dispersion plots of the nominal cube deployment to Itokawa.	206
4.35 Mollweide projection of the dispersion plots of the nominal Itokawa deployment.	207
4.36 The eight considered probe shapes.	208
4.37 The interior angles β_i used to determine the vertex defect δ	210

4.38	Settling position density of the considered shapes.	211
4.39	Mollweide projection of the settling position density of the considered shapes.	211
4.40	Mollweide projection of the dispersion areas of the considered shapes.	212
4.41	Dispersion statistics of the considered shapes.	212
4.42	Settling time PDFs of the considered shapes.	213
4.43	The minimum and maximum aspect angle α of various regular, convex polygons.	214
4.44	Geometry of a regular, convex polygon.	215
4.45	Distribution of collision vector composition for random 2D polygon impacts.	216
4.46	Selected geometries as marked in Fig. 4.45.	216
4.47	Sample impacts of a tetrahedron with opposing friction force direction.	217
4.48	Distribution of aspect angles for the eight probe shapes, for random impact geometries.	218
4.49	Example aspect angles of the tetrahedron.	219
4.50	Illustration of the two modes of motion of the MINERVA-II-1 rover.	220
4.51	Sample MINERVA-II-1 simulation in which Mode 2 dominates.	221
4.52	Sample MINERVA-II-1 simulation in which Mode 1 dominates.	221
4.53	Bouncing of a sphere on (left) flat and (right) irregular terrain.	222
4.54	Settling time PDFs of considered shapes, with and without continued contact.	223
4.55	Changes in settling position and time when ignoring continued contact.	224
4.56	Atypical deployment of a cube, in which the inclusion of contact motion results in significant additional motion.	225
4.57	Illustration of proposed (left) plexiglass and (right) air-bearing table experiments.	227
4.58	Mean deployment t_f for varying (left) restitution and (right) friction coefficient.	228
4.59	Illustration of different mass configurations of a triangle.	230
4.60	Settling time distribution with varied spacecraft inertias.	232
4.61	Mean settling time for varying f and j of the tetrahedron, cube, and dodecahedron.	232
4.62	Tangential contact point velocity statistics for cube variants.	233
4.63	Tangential contact point rate for sample cube collision.	235

4.64	Definition of the α and β angles.	238
4.65	Variation of α and β in impacts on (left) a large and (right) a small surface feature.	239
4.66	Variation in the α and β angles for impact with different noise levels.	241
4.67	Impact angle distributions for flat and noisy surfaces, (left) for untessellated and (right) finely tessellated spacecraft shapes.	243
4.68	Settling time PDFs of deployments to various Itokawa surfaces.	244
4.69	Mean settling time trends of deployments to various Itokawa surfaces.	246
4.70	Settling time PDFs of deployments to smooth and noisy Itokawa.	246
4.71	Sample rock populations with varying k_0	248
4.72	Mean settling time of various shapes in the presence of PSRs.	248
4.73	Surface dispersion of various shapes in the presence of PSRs.	249
4.74	Mean settling time time for various impact angles and rock densities.	251
4.75	Critical impact angle variation for the three tested shapes.	253
4.76	Restitution modulation in the two tested regolith surfaces.	253
4.77	Definition of the incidence angle γ	254
4.78	Initial impact angles for different shapes and incidence angles.	255
4.79	Successive impact angles for different shapes and γ , for deployment with fixed e	256
4.80	Settling time statistics for relevant deployments to hard and regolith surfaces.	257
4.81	Successive impact angles for different shapes and γ , deployment with fixed e	259
4.82	Settling time statistics for relevant deployments to hard and regolith surfaces.	259
5.1	The formally best-fit shape model of Ryugu, taken from Müller et al. [76].	263
5.2	The Ryugu training model.	264
5.3	(left) Rear and (right) side view illustrations of the Hayabusa2 spacecraft at home position.	265
5.4	Applied shape models of the MINERVA-II rovers.	266
5.5	Variation of the sub-solar and sub-Earth latitudes on the Ryugu training model.	267

5.6	Hayabusa2 quasi-orbit about the Ryugu training model, at both release epochs. . . .	268
5.7	(left) High-resolution and (right) reduced-resolution training model used for gravity evaluations.	269
5.8	Sample simulation batch to the Ryugu reference sphere, and corresponding settling time statistics.	270
5.9	Surface dispersion of the sample simulation batch.	271
5.10	Deployment statistics for e variation on Ryugu reference sphere.	272
5.11	Altitude profile of the Hayabusa spacecraft during touchdown [145].	274
5.12	Relevant estimated coefficient of restitution values in literature [64, 145, 13, 12]. . .	274
5.13	Deployment statistics for f variation on Ryugu reference sphere.	275
5.14	(left) Mean t_f and (right) $3\sigma A_d$ for (e, f) variation on the Ryugu reference sphere. .	276
5.15	Settling time statistics for (left) ϕ and (right) α variation on the Ryugu reference sphere.	277
5.16	Surface dispersion of rover releases due east and due west.	277
5.17	Variation of the tangential rover impact velocity v_T for $V_0 = 20$ cm/s.	278
5.18	Magnitude and direction of the horizontal pre-release maneuver for $V_0 = 20$ cm/s. .	280
5.19	Surface dispersion of east/west release scenarios, with and without PRM.	280
5.20	Settling time of east/west release scenarios, with and without PRM.	280
5.21	Surface dispersion for various release α and V_0 when including the nominal PRM. . .	282
5.22	Sample MINERVA-II release sequence without (dashed line) and with (solid line) PRM.	284
5.23	Selected release longitudes λ for Epoch 1 simulations, with nominal descent trajectory.	285
5.24	Release, first impact, and final settling positions of Epoch 1 deployment simulations to the Ryugu training model.	286
5.25	Detailed view of the $\lambda = 30$ deg Epoch 1 deployment to plains.	289
5.26	Detailed view of the $\lambda = 150$ deg Epoch 1 deployment to a crater.	289
5.27	Surface dispersion statistics of Epoch 1 deployment to the Ryugu training model. . .	289

5.28	Settling time statistics of Epoch 1 deployment to the Ryugu training model.	290
5.29	Release, first impact, and final settling positions of Epoch 2 deployment simulations to the Ryugu training model.	291
5.30	Detailed view of the $\lambda = 240$ deg Epoch 2 deployment to plains.	292
5.31	Detailed view of the $\lambda = 90$ deg Epoch 2 deployment to a crater.	292
5.32	Settling time statistics of Epoch 2 deployment to the Ryugu training model.	292
5.33	Surface dispersion statistics of Epoch 2 deployment to the Ryugu training model. . .	293
5.34	Comparison of Ryugu training model ρ with relevant small bodies [34, 20, 146, 88]. .	294
5.35	Worst- and best-case Epoch 1 training model deployments, with varying μ	295
5.36	Worst- and best-case Epoch 2 training model deployments, with varying μ	296
5.37	Surface dispersion statistics of deployment with varying μ	297
5.38	Settling time statistics of deployment with varying μ	297

Chapter 1

Introduction

1.1 Small bodies

Comets The scientific interest in comets as targets for large space missions is not surprising, given their rich observational history among which the earliest records date back more than two millennia. Since comets are visible with the naked eye, prehistoric man must surely have wondered about these bright-tailed objects crossing the sky. The first *recorded* observation of a comet is found on Chinese oracle bones from the fourth century B.C. and suggests that early civilizations primarily interpreted comets as astrological phenomena [152]. Ever since, comet apparitions are not uncommon in written history and are frequently associated with important events. Halley's comet (1P/Halley) is easily the most (in)famous of the comets, as it is the only known short-period comet that is visible to the naked eye at perihelion, roughly every 75 years. Though orbital perturbation prevent reliable back-propagation beyond roughly a thousand years, Halley's apparitions can be traced back several millennia in the written record [147]. The most dramatic association made with Halley's apparitions may be that of 1066 A.D. as represented on the famous *Bayeux Tapestry*. This 230-foot long tapestry (see Fig. 1.1) depicts the events of the Norman conquest of England and shows Halley's comet in the sky during the coronation of King Harold II of England. It is included as an omen, heralding Harold's defeat at the infamous Battle of Hastings [83].

It was not until several hundred years later that comets started to be recognized as *astrominical*, rather than *astrological* objects, thereby paving the way for exploration with spacecraft. The International Cometary Explorer (ICE/ISEE-3) was the first to intercept a comet when it



Figure 1.1: Halley’s comet, as (*left*) shown on the 11th-century Bayeux Tapestry and (*right*) imaged by the Giotto spacecraft in 1986 [52].

passed through the plasma tail of Giacobini-Zinner (21P/Giacobinni-Zinner) in 1985 [137]. It verified the “dirty iceball” hypothesis that assumed comet nuclei to be a mix of dust, rocks, water ice, and other frozen gases. Various spacecraft studied Halley’s comet during its 1986 apparition, with ESA’s Giotto spacecraft being the first to image a comet nucleus, see Fig. 1.1 [38]. The contrast between Halley appearing as an omen on the Bayeux tapestry and as an object of scientific study in Giotto’s observation is staggering, and is an exhilarating example of man’s innate drive for exploring the unknown. Two other noteworthy comet explorers from the early 2000’s are NASA’s Deep Impact and Stardust missions. Deep Impact sent an impactor onto the nucleus of comet Tempel 1 (9P/Tempel) in 2005, and imaged the impact with a flyby spacecraft, see Fig. 1.3 [2]. Stardust returned samples of the coma of comet Wild 2 (81P/Wild) in 2006 [15].

Easily the most comprehensive case of comet exploration was ESA’s Rosetta mission, which explored comet 67P/Churyumov-Gerasimenko between 2014 and 2016 and obtained a wealth of new observations [18]. Perhaps the most notable measurement made by Rosetta was that of the isotope ratio of outgassed water: the deuterium-to-hydrogen ratio of its water was found to be three times that of terrestrial water, disproving the hypothesis that the water on Earth was delivered through bombardment by comets. The spacecraft also deployed the Philae lander, which performed

a swath of surface measurements revealing, among other things, a hard ice layer covered with a regolith layer approximately 25 cm thick [11]. Some highlights from the Rosetta mission are shown in Fig. 1.2 which shows, from top left to bottom right: comet 67P/C-G in August 2014, one of Rosetta’s final descent images in September 2016, Philae on the comet surface in September 2016, and Philae’s landing leg in November 2014.

In an interesting turn of events, NASA is currently investigating a return to 67P/C-G with the CAESAR spacecraft, which proposes to return samples from the comet surface to Earth as part of the New Frontiers program [108]. In the spring of 2019, NASA will decide whether or not to pursue the mission.




Figure 1.2: Highlights of the Rosetta mission (image courtesy ESA, see rosetta.esa.int).

Asteroids The observational history of asteroids is far shorter than that of comets, with Ceres being discovered as the first asteroid by Giuseppe Piazzi on the 1st of January, 1801. In the first half of the nineteenth century, astronomers realized that an entire ‘belt’ of asteroids exists

between the orbits of Mars and Jupiter, and the number of observations has skyrocketed since [28]. Several hundred thousand asteroids, spread throughout the entire solar system, have been categorized. The Galileo probe was the first to closely image an asteroid, providing photos of asteroid 951 Gaspra in 1991 while en-route to Jupiter [9]. Galileo also visited 243 Ida and revealed the existence of the small satellite (243 Ida I) Dactyl, thereby discovering the first ever *binary* asteroid system. Several missions to asteroids have been flown since and are planned for the future.

The NEAR-Shoemaker spacecraft explored asteroid 433 Eros as the first mission in NASA’s Discovery program [135]. It investigated the composition, mineralogy, mass distribution, and magnetic field of Eros, eventually landing on its surface at the end of the mission in 2001 (see Fig. 1.3). The Japanese Hayabusa probe performed a similar study of asteroid 25143 Itokawa and returned samples from its surface to Earth in 2010. It revealed that Itokawa is a “rubble-pile” asteroid: an aggregate of many large boulders, smaller rocks, and tiny regolith particles, see Fig. 1.3 [34]. Asteroid exploration is steadily continuing beyond these first missions: the Japanese Hayabusa-2 and American OSIRIS-REx spacecraft successfully launched and are en-route to the carbonaceous asteroids 162173 Ryugu and 101955 Bennu, respectively [125, 56]. Both missions will extensively study their targets and return surface samples to Earth in the early 2020s. Hayabusa2 also plans to deploy three small “*Micro/Nano Experimental Robot Vehicle for Asteroid - II*” (MINERVA-II) rovers and one “*Mobile Asteroid Surface Scout*” (MASCOT) lander to Ryugu’s surface [125, 43].

Additionally, both the Lucy and Psyche missions were recently selected as part of NASA’s Discovery program. Lucy plans to perform flybys at several of the Trojan asteroids following launch in 2021 [58]. Psyche will explore the metallic asteroid 16 Psyche, thought to be the exposed core of a proto-planet, following a planned launch in 2022 [82]. NASA also plans to launch the Double Asteroid Redirect Test (DART) in late 2020 or early 2021, which will impact the small secondary of the binary asteroid system 65803 Didymos in 2022 at a velocity of approximately 6 km/s [19]. Through observations of the orbital changes of the binary body with Earth-based telescopes, the mission enables validation of kinetic impactors for planetary defense. ESA is also currently investigating its proposed Hera mission, which would orbit Didymos in order to better

examine the resulting crater and asteroid orbit change [66]. At the time of writing, Hera has not been fully funded yet.




Figure 1.3: (*left*) The impact of Deep Impact on comet Tempel 1 [1], (*middle*) the surface of Eros as seen during the descent of NEAR-Shoemaker [136], and (*right*) Itokawa as seen by Hayabusa [33].

Small-body exploration Together, asteroids and comets are known as the *small bodies* of our solar system. They are considered interesting targets for exploration with space missions for three main reasons:

- (1) It is commonly accepted that asteroids and comets are leftover ‘building blocks’ from the early solar system. They have experienced a rich history over the past 4.5 billion years, with (relatively) frequent collision, fission, and re-accretion events [67, 50]. Small bodies therefore provide a window through time into these early conditions and can teach us lessons about the formation and evolution of our solar system. All of the mentioned missions have contributed to this goal in some form.
- (2) The asteroid impact risk is regarded as one of the major long-term threats to human civilization. Consider for example the Chicxulub impact that is thought to have caused the mass extinction of dinosaurs around 66 million years ago, or the Tunguska event. The latter is thought to be the result of the atmospheric disruption of a stony asteroid, amounting to the largest impact event in recorded history and flattening a forest area of approximately 2,000 km² in 1908 [21, 126]. More recently, the Chelyabinsk superbolide injured about

1,500 Russian residents in 2013, mostly through damage from glass that shattered when the bolide exploded [60]. Events like these are a stark reminder of the damage that can be caused by such impacts; small-body exploration allows for the development and validation of planetary defense strategies that mitigate the impact risk. This was demonstrated most notably by the Deep Impact spacecraft and is planned to be explored in further detail by the DART impactor.

- (3) Asteroids and comets contain large amounts of water, metals, and complex carbon molecules; resources which could be harvested and used *in situ*. The collection of water is especially promising, as it can be electrolyzed into rocket propellant using solar energy. A small body placed in Earth orbit could thus become an orbital fuel station and significantly reduce the cost of spaceflight. Exploratory missions to small bodies provide a proving ground for these (exotic) new technologies [59]. Although the proper ‘mining’ of asteroids remains a far-off goal, it is receiving increasing attention from private companies such as Planetary Resources and Deep Space Industries, both of which have received notable support from the government of Luxembourg [16].

1.2 Small-body surface exploration

The missions mentioned above have been successful in establishing our current understanding of the origins, characteristics, and evolution of asteroids and comets. They primarily addressed the three goals listed above through remote sensing observations with orbiting spacecraft. The return of such a mission can be significantly increased by including lander/rover probes capable of sampling the small-body (sub-)surface *in situ*. This was demonstrated most notably by Rosetta’s Philae lander, which investigated the surface of comet 67P/C-G. As mentioned, Philae revealed the existence of a high-strength ice layer covered with soft regolith [13]. The lander was equipped with thrusters and harpoons to anchor itself to the comet’s surface. However, both of these active landing systems did not activate, which caused the lander to rebound off the comet’s surface

following an initial touchdown. Subsequent bounces eventually brought the lander to rest far away from the targeted landing site. It settled at an angle of approximately 30 deg against a cliff or crater wall that significantly restricted insolation. As a result, Philae's secondary batteries could not be recharged and contact with the lander was lost after three days, though approximately 80% of the planned science operations were successfully performed [11]. In a more dramatic turn of events, the MINERVA hopper onboard the Hayabusa mission was improperly released while the mothership ascended from asteroid Itokawa. The craft was subsequently lost to space [150]. These two examples illustrate the challenges inherent in the delivery of landers to the small-body surface. This is one reason why they have often been perceived as high-risk mission elements, despite their potential for increasing the scientific return of a mission.

Fortunately, small scientific packages are now more commonly being considered as secondary payloads for small-body exploration missions. Most notably, the Japanese Hayabusa2 spacecraft that is set to arrive at asteroid Ryugu in June/July of 2018 carries a total of four deployable payloads. These include the three Japanese MINERVA-II rovers, each with a mass of approximately 1 kg and carrying simple scientific instruments. The rovers are also equipped with momentum exchange mechanisms that, if operated successfully, will allow them to demonstrate (intentional) surface mobility on the surface of an asteroid [125]. Hayabusa2 also carries the European MASCOT lander with a mass of approximately 10 kg, which contains a multitude of instruments for surface exploration as well as a swinging arm that enables surface mobility [43]. A similar craft (MASCOT-II) was investigated for ESA's AIM mission, though its cancellation put an end to this study [19]. Nevertheless, MINERVA-II and MASCOT are testament to the increasing interest in small-body surface exploration. Though they were not selected, the BASiX and MarcoPolo-R also proposed to deploy payloads to the surfaces of their targets [6, 41].

In order to further advance lander/rover spacecraft from small secondary payloads to primary missions systems, mission designers must be convinced of their reliability and low inherent risk to a mothership. This confidence can be (partially) provided by high-fidelity simulation software, in which the many complexities of the small-body environment are accounted for, and which can

predict the settling distribution of a lander/rover spacecraft. Since the complexity of the small-body environment renders the ballistic bouncing motion of a spacecraft chaotic, it is impossible to *precisely* predict its settling behavior. However, by performing many simulations while accounting for relevant uncertainties, it is possible to provide an *accurate* prediction of the spacecraft settling statistics. These simulations have two major applications. First, they allow for the direct application to relevant mission scenarios, such as the deployment of the rover/lander spacecraft onboard Hayabusa2. By taking into account all relevant uncertainties in the mothership navigation and rover release, combined with high-fidelity models of Ryugu's shape, gravity, and surface properties, it is possible to predict the rover settling statistics. This allows for the development of a deployment strategy, *i.e.*, a strategy for where and how the rovers can be released in order to target some landing zone, while complying with relevant mission requirements. The latter may include line-of-sight requirements, restrictions to exclude certain surface regions, or ensuring sufficient insolation. High-fidelity simulation software enables informed discussion between spacecraft operators (that set mission constraints) and scientists (that wish to maximize the scientific return of a deployed payload). Second, if such software is capable of performing a large number of simulations in sufficiently small amount of time, it allows for parametric studies to be carried out. In this, a single parameter of the lander spacecraft or small-body environment is varied and its effect on the resulting spacecraft motion investigated. This allows for a quantification of the effects of, *e.g.*, surface interaction coefficients, lander shape, or lander mass distribution. The resulting insights can provide design guidelines for lander/rover probes meant to, *e.g.*, settle on the small-body surface with minimal surface dispersion.

The goal of this thesis is to establish an efficient methodology for such high-fidelity simulation of small-body surface exploration probes. This will be applied to perform broad parametric studies that establish the factors governing lander motion, as well as a pre-arrival deployment analysis of rovers onboard the Hayabusa2 spacecraft. Before providing a detailed outline of the thesis, the literature relevant to this topic is first reviewed.

1.3 Literature review

Several authors have investigated the bouncing motion of payloads in the small-body environment. A brief overview highlighting the contributions of the relevant works follows:

- (1) The first author to consider this problem are Sawai et al. In their 2001 paper “*Development of a target marker for landing on asteroids,*” they simplify the problem by studying the motion of a simple pointmass, bouncing on a rotating ellipsoid [100]. This work assisted in the mechanical design of the target markers deployed by Hayabusa to asteroid Itokawa, but is too limited in its modeling techniques to allow for reliable landing predictions or broader analyses.
- (2) Yoshimitsu et al. detail their design of the MINERVA-I rover onboard Hayabusa in the 1999 and 2004 papers “*Hopping rover “MINERVA” for asteroid exploration*” [151] and “*Development of autonomous rover for asteroid surface exploration*” [148]. This work focuses mostly on the mechanical design and experimental testing of the hopping mechanism. Since the rover was designed to operate without knowledge of its location on the surface, the authors do not attempt to perform simulations of its motion across the small-body surface. For a discussion on MINERVA’s failure to reach Itokawa’s surface, the reader is referred to [150].
- (3) Ogawa et al. extended the work by Sawai et al. on the target marker design to the Hayabusa2 mission in their 2010 paper “*Precise landing of space probes on asteroid using multiple markers*” [81]. In this, they consider the use of multiple target markers, but continue to use a highly simplified planar model for the asteroid surface and pointmass probe model.
- (4) Tardivel et al. were the first to consider the *global* motion of a (spherical) lander probe bouncing on the polyhedron shape of a small body, in their 2014 paper “*Contact motion on surface of asteroid*” [118]. They observed chaotic probe dynamics and linear-angular

velocity couplings that cannot be captured by simplified dynamical models. In addition to this first step towards the accurate modeling of surface interactions, they also gave significant attention to the selection of initial descent orbits that guarantee the surface impact of a lander [116, 114, 115]. Finally, the authors performed a preliminary investigation of the effects of rocks on the small-body surface, noting that they may significantly effect the deployment of spherical probes. The rocks were accounted for using a simplified stochastic impact model. This work culminated in a Ph.D. thesis carried out at the University of Colorado Boulder [120].

- (5) Hockman et al. approach the challenge of operating a hopping payload from a different direction and focus on the mechanical design and stochastic motion planning of a cube-shaped hopper, in their 2017 papers “*Design, Control, and Experimentation of Internally-Actuated Rovers for the Exploration of Low-gravity Planetary Bodies*” [45] and “*Stochastic Motion Planning for Hopping Rovers on Small Solar System Bodies*” [44]. They account for the small-body shape and gravity using the polyhedron model, but do not use high-fidelity models for the surface interactions of an arbitrary-shaped lander.
- (6) The deployment of the Philae lander was only studied down to the point of first impact, since the probe was designed to anchor itself to the surface of comet 67P/Churyumov-Gerasimenko, see for example the 2012 paper “*Philae Descent Trajectory Computation and Landing Site Selection on Comet Churyumov-Gerasimenko*” by Canalias et al [17]. This analysis is therefore not relevant to the problem at hand, although the reconstruction of Philae’s trajectory sheds some light on the dispersive properties of the comet surface [13].
- (7) Finally, some studies have been performed regarding the deployment of the European MAS-COT lander onboard Hayabusa2. In the 2017 paper “*Mascot: Analyses of the descent and bouncing trajectories to support the landing site selection,*” Lorda et al. perform a deployment analysis of MASCOT to the coarse, pre-arrival shape model of asteroid Ryugu [62]. Again, this work is restricted to using pointmasses to represent the more complex lan-

der shape, with a statistical model for bouncing. For a detailed discussion of MASCOT's mechanical design, the reader is referred to [43].

In summary, the problem of bouncing payload motion in the small-body environment has been investigated by various authors. However, all published works are limited to the motion of either pointmass or spherical lander models. This is a severe limitation; non-spherical shapes exhibit fairly chaotic bouncing motion akin to a die being thrown on a table. Furthermore, none of the studies available in literature are implemented in a manner that allows fast simulation of a large number of trajectories. This is necessary in order to perform statistically sound mission analysis or any type of broad parametric study. Finally, only one author considered the inclusion of statistical surface features. Even though this used a fairly limited stochastic model, it was able to indicate the importance of including such features in simulation.

1.4 Thesis outline

The above literature review reveals the lack of a methodology to simulate the bouncing motion of payloads in the small-body environment. It must be able to handle both targets and probes with an arbitrary shape, but also account for features such as rocks on the target surface. It should be integrated in a simulation framework that allows for the fast simulation of many trajectories, to enable broad studies and application to realistic mission scenarios. This is the goal of the thesis, which builds upon the work of [118] by extension to complex lander shapes and surfaces. The resulting framework is used to perform investigations into lander motion and is applied to a deployment case study of rovers onboard Hayabusa2. This thesis consists of 5 chapters.

Chapter 2 presents the modeling methods applied in the thesis. This first consists of a presentation of the general problem set-up and the equations of motion that describe the motion of a spacecraft in the small-body environment. The methods for modeling the shape and gravitational field of a small body are then discussed, as well as the contact interactions between a spacecraft and the surface. Next, the procedural generation of statistical features on the surface is discussed,

as well as the presence of a regolith layer. Finally, the numerical implementation of these models in a parallelized computation environment is presented.

Chapter 3 discusses the surface motion of a vehicle and is specifically concerned with establishing speed limits for tangential motion of the vehicle, called the *lift-off velocity*. A general derivation for this velocity is first presented, as well as provide methods for the computation of the surface curvature of a body, which is a primary determinant of the lift-off velocity. This is then applied to some generic bodies such the sphere, ellipsoid, and plane in order to investigate some general properties of lift-off behavior. Finally, lift-off velocity distributions are provided and discussed for asteroids 1999 KW₄ Alpha, Bennu, Eros, as well as the Mars moon Phobos.

Chapter 4 provides the first application of the simulation framework, namely, parametric studies of lander motion. This begins with an investigation of how a spherical lander is affected by the surface interaction coefficients, its mass distribution, and the presence of rocks on the surface. This is followed by a similar investigation for landers with a more complex shape, that discusses how the shape and mass distribution of a complex lander affects its motion. This is compared and contrasted with the trends observed for spherical probes. Finally, the effects of surface topography on lander motion are investigated. This consists of the inclusion of realistic variations to the ‘hard’ surface layer as well as a model to account for the existence of surface regolith.

Chapter 5 then presents the second application, that is, the investigation of a specific mission scenario. The simulation framework is applied to perform a pre-arrival deployment analysis of the MINERVA-II rovers onboard the Hayabusa2 spacecraft. The mission set-up is discussed first, followed by some simple simulations to the spherical Ryugu model. This reveals some challenges in the nominal deployment sequence; modifications that increase the chances of successful deployment are suggested. Those lessons are then applied when performing simulations to the full Ryugu training model, which reveal the effect of variations in the terrain on corresponding deployment trends. This results in further recommendations for the rover release. Finally, the effects of changes in the asteroid mass on the rover dynamics are briefly investigated. The conclusions of this thesis are provided in Chapter 6, which also includes suggested pathways for future work.

Chapter 2

Modeling

This chapter presents the modeling applied in our simulation framework and is structured as follows. It begins by establishing the relevant reference frames with respect to which the motion of the lander/rover spacecraft is tracked and develop the corresponding equations of motion in Section 2.1. This is followed with a discussion of the three different small-body shape models applied in this work in Section 2.2. First is the spherical harmonic shape model, used primarily in the lift-off velocity chapter. Second, the polyhedron model, which is the classical small-body shape model. Finally, the signed distance field, an *implicit* shape model that enables fast distance evaluations. These shape models represent relatively smooth surface devoid of features such as rocks and boulders. Section 2.3 presents methods to efficiently generate such features. The chapter then goes on to review our gravity field modeling in Section 2.4. After a brief overview of some commonly-used models, the constant-density polyhedron model is first reviewed. This is followed by a discussion of the voxelization of this model, which allows a body's gravity field to be pre-computed and enables cheap run-time interpolations. Next, a methodology for modeling contact interactions between a spacecraft and the small-body surface is presented, in Section 2.5. This begins with an overview of the contact geometry and then goes into detail on the evaluation of (impulsive) collisions. This is followed by the derivation of a corresponding model for contact motion, in which the spacecraft is in continued contact with the surface. A simple model for regolith deformation is then presented in Section 2.6. The chapter concludes with a discussion our parallelized GPU implementation that allow for fast simulation of a large number of trajectories, in Section 2.7.

2.1 Equations of motion

The motion of some rigid *spacecraft* with arbitrary shape and inertia is modeled. This spacecraft moves in the neighborhood of a targeted small body (the “*target*,” assumed to be an asteroid, comet, or small moon), also with arbitrary shape, in uniform rotation. In order to describe the motion of the spacecraft, the three reference frames illustrated in Fig. 2.1 are used. The first is the inertial frame \mathcal{N} with orthonormal axes $\{\hat{\mathbf{n}}_1, \hat{\mathbf{n}}_2, \hat{\mathbf{n}}_3\}$. Second, the rotating target frame \mathcal{T} , with its origin at the target center of mass and orthonormal axes $\{\hat{\mathbf{t}}_1, \hat{\mathbf{t}}_2, \hat{\mathbf{t}}_3\}$ aligned with the target principal axes of inertia. This frame is assumed to have zero translational velocity relative to the inertial \mathcal{N} -frame. Finally, the rotating spacecraft frame \mathcal{S} , with origin at the spacecraft center of mass and orthonormal axes $\{\hat{\mathbf{s}}_1, \hat{\mathbf{s}}_2, \hat{\mathbf{s}}_3\}$ that are aligned with the spacecraft principal axes of inertia.




Figure 2.1: The applied inertial (\mathcal{N}), target (\mathcal{T}), and spacecraft (\mathcal{S}) frames.

The state \mathbf{X} of the spacecraft-target system is represented using six vectors that add to a total of 20 variables:

$$\mathbf{X} = \left[\mathbf{x} \quad \mathbf{v} \quad \mathbf{q}_{S/T} \quad \boldsymbol{\omega}_{S/T} \quad \mathbf{Q}_{T/N} \quad \boldsymbol{\Omega}_{T/N} \right]^T \quad (2.1)$$

The \mathbf{x} and \mathbf{v} vectors are the spacecraft center of mass position and velocity relative to the target, and together capture the motion of that center of mass. The spacecraft attitude is expressed

relative to the target using the quaternion $\mathbf{q}_{S/\mathcal{T}}$ and corresponding angular velocity $\boldsymbol{\omega}_{S/\mathcal{T}}$. Both \mathbf{v}_S and $\boldsymbol{\omega}_{S/\mathcal{T}}$ express velocities relative to a rotating, non-inertial reference frame. As a result, the corresponding equations of motion will contain virtual acceleration components. This type of expression was chosen because the contact forces exerted on the spacecraft drive these relative velocities to zero. If the velocities were expressed relative to, *e.g.*, the inertial \mathcal{N} -frame, their values ‘at rest’ would be non-zero and location-dependent. The chosen approach allows the spacecraft to intuitively come to rest at $\mathbf{v} = \boldsymbol{\omega} = \mathbf{0}$. In addition to these four spacecraft-related vectors, the attitude of the target is also tracked relative to the inertial \mathcal{N} -frame, using its quaternion $\mathbf{Q}_{\mathcal{T}/\mathcal{N}}$ and corresponding angular velocity $\boldsymbol{\Omega}_{\mathcal{T}/\mathcal{N}}$. Note that the capitalized symbols (\mathbf{q} and $\boldsymbol{\omega}$) refer to the target body, while the non-capitalized symbols (\mathbf{Q} and $\boldsymbol{\Omega}$) refer to the spacecraft. The time rate of change of the system state is described by the following equations of motion:

$$\dot{\mathbf{X}} = \frac{d\mathbf{X}}{dt} = \begin{bmatrix} \dot{\mathbf{x}} \\ \dot{\mathbf{v}} \\ \dot{\mathbf{q}}_{S/\mathcal{T}} \\ \dot{\boldsymbol{\omega}}_{S/\mathcal{T}} \\ \dot{\mathbf{Q}}_{\mathcal{T}/\mathcal{N}} \\ \dot{\boldsymbol{\Omega}}_{\mathcal{T}/\mathcal{N}} \end{bmatrix} = \begin{bmatrix} \mathbf{v} \\ \mathbf{g} - [\tilde{\boldsymbol{\Omega}}_{\mathcal{T}/\mathcal{N}}][\tilde{\boldsymbol{\Omega}}_{\mathcal{T}/\mathcal{N}}]\mathbf{x} - 2[\tilde{\boldsymbol{\Omega}}_{\mathcal{T}/\mathcal{N}}]\mathbf{v} + \frac{1}{m}\mathbf{F}_c \\ \frac{1}{2}[\mathbf{B}(\mathbf{q}_{S/\mathcal{T}})] \begin{bmatrix} 0 & \boldsymbol{\omega}_{S/\mathcal{T}} \end{bmatrix}^T \\ \text{see Eq. 2.5} \\ \frac{1}{2}[\mathbf{B}(\mathbf{Q}_{\mathcal{T}/\mathcal{N}})] \begin{bmatrix} 0 & \boldsymbol{\Omega}_{\mathcal{T}/\mathcal{N}} \end{bmatrix}^T \\ -[\mathbf{I}_{\mathcal{T}}]^{-1}[\tilde{\boldsymbol{\omega}}_{\mathcal{T}/\mathcal{N}}][\mathbf{I}_{\mathcal{T}}]\boldsymbol{\omega}_{\mathcal{T}/\mathcal{N}} \end{bmatrix} \quad (2.2)$$

The quaternion equation of motion makes use of the classical $[\mathbf{B}(\mathbf{q})]$ matrix:

$$[\mathbf{B}(\mathbf{q})] = \begin{bmatrix} q_0 & -q_1 & -q_2 & -q_3 \\ q_1 & q_0 & -q_3 & q_2 \\ q_2 & q_3 & q_0 & -q_1 \\ q_3 & -q_2 & q_1 & q_0 \end{bmatrix} \quad (2.3)$$

In these equations, $[\tilde{\cdot}]$ is the cross-product tensor notation. Furthermore, \mathbf{g} is the gravitational attraction of the target on the spacecraft, \mathbf{F}_c is the net contact force on the spacecraft, and m is the spacecraft mass. In the equation for the spacecraft acceleration, notice the $-[\tilde{\boldsymbol{\Omega}}_{\mathcal{T}/\mathcal{N}}][\tilde{\boldsymbol{\Omega}}_{\mathcal{T}/\mathcal{N}}]\mathbf{x}$ and $-2[\tilde{\boldsymbol{\Omega}}_{\mathcal{T}/\mathcal{N}}]\mathbf{v}$ accelerations, which are respectively the (virtual) centrifugal and Coriolis accelerations

that result from the expression of the spacecraft position and velocity vectors relative to the rotating \mathcal{T} -frame. For completeness, Eq. 2.2 includes the equation of motion for $\dot{\mathbf{\Omega}}_{\mathcal{T}/\mathcal{N}}$ that should be applied when the target is not in a uniform rotation. Since uniform target rotation *is* assumed throughout this thesis, it holds that $\dot{\mathbf{\Omega}}_{\mathcal{T}/\mathcal{N}} = \mathbf{0}$.

In order to derive the equation of motion for the spacecraft angular velocity $\boldsymbol{\omega}_{S/\mathcal{T}}$, this relative angular velocity is first defined as:

$$\boldsymbol{\omega}_{S/\mathcal{T}} = \boldsymbol{\omega}_{S/\mathcal{N}} - \boldsymbol{\omega}_{\mathcal{T}/\mathcal{N}} \quad (2.4)$$

Applying the transport theorem [101], the rate of change of this angular velocity is found as:

$$\dot{\boldsymbol{\omega}}_{S/\mathcal{T}} = -[\mathbf{I}_S]^{-1}[\tilde{\boldsymbol{\omega}}_{S/\mathcal{N}}][\mathbf{I}_S]\boldsymbol{\omega}_{S/\mathcal{N}} + [\mathbf{I}_S]^{-1}\mathbf{L}_c(\mathbf{X}) - [\tilde{\boldsymbol{\omega}}_{S/\mathcal{N}}]\boldsymbol{\omega}_{S/\mathcal{T}} \quad (2.5)$$

in which $[\mathbf{I}_S]$ and $[\mathbf{I}_T]$ are the inertia matrices of respectively the spacecraft and the target, and \mathbf{L}_c is the net contact torques on the spacecraft. The angular velocity $\boldsymbol{\omega}_{S/\mathcal{N}}$ can be computed using Eq. 2.4. Since the target is assumed to be in uniform rotation, the fact that $\dot{\mathbf{\Omega}}_{\mathcal{T}/\mathcal{N}} = \mathbf{0}$ was used.

It is noted that the uniform rotation assumption is not true for all small bodies; indeed, various bodies such as 4179 Toutatis and 253 Mathilde are observed to be non-principal axis rotators [46, 90]. In order to study deployment to such *complex rotators*, one should account for this rotation and apply Euler's rigid body rotational equations to account for the $\dot{\mathbf{\Omega}}_{\mathcal{T}/\mathcal{N}}$ term. For our purposes, it is assumed that any nutation in the small-body rotation is small enough to have a negligible effect on the statistics of the considered lander deployment scenarios. This assumption likely holds for slow complex rotators such as Toutatis, whose effective slope changes by less than 0.1 degree over the course of a rotation period, but may be invalid for fast-rotating bodies [103].

2.2 Shape

Small bodies such as asteroids and comets are not sufficiently large to be shaped into spheres by self-gravitation [121], such that a more complex model is required to capture these bodies' irregular shape. Here, the three models used to capture small-body shapes in this work are reviewed.

2.2.1 Spherical harmonics

The spherical harmonic expansion uses the spherical coordinate set (r, ϕ, λ) to represent the position of some surface point \mathbf{r} . In this, r is the radius, ϕ is the latitude, and λ is the longitude of \mathbf{r} . The transformation between Cartesian and spherical coordinates can be performed with:

$$\begin{cases} r = \sqrt{x^2 + y^2 + z^2} \\ \phi = \arcsin\left(\frac{z}{\rho}\right) \\ \lambda = \arctan\left(\frac{y}{x}\right) \end{cases} \Leftrightarrow \mathbf{r}(\phi, \lambda) = r(\phi, \lambda) \cdot \begin{bmatrix} \cos \phi \cos \lambda \\ \cos \phi \sin \lambda \\ \sin \phi \end{bmatrix} = \begin{bmatrix} x \\ y \\ z \end{bmatrix} \quad (2.6)$$

In the spherical harmonics model, the radius r of a surface point is expressed as a Fourier series function of its latitude and longitude, as [142]:

$$r(\phi, \lambda) = \sum_{j=0}^J \sum_{k=0}^j (A_{jk} \cdot \cos(k\lambda) + B_{jk} \cdot \sin(k\lambda)) \cdot P_{jk}(\sin \phi) \quad (2.7)$$

in which A_{jk} and B_{jk} are the degree j and order k spherical harmonic coefficients, and J is the maximum degree of the considered shape model. Furthermore, P_{jk} is the degree j order k associated Legendre function. These functions are most effectively computed using recursion, where the first three terms are given as:

$$P_{0,0}(x) = 1 \quad \text{and} \quad P_{1,0}(x) = x \quad \text{and} \quad P_{1,1} = \sqrt{1-x^2} \quad (2.8)$$

The other terms can be computed with the following recursion rules [128, p. 593]:

$$\begin{cases} P_{n,0} = \frac{(2n-1)xP_{n-1,0} - (n-1)P_{n-2,0}}{n} & \text{if } n \geq 2 \\ P_{n,m} = P_{n-2,m} + (2n-1)\sqrt{1-x^2}P_{n-1,m-1} & \text{if } m \neq 0, m < n \\ P_{n,n} = (2n-1)\sqrt{1-x^2}P_{n-1,m-1} & \text{if } n \neq 0 \end{cases} \quad (2.9)$$

It is noted that this spherical harmonic expansion is a continuous representation of a surface; it can be evaluated at any arbitrary (ϕ, λ) and thus produce a shape of arbitrary resolution (though the degree J of the model is limited by numerics). As an example, the $J = 45$ spherical harmonic model of the Mars moon Phobos, as developed by [143], is shown on the left side of Fig. 2.2.




Figure 2.2: (left) Spherical harmonic and (right) polyhedron shape model of Phobos

The spherical harmonic shape model works well for small bodies that are fairly spherical in shape; in this case, relatively few A_{jk} and B_{jk} coefficients are sufficient to capture the global body shape, and relatively little memory is required. Note that changes in the value of a single A_{jk} or B_{jk} *globally* affect the shape of a body. Therefore, the inclusion of an additional surface feature on top of an existing model generally requires reconfiguration of all coefficients. In other words, when updated measurements need to be included into a spherical harmonic model, this generally requires a re-fitting of all measurement data. This is a disadvantage of the spherical harmonic model for shape model representation.

2.2.2 Polyhedron

The polyhedron model captures the shape of a small body as a collection of vertices with three-dimensional coordinates, connected by triangular facets. Each facet consists of three edges, with each edge being shared by two facets. This model has been used extensively in the field of graphics and visualization, but has also become the most commonly used small-body shape model. One of the model's advantages is that it can be used to evaluate the gravitational field of a small

body; a discussion of these techniques is provided in Section 2.4. For a quick visualization, the right side of Fig. 2.2 illustrates a sample polyhedron shape model of Phobos and clearly shows the model’s triangular facets. The resolution of a polyhedron model depends on the number of vertices and facets of the considered model, which are likely to differ depending on the purpose of the model. Regardless of the model resolution, *Euler’s criterion* states that for a closed polyhedron:

$$n_V + n_F - n_E = 2 \quad (2.10)$$

in which n_V , n_F , and n_E are respectively the number of vertices, facets, and edges of a considered polyhedron [93, pp. 66]. As noted above, each facet of a *closed* polyhedron consists of three edges, with each edge shared by two facets. Therefore, the number of *unique* edges of a closed polyhedron with triangular faces is equal to $n_E = \frac{3}{2}n_F$. Substituting this into Eq. 2.10 results in:

$$\begin{cases} n_F = 2n_V - 4 \\ n_E = 3n_V - 6 \end{cases} \quad (2.11)$$

This is a powerful formula that can be used to verify the consistency of a polyhedron model; this should be done whenever an externally-provided shape model is to be used in simulation. If a given model does not satisfy Eq. 2.11, it may contain holes, duplicate vertices, or non-manifold vertices and/or edges. These defects can be repaired using the open-source **MeshLab** software [22].

The polyhedron model of Phobos previously shown in Fig. 2.2 is of relatively low resolution, so as to allow for individual facets to be visible. Although low-resolution models are useful for *gravity* representation (as will be discussed in Sec. 2.4), it is necessary to use high-resolution models for *shape* representation. This is because low-resolution models may not sufficiently capture terrain details, such as craters or hills, that can create basins of attraction or rejection for lander/rover spacecraft. As an illustration, consider the high- and low-resolution models of asteroid Itokawa shown in Fig. 2.3. Many of Itokawa’s smaller hills and boulders have clearly disappeared in the low-resolution models; use of this model in simulation may bias the motion of deployed payloads.

The use of the polyhedron model for shape representation in lander/hopper simulations requires computation of the minimum distance between the two, in order to detect collisions and con-




Figure 2.3: (left) High- and (right) low-resolution polyhedron shape model of asteroid Itokawa.

tact between the two. This involves iteration over all features of that shape (vertices/edges/facets), which is numerically burdensome when naïvely applied to a high-resolution model. This burden can be reduced significantly through the application of two techniques: the division of the ‘global’ target surface into a number of smaller ‘local worlds’ and the use of bounding spheres. Although the former was developed by Tardivel et al. [118, 120], it will play a central role in the generation of polyhedral rocks on the small-body surface; it is therefore discussed in some detail. Our use of bounding spheres is a novel development that significantly accelerates the collision detection between a deployed probe and the (possibly rocky) small-body surface.

Atlas Prior to performing simulations, the *atlas* of the targeted small body is constructed. This is done by creating a latitude-longitude grid where individual cells contain those target surface features that fall within the cells’ coverage, taking into account some small additional angular margin. In order to establish a consistent grid size across the surface, one atlas is created relative to the rotation axis of the target (‘Z-atlas’) and a second relative to an axis in its equatorial plane (‘X-atlas’). Each of the two atlases spans $\pm 45^\circ$ of respective latitude, thus covering the entire small body surface with cells of similar size.

Fig. 2.4 illustrates the atlas, showing the coverage of the Z-atlas on the left, of the X-atlas in the middle, and the features included in a sample local world on the right. In this sample world, the solid black lines mark the latitude-longitude grid, the dashed line shows the margin applied to




Figure 2.4: Illustration of the grids that comprise (*left*) the Z-atlas, (*middle*) the X-atlas, and (*right*) the features included in a sample local world.

the central cell, the white triangles are the underlying faceted surface, and the red triangles are those facets included in the local world of the central cell. Assuming a fixed probe shape, this atlas needs to be pre-computed only once and can simply be loaded prior to running simulations, thus decreasing the necessary overhead (though this does carry an additional memory cost).

When propagating probe motion, the simulation continuously updates the ‘active’ local world based on the current probe latitude and longitude¹. The distance computation necessary to achieve collision detection then only requires iteration over the handful of features in the active local world, rather than iteration over all features of the entire small-body surface. This significantly reduces the computational cost of performing collision detection with a high-resolution surface. Fig. 2.5 illustrates a simple two-dimensional example of how the active local world is updated throughout a ballistic arc, with the simulation eventually converging on a probe-surface collision at time t_c . The collision convergence is carried out using a simple bisection method.

Note that although our twin-atlas method ensures consistent latitude-longitude grid sizes, the number of facets per local world is not necessarily uniform. This is especially true for strongly elongated or strongly non-convex targets. As an illustration of this fact, consider the three sample worlds of asteroid Itokawa shown on the left side of Fig. 2.6. Although the angular size of all worlds is the same, their surface coverage is significantly different, in particular for ‘world C’ on Itokawa’s

¹ If the probe position corresponds to local worlds in both atlases, either can be used.

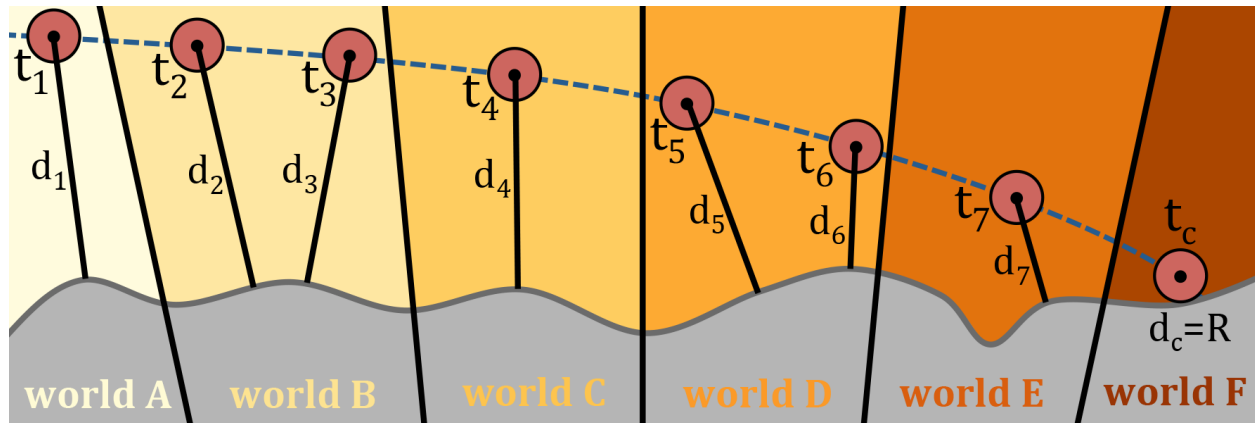


Figure 2.5: Illustration of local world switching during collision convergence.

neck region. This is further illustrated on the right side of Fig. 2.6, where the number of facets on the local worlds of Itokawa’s Z-atlas is plotted. Please note that, in this example, the size of the local worlds has been exaggerated for visual purposes; local worlds typically contain about a dozen facets. The figure shows that local worlds containing parts of the neck region indeed contain more facets than those corresponding to other regions on the body. Although this phenomenon does not hamper the collision detection process on asteroid Itokawa, it may reduce the computational gain of the atlas on strongly non-convex targets, such as comet 67P/C-G. On such targets, a different shape modeling approach may be considered, as will be discussed shortly.




Figure 2.6: Illustration of local world sizes on the polyhedron model of asteroid Itokawa.

Bounding spheres The efficiency of collision detection with respect to a polyhedral shape model can be increased even further through the use of *bounding spheres*; one such sphere is

defined for each local world, and fully encompasses that world with some radius R_i . Following this definition, collisions between a probe and the active world are possible only when the probe has crossed into the respective bounding sphere. Therefore, it is sufficient to simply compute the ‘coarse’ distance between the probe and the bounding sphere in most cases; the ‘fine’ distance computation between the probe and all features of the local world is only carried out when the bounding sphere has been crossed. This process is illustrated for a two-dimensional case in Fig. 2.7, where the coarse bounding sphere computation is sufficient at t_2 , but where the fine feature computation is necessary at t_7 . As this technique significantly reduces the number of times the computationally expensive fine distance computation has to be performed, it reduces the overall computational cost of collision detection. Note that the bounding spheres shown in Fig. 2.7 have their aforementioned margins shown.




Figure 2.7: Illustration of bounding sphere usage in collision detection with a polyhedron model.

2.2.3 Signed distance field

Despite the extensive use of the polyhedron model in small-body research, it suffers from some limitations that largely stem from its discrete and explicit nature. As previously mentioned, collision detection and contact evaluation between a spacecraft and some surface requires distance computations between the two. When performed naively with a polyhedron, the numerical burden of these computations may be significant due to the resolution of the involved models. Although local subdivisions of the surface (such as the longitude-latitude atlas discussed above) may reduce

some of this burden, collision detection using polyhedra remains a significant cost in the simulation of small-body probes. Furthermore, even high-resolution models are limited in terms of the smallest surface features that can be captured and cannot account for smaller rocks and boulders on the surface, even though these may significantly affect the motion of a lander. Although it is possible to populate the surface with polyhedral rocks (as discussed further on), this inevitably further increases the computational cost of collision detection.

As a third and final shape model, consider now the *signed distance field* (SDF), which has seen frequent use for general visualization purposes but (to the best of our knowledge) has not been used for small-body shape modeling. Prior to discussing the SDF, some general properties of signed distance *functions* are reviewed, of which the SDF is one.

Signed distance functions Given some closed and bounded set $\Omega \subset \mathbb{R}^3$, one can distinguish between the interior and exterior portions of the domain \mathbb{R}^3 , denoted respectively Ω^- and Ω^+ . The border between the interior and exterior portions is known as the boundary or *interface*, $\partial\Omega$ [84]. These regions are illustrated in Fig. 2.8.




Figure 2.8: The signed distance function of a closed two-dimensional object.

When using the polyhedron model, this interface is represented *explicitly* as a collection of vertices, edges, and facets. In contrast, a signed distance function represents the interface $\partial\Omega$

implicitly through some $d(\mathbf{x})$ with the following properties:

$$d(\mathbf{x}) = \text{sgn}(\mathbf{x}) \cdot \min_{\mathbf{c} \in \partial\Omega} \|\mathbf{x} - \mathbf{c}\| \quad \text{with} \quad \text{sgn}(\mathbf{x}) = \begin{cases} + & \text{if } \mathbf{x} \in \Omega^+ \\ 0 & \text{if } \mathbf{x} \in \partial\Omega \\ - & \text{if } \mathbf{x} \in \Omega^- \end{cases} \quad \text{and} \quad \|\nabla d(\mathbf{x})\| = 1 \quad (2.12)$$

in which $\|\mathbf{x} - \mathbf{c}\| = \sqrt{(\mathbf{x} - \mathbf{c})^T(\mathbf{x} - \mathbf{c})}$ is the Euclidean or L^2 norm. In words, this function provides the *signed minimum distance* between a point \mathbf{x} and the interface $\partial\Omega$. The distance is signed because the function indicates whether \mathbf{x} lies in the interior Ω^- or exterior Ω^+ portion of the domain. The interface point corresponding to this minimum distance is denoted \mathbf{c}^* , which can be found from \mathbf{x} and the signed distance function as [84]:

$$\mathbf{c}^* = \mathbf{x} - d(\mathbf{x}) \cdot \nabla d(\mathbf{x}) \quad (2.13)$$

In other words, the gradient $\nabla d(\mathbf{x})$ of the signed distance function corresponds to the surface normal $\hat{\mathbf{N}}$ through both \mathbf{c}^* and \mathbf{x} . As a simple example of a signed distance function, consider the sphere with radius R and center at $(0, 0, 0)$. The signed distance d_s from a point $\mathbf{x} = (x, y, z)$ to the surface of the sphere is easily found as:

$$d_s(\mathbf{x}) = \sqrt{x^2 + y^2 + z^2} - R = \sqrt{\mathbf{x}^T \mathbf{x}} - R \quad (2.14)$$

By differentiating Eq. 2.14, its gradient ∇d_s is found as:

$$\nabla d_s(\mathbf{x}) = \frac{\mathbf{x}}{\sqrt{\mathbf{x}^T \mathbf{x}}} \quad (2.15)$$

This gradient indeed has norm $\|\nabla d_s(\mathbf{x})\| = 1$. The behavior of this signed distance function is illustrated in Fig. 2.9, which shows three arbitrary sections of $d_s(\mathbf{x})$ around a sphere with radius $R = 5$ m. As expected, the function has value $d_s = 0$ at the surface of the sphere and increases in value when moving away from the surface. Though not shown, the inverse behavior takes place on the interior of the sphere, where the signed distance function takes on negative values.




Figure 2.9: Sections of the signed distance function $d_s(\mathbf{x})$ of a sphere.

The signed distance field When using a simple shape such as the sphere shown above, the corresponding signed distance function is relatively simple. The spherical harmonics shape model discussed above is an example of a more complex signed distance function. Although this can be used to capture the shape of more irregular targets, it holds that arbitrary-shaped objects such as comet 67P/Churyumov-Gerasimenko (shown in Fig. 2.11) do not have a simple analytic surface expression that can be molded into a distance function. In this case, one can make use of a *signed distance field* (SDF). An SDF is a three-dimensional grid around an object with its mesh points storing the minimum signed distance between that object and the respective points [51]. An interpolation method can then be used to sample the distance field at intermediate points. Here, a simple tri-linear interpolator is used. To illustrate the application of this interpolator, generic $4 \times 3 \times 4$ *voxel* (volumetric pixel) field is shown on the left side of Fig. 2.10.

A few points $\mathbf{x}_{i,j,k}$ of the field are highlighted. In principle, the voxels can be cuboids with different dimensions along the X, Y, and Z axes. Here, it is assumed that the voxels are cubes with side Δx . It is further assumed that the signed distance field $d(\mathbf{x})$ has been evaluated at all




Figure 2.10: Illustration of the tri-linear SDF interpolation.

$\mathbf{x}_{i,j,k}$ and simply wish to linearly interpolate the field at some point $\mathbf{x} = (x, y, z)$. To perform this interpolation, the eight grid values that surround \mathbf{x} with a single cube are selected, as illustrated on the right side of Fig. 2.10. Within this cube, the grid point with minimum coordinates along the X, Y, and Z axes (on the ‘bottom left’ of the cube) defines the (x_0, y_0, z_0) coordinates. To evaluate $d(\mathbf{x})$, the non-dimensionalized coordinates $\hat{\mathbf{x}}$ within the selected voxel are first computed:

$$\hat{\mathbf{x}} = \begin{bmatrix} \hat{x} \\ \hat{y} \\ \hat{z} \end{bmatrix} = \begin{bmatrix} \frac{x-x_0}{x_1-x_0} \\ \frac{y-y_0}{y_1-y_0} \\ \frac{z-z_0}{z_1-z_0} \end{bmatrix} \quad (2.16)$$

in which the subscripts 0 and 1 refer to respectively the minimum and maximum values encountered among the eight corners of the cuboid. This ensures that $0 \leq \{\hat{x}, \hat{y}, \hat{z}\} \leq 1$. Using the tri-linear interpolator, $d(\mathbf{x})$ can then be found as [91, 14]:

$$d(\mathbf{x}) = \kappa_0 + \kappa_1\hat{x} + \kappa_2\hat{y} + \kappa_3\hat{z} + \kappa_4\hat{x}\hat{y} + \kappa_5\hat{y}\hat{z} + \kappa_6\hat{x}\hat{z} + \kappa_7\hat{x}\hat{y}\hat{z} \quad (2.17)$$

in which the κ_i values are equal to:

$$\left\{ \begin{array}{l} \kappa_0 = d(0, 0, 0) \\ \kappa_1 = d(1, 0, 0) - d(0, 0, 0) \\ \kappa_2 = d(0, 1, 0) - d(0, 0, 0) \\ \kappa_3 = d(0, 0, 1) - d(0, 0, 0) \\ \kappa_4 = d(0, 0, 0) - d(1, 0, 0) - d(0, 1, 0) + d(1, 1, 0) \\ \kappa_5 = d(0, 0, 0) - d(0, 1, 0) - d(0, 0, 1) + d(0, 1, 1) \\ \kappa_6 = d(0, 0, 0) - d(1, 0, 0) - d(0, 0, 1) + d(1, 0, 1) \\ \kappa_7 = -d(0, 0, 0) + d(1, 0, 0) + d(0, 1, 0) + d(0, 0, 1) \\ \quad -d(1, 1, 0) - d(0, 1, 1) - d(1, 0, 1) + d(1, 1, 1) \end{array} \right. \quad (2.18)$$

Additionally, the partial derivatives of Eq. 2.17 yield the gradient of the scalar field $d(\mathbf{x})$, as:

$$\nabla d(\mathbf{x}) = \frac{1}{\Delta x} \begin{bmatrix} \kappa_1 + \kappa_4 \hat{y} + \kappa_6 \hat{z} + \kappa_7 \hat{y} \hat{z} \\ \kappa_2 + \kappa_5 \hat{z} + \kappa_4 \hat{x} + \kappa_7 \hat{x} \hat{z} \\ \kappa_3 + \kappa_6 \hat{x} + \kappa_5 \hat{y} + \kappa_7 \hat{x} \hat{y} \end{bmatrix} \quad (2.19)$$

The $1/\Delta x$ factor is included in order to obtain the gradient of the field in *dimensional* coordinates, rather than the normalized \hat{x} coordinates. Finally, the Hessian (or second-order gradient) of the field can be found from Eq. 2.19 as:

$$\nabla \nabla d(\mathbf{x}) = \frac{1}{\Delta x^2} \begin{bmatrix} 0 & \kappa_4 + \kappa_7 \hat{z} & \kappa_6 + \kappa_7 \hat{y} \\ \kappa_4 + \kappa_7 \hat{z} & 0 & \kappa_5 + \kappa_7 \hat{x} \\ \kappa_6 + \kappa_7 \hat{y} & \kappa_5 + \kappa_7 \hat{x} & 0 \end{bmatrix} \quad (2.20)$$

By using Eqs. 2.16 through 2.19, the pre-computed values of the SDF can be interpolated at some intermediate \mathbf{x} point, yielding both the minimum distance between \mathbf{x} and the small-body surface, as well as the corresponding normal. Although it was stated in Eq. 2.12 that signed distance *functions* satisfy $\|\nabla d(\mathbf{x})\| = 1$, the application of a linear interpolator breaks this property. The

surface gradient obtained with Eq. 2.19 must therefore be normalized in order to obtain the surface normal $\hat{\mathbf{N}}$ at \mathbf{x} , *i.e.*:

$$\hat{\mathbf{N}} = \frac{\nabla d(\mathbf{x})}{\|\nabla d(\mathbf{x})\|} \quad (2.21)$$

Similarly, Eq. 2.13 must be iterated in order to converge on the surface crossing \mathbf{c}^* along some surface normal, when using the SDF.

Discussion One important detail that has not yet been addressed is the generation of the initial SDF mesh point values $d(i, j, k)$. Fortunately, extensive literature exists on the conversion of polyhedron models into signed distance fields, see for example Osher and Fedkiw [84], Sigg, Peikert, and Gross [107], or Jones, Bærentzen, and Šráníek [51]. Here, the open-source **SDFGen** tool by Batty [8] is used. This allows for a sampling of the signed distance *function* of some polyhedron model at the desired mesh points, and the construction of the corresponding signed distance *field*. The resulting SDF can then be interpolated during simulation. As an example, slices of the equivalent SDF model of 67P/C-G are shown in Fig. 2.11.




Figure 2.11: Sections of the SDF of comet 67P/C-G.

At first glance, this approach may appear to be a simple recasting of the polyhedron model. However, consider the following: when using a polyhedron shape model, distance computation requires iteration over a set of surface features, either the entire shape or some subset of the shape. Generally speaking, the numerical burden of this iteration scales with r^3 when the resolution r of the underlying shape is increased. In contrast, when using the SDF, a distance computation involves only a single, cheap evaluation of the tri-linear interpolant with Eq. 2.17. The cost of this computation is *independent* of the resolution of the SDF. Although the use of a higher resolution SDF will increase the pre-computational cost, the simulation-time effort of a low- and high-resolution model will be the same, assuming no procedural noise is added.

On a final note, consider briefly the physical shape of the surface that is represented by an SDF. For a polyhedron, the shape is obvious: it is simply a collection of flat, triangular faces. For the SDF, the shape is not obvious. Indeed, Lopes and Brodlie point out that the use of a tri-linear interpolator may result in complex shapes [61]. The particular shape of the isosurface within any given voxel will differ depending on which of its eight $d_{i,j,k}$ values are positive and which are negative. As an example, consider Fig. 2.12, taken from [61]. This shows that the SDF is indeed capable of representing quite complex shapes. However, given the relative smoothness of global small-body shape models, most of the voxels in a small-body SDF take the form of examples 1, 2, 5, or 8 from Fig. 2.12, and are relatively (though not perfectly) flat. This makes sense, since the SDF shape models are constructed from a base polyhedron shape with flat facets.

2.3 Statistical features

The small-body shape models discussed in Section 2.2 are able to capture the complex and irregular shape that many of these bodies have. Any such model has a corresponding resolution r that relates to the smallest surface detail resolved by the model. The resolution of a polyhedron model is quantified by the size of its facets, though it is possible for this facet size to vary significantly across any given shape. For an SDF, the resolution is simply expressed by the size of its voxels. At low resolutions, these models must have smooth surfaces without small features such as rocks,




Figure 2.12: Surface shapes captured by a tri-linear interpolator. Taken from [61].

boulders, or cliffs. In order to resolve these features, the model resolution must be decreased. When the model resolution is halved, it generally holds that the number of model elements multiplies eightfold. This is most clear for the SDF, where the number of voxels $r \propto r^3$. As a result, the memory storage requirements of a small-body shape model grow strongly as it is built to resolve smaller and smaller features. At some limit feature size, the model will occupy so much computer memory that simulation of surface exploration probes is no longer possible. A second, more external limit on the best-available model resolution also exists: in order to resolve small surface features such as rocks, a spacecraft must image the small-body surface from low altitude. This operation is often perceived as risky, such that high-resolution imagery is only taken at a few sites of interest and is not available globally; shape models therefore cannot include small rocks on the target surface. Even if the high-resolution imaging could be performed risk-free, the sheer physical size of the produced image files could not be returned to Earth in any practical amount of time.

Nevertheless, the presence of small surface features may significantly affect the motion of a deployed payload. This is obvious from intuition on Earth: a basketball thrown onto a flat surface displays distinctly different bouncing motion compared to a ball thrown onto a boulder field. This poses a challenge: how can the presence of rocks, boulders, and similar features on the small-body surface be accounted for during simulation, without exceeding the available computer memory, or without information on the precise location of such features?

One solution to this problem is the use of a *stochastic* rock model, as done by Tardivel et al. [118] for spherical probes. In this model, a large number of test simulations on a small, flat, and rocky surface are used to generate statistics of spacecraft-rock collisions, such as the distribution of the impact point across the probe shell. These statistics are then used in deployment simulations to impose a justified degree of randomness on surface collisions. Using this strategy, it is possible to account for the presence of surface rocks without requiring creation and storage of any rocks. In other words, the stochastic model uses *virtual* rocks that impose a justified degree of randomness on the impulses imparted on the lander during collision. A similar technique can be used to account for the presence of rocks during continued contact motion, using the mean free path traveled by a surface exploration probe before impacting a rock. Unfortunately, stochastic models have a number of drawbacks, the most important of which relates to the conditions under which a probe can settle on the small body surface; Fig. 2.13 is used to illustrate this point.

When instead using a *persistent* rock model that spawns rocks onto the small-body surface, it is possible for a probe to become lodged against a rock and settle on some steep slope δ , as seen on the left side of Fig. 2.13. When using a stochastic model instead, the probe does not encounter a solid rock, but is instead subjected to a random impulse that will likely send it on a brief ballistic arc, as shown on the right side of Fig. 2.13. The probe cannot settle on the steep slope, as there is no surface against which it can become lodged; the simulation effectively ‘forgets’ about the existence of the rock as soon as the impulse has been applied. As a result, the stochastic model is biased towards settling regions with a low slope. Similarly, it is possible for grazing impacts to be ignored when using a stochastic model, whereas a collision with a rock would be detected if

instead using the persistent rock model. This is illustrated in Fig. 2.14. These limitations should be most prominent when the range of motion of the probe, *i.e.*, its velocity is small, but may significantly affect the settling statistics of a batch of simulations. This will be shown further on, in Section 4.1. In order to generate the proper level of simulation fidelity, it is thus considered necessary to use some form of persistent rock modeling in our simulations. An additional reason to prefer the use of a persistent rock model over a stochastic one relates to the manner in which the latter is constructed. Given that stochastic rock models are defined using results from some sample set of simulations with a specific probe and rock field, the model cannot be easily adjusted to account for other probes and rock fields. Instead, the sample simulations must be repeated and statistically fitted to some distribution. This is an impractical process when one attempts to do broad studies in which both the probe shape and rock field are continuously adjusted.




Figure 2.13: Stochastic rock models bias lander settling towards low-slope regions.




Figure 2.14: Stochastic rock models may fail to capture grazing impacts.

2.3.1 Polyhedral rocks

As mentioned, it is numerically infeasible to generate and store rocks on the entire polyhedron model of a small body. Fortunately, our subdivision of the global surface into smaller local worlds using the atlas discussed in Section 2.2.2 may be exploited to enable a suitable rock generation method. Recall that a probe can only collide with the *active* local world; it is therefore sufficient to only create and place rocks on the active world. Rock distributions on Earth typically follow a power law; this relationship was also observed by the Hayabusa spacecraft on asteroid Itokawa [68]. The total number of rocks K larger than some size s_{\min} per unit surface area is given by a power law with power index α as [118]:

$$K = k_0 \left(\frac{s_0}{s_{\min}} \right)^\alpha \quad (2.22)$$

in which k_0 is the number of rocks larger than some size s_0 ; these two numbers effectively ‘initialize’ the distribution governed by α . Given the surface area of each triangular facet in the active local world, Eq. 2.22 determines the number of rocks larger than some size s on that facet. The size of each rock is described by a Pareto random number, which can be computed from a uniform random number $y = \mathcal{U}(0, 1)$, as [118]:

$$s = \frac{s_{\min}}{(1 - \beta y)^{1/\alpha}} \quad \text{with} \quad \beta = 1 - \left(\frac{s_{\min}}{s_{\max}} \right)^\alpha \quad (2.23)$$

in which s_{\min} and s_{\max} are the minimum and maximum allowed size of the rocks. Each rock is initially shaped as a regular icosahedron with diameter s , whose vertices are then randomly deformed following a Gaussian deformation with zero mean and standard deviation σ along each Cartesian axis. The rocks are then randomly rotated and randomly placed on the considered facet. Finally, the rocks are sunk to some random depth, uniformly generated between r_{\min} and r_{\max} , which are expressed as a ratio of the corresponding rock size s . By repeating this process for all rocks on all facets of the active local world, a local world can be populated with rocks following a particular distribution. The creation process is illustrated for a single rock in Fig. 2.15; a sample rock-populated local world is shown on the right side of Fig. 2.16.




Figure 2.15: Illustration of the rock creation algorithm.




Figure 2.16: (*left*) Itokawa's rocky surface [74] and (*right*) active local world with procedurally generated rocks.

This technique is effective at quickly populating a local world on a polyhedron model, based on a handful of parameters and a series of random numbers. In order to ensure consistent and repeatable generation of rocks, the random number generator (RNG) is reinitialized before generating rocks on each facet. The seed of this initialization is equal to the identifying number of the considered facet. As that number is unique to each facet, this RNG initialization ensures that the same set of rocks will always be generated on the same facet, across all simulated trajectories, without needing to ‘remember’ previously generated rocks in memory. The rock distribution is effectively stored in the algorithm used to generate the (pseudo-)random numbers.

In order to minimize the computational effort required to detection collisions between a probe and rocks on a local world, each rocks given its own bounding sphere, in analogy with the local world bounding spheres discussed in Section 2.2.2. These bounding spheres are again used to switch between coarse and fine distance checks. The use of bounding spheres is extremely important in

this case, as a local world may contain a large number of rocks, such that performing fine distance checks to each individual rock would be resource-consuming.

2.3.2 Fractal noise

The procedural rock creation algorithm discussed above can be used to create *explicit*, polyhedral rocks. Although this technique is successful, it carries a significant computational load, in particular when the considered rock distribution includes many features. When discussing small-body shape models in Section 2.2, the use of the *implicit* SDF shape model was found to provide significant advantages over the commonly-used explicit polyhedron shape model. In a similar sense, it is possible to instead use *implicit* models to account for small features on a given surface. For this, *procedural modeling techniques* are used. We quote Ebert et al. from their book on procedural modeling and repeat the following: “*One of the most important features of procedural techniques is abstraction. In a procedural approach, rather than explicitly specifying and storing all the complex details of a scene, they are abstracted into a function or algorithm (i.e., a procedure) and evaluate that procedure when and where needed. We gain a storage savings, as the details are no longer explicitly specified but implicit in the procedure, and the time requirements for specification of details are shifted from the programmer to the computer. This allows us to create inherent multi-resolution models and textures that we can evaluate to the resolution needed.*” [30, p. 2].

Procedural techniques are used extensively for the generation of a variety of digital features, including, but not limited to: textures, flowers, trees, and terrain. The application of procedural techniques for the generation of one- or two-dimensional terrain features typically takes the form of a *height field*. In this, we start with some externally-defined base terrain; in the simplest form, this is a flat plane with height $z = 0$ everywhere. Some procedural technique is then employed to modify the height field and generate some $z(x, y)$ to introduce certain features onto the plane. A variety of such procedures exist, though the some of the most commonly applied ones make use of *fractals*. This is because fractals allow for the creation of realistic terrain features, since most natural features on Earth are of fractal nature and display a high degree of self-similarity.

For example, the relative roughness present in large ridges on a mountainside is also present at a smaller scales in boulders on that mountain, and again in smaller rocks, and again in grains of sand. Effectively, this kind of fractal geometry consists of a repetition of some pattern at different scales. To illustrate how these fractals may be used to generate surface features on a small body, a simple one-dimensional example is first examined.

Uniform noise field Consider a simple flat line, as marked with dots on the top of Fig. 2.17. Sixteen uniform random numbers $\mathcal{U}(-0.5, 0.5)$ are generated, which provide height values above the line of 16 points spaced $\Delta x = 1$ m apart. When one linearly interpolates between these points along the x -direction, a series of jagged peaks appears, as done on the top of Fig. 2.17. In order to generate smoother features, a higher-order interpolant could be used. Although this is relatively straightforward in this one-dimensional example, higher-order interpolation in three dimensions becomes quite mathematically involved and may significantly slow down collision detection during simulation. Instead, the interpolator is therefore modified with a quintic function that smooths the values of the height field [48]. This results in the smoother terrain features shown in the middle of Fig. 2.17. For clarity, the equations used to linearly interpolate the random numbers are provided, both without (*lerp*) and with (*qlerp*) the quintic smoother:

$$\begin{cases} \bar{h}_{lerp}(x) = h_0 + t \cdot (h_1 - h_0) & \text{with } t = \frac{x-x_0}{x_1-x_0} \\ \bar{h}_{qlerp}(x) = h_0 + t_q \cdot (h_1 - h_0) & \text{with } t_q = t^3 \cdot (t \cdot (6t - 15) + 10) \end{cases} \quad (2.24)$$

Note that the application of the quintic curve does not change the extremum values of this random ‘noise’ field; it merely smooths the peaks and values such that they do not appear as jagged. Finally, also note that in the top two plots of Fig. 2.17, the vertical scale has been exaggerated. The true scale is visible in the plot at the bottom of the figure. In summary, this simple example shows how a small set of random numbers can be used to perturb a flat line into a noisy, bumpy surface. Using the modulo operator, this piece of surface can be easily repeated without needing to store any additional random numbers. This is also illustrated in Fig. 2.17, where blue lines indicate a repetition of the original 16-meter domain covered by the 16 random numbers.




Figure 2.17: Interpolation of a one-dimensional noise field.

Amplitude and frequency modulation Fig. 2.17 illustrated a base noise field in which random numbers were placed $\Delta x = 1$ m with a maximum height different of $\Delta h = 1$ m. This fixes the horizontal and vertical sizes of the created features. The same set of 16 random numbers may also be used to create features of different sizes, by modulating the original noise field with some amplitude A and frequency f , as:

$$h_{qlerp}(x) = A \cdot \bar{h}_{qlerp}(f \cdot x) \quad (2.25)$$

Figs. 2.18 and 2.19 illustrate the effect of using, respectively, different A and f values on the resulting noise field. In other words, the size of the generated features can be controlled by adjusting A and f , in order to match some realistic distribution observed on a small body.

Fractional Brownian motion The above discussion illustrated independent variation of the amplitude A and frequency f in order to modulate some source noise field. By adjusting these two parameters at the same time, it is possible to generate features of varying sizes. In Fig. 2.20, four (individual) *octaves* of noise corresponding to different sets of (A, f) are illustrated.




Figure 2.18: Interpolation of a one-dimensional noise field for varying A , with fixed $f = 1 \text{ m}^{-1}$.




Figure 2.19: Interpolation of a one-dimensional noise field for varying f , with fixed $A = 1 \text{ m}$.

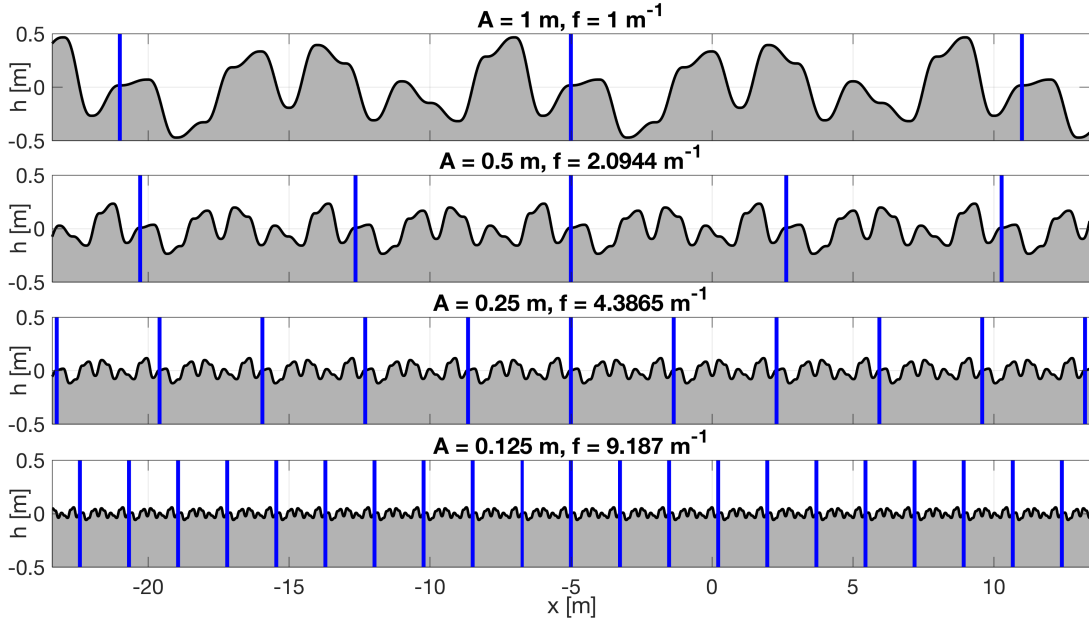


Figure 2.20: Interpolation of a one-dimensional noise field for varying (A, f) pairs.

Although it is possible to arbitrarily change A and f between successive octaves, the commonly-used *fractional Brownian motion* algorithm specifies that the amplitude A_i and frequency f_i of the i -th octave of noise are given by:

$$\begin{aligned}
 A_i &= A_0 \cdot L^{-Hi} \\
 f_i &= f_0 \cdot L^i
 \end{aligned}
 \tag{2.26}$$

in which L is called the *lacunarity* and H is the fractal increment. Barnsley et al. note that “the value $H \sim 0.8$ is empirically a good choice for many natural phenomena” such as mountain ranges, coastlines, cliffs, boulders, rocks, and pebbles [31]. Furthermore, Ebert et al. note that a lacunarity of $L \sim 2$ works well for most purposes [30]. In order to avoid overlapping the frequencies, an irrational lacunarity of $L = 2\pi/3 \approx 2.0744$ is chosen. Using these values restricts the number of free parameters to the initial amplitude A_0 , the initial frequency f_0 , and the number of octaves n to consider. When multiple of these octaves are added together of noise, as is illustrated in Fig. 2.21, the result is a random-looking surface that does not repeat, but is nevertheless parametrizable with only three numbers. In terms of memory storage, it only requires storage of the 16 initial random numbers to be modulated (this number can be increased at will to further increase the degree of randomness of the surface).




Figure 2.21: Interpolation of a one-dimensional noise field for increasing number of octaves.

Extension to 2D and 3D Fractional Brownian motion can also be applied in a straightforward manner to generate surface detail in two dimensions. Starting with a flat plane, a two-dimensional grid of random numbers can be modulated and interpolated, analogous to the one-dimensional example shown above, to generate a height field above the plane. This is illustrated in Fig. 2.22, where a standard basketball is included for scale.




Figure 2.22: Interpolation of a two-dimensional noise field for increasing number of octaves.

Finally, in order to apply fractional Brownian motion to a three-dimensional object, the signed distance field can be used. A three-dimensional grid of random numbers is first generated, which is then interpolated linearly in analogy with our previous discussions. The interpolated noise field is then superimposed onto the SDF, such that the signed distance to the noisy surface is computed as:

$$d(\mathbf{x}) = d_{SDF}(\mathbf{x}) + d_{fBm}(\mathbf{x}) \quad \text{with} \quad d_{fBm}(\mathbf{x}) = \sum_{i=0}^n A_i \cdot u(f_i \cdot \mathbf{x}) \quad (2.27)$$

in which $u(\dots)$ is the function that samples and linearly interpolates the three-dimensional random number grid. If this grid is built on random numbers $\mathcal{U}(-0.5, 0.5)$, the mean noise perturbation must be zero. In other words, the expected value of the target volume remains unchanged by this creation of random surface features. As a final illustration, Fig. 2.23 shows how the small-body surface changes as an increasing number of noise octaves are added.




Figure 2.23: Interpolation of a three-dimensional noise field for increasing number of octaves..

Noise resolution Using fractional Brownian motion, it is possible to generate surface detail of arbitrarily large or small size. It is therefore natural to ask: what is the size range across which features should be generated on a small body? An upper size limit can be established by determining the resolution of the small-body shape model on which features are to be generated. Given that our SDF implementation is generated from polyhedron shape models, it is instructive to examine the resolution statistics of such a model. The surface area A_f of a triangular facet can

be computed as:

$$A_f = \frac{1}{2} (\mathbf{P}_1 - \mathbf{P}_0) \times (\mathbf{P}_2 - \mathbf{P}_0) \quad (2.28)$$

in which \mathbf{P}_0 , \mathbf{P}_1 , and \mathbf{P}_2 are the three vertices of the considered facet. Assuming that the facets are equilateral triangles, their corresponding side length s_f is found as:

$$s_f = 2\sqrt{\frac{A_f}{3}} \quad (2.29)$$

As an example, Fig. 2.24 provides histograms of the facet area and side length distribution of Itokawa’s 320,000-facet model. The figure also graphically shows the distribution of the facet areas. The mean side length is $\bar{s} = 1.61$ m, with a standard deviation of $\sigma_s = 0.50$ m. The facets tend to be larger in the relatively flat Sagamihara and Muses Sea regions. These statistics suggest that the SDF of asteroid Itokawa will approximately match the resolution of the source polyhedron model when using a voxel resolution of $\Delta x = 1.6$ m.




Figure 2.24: Facet resolutions on Itokawa.

Based on this resolution, one may be inclined to generate surface features up to 1.6 m in size on Itokawa’s surface. However, the shape models of most small bodies only account for surface features several times larger than their actual spatial resolution. Hudson et al. state this in their paper on the development of a shape model of asteroid Toutatis, by saying that “*It is much more likely that Toutatis does not have structure that is evident in our model: we have chosen to prefer understating relief rather than to risk overstating it.*” [47] In other words, small-body shape models tend to have an actual resolution several times greater than what is implied by the facet size. For

Itokawa, features up to $3 \cdot s = 4.8$ m in height and width are therefore included onto the asteroid's base shape model. In order to select the fBm parameters that generate features up to some h_{max} in size, recognize that this h_{max} is a simple geometric series:

$$h_{max} = A_0 + A_0 \cdot L^{-H} + A_0 \cdot L^{-2H} + A_0 \cdot L^{-3H} + \dots \quad (2.30)$$

This series can be rewritten as:

$$h_{max} = A_0 \sum_{i=0}^n L^{-Hi} = \frac{A_0 \cdot L^{-Hi} (L^{H(n+1)} - 1)}{L^H - 1} \quad (2.31)$$

Rewriting this expression, the initial amplitude A_0 that should be used to generate a maximum height of h_{max} in n octaves can be determined, given some lacunarity L and fractal increment H :

$$A_0 = \frac{h_{max} (L^H - 1)}{L^{-Hi} (L^{H(n+1)} - 1)} \quad (2.32)$$

Probe resolution Section 2.5 detailed the contact model used to evaluate collisions and continued contact motion between a surface exploration probe and the small-body surface. The contact forces exerted on the probe by the surface act at the vertices of the former. In other words, the surface effectively only ‘sees’ the vertices of the probe shape model. This is not so much a choice made in the development of the applied contact model, but rather a result from the use of the implicit SDF shape model. As previously illustrated in Fig. 2.12, the surface expressed by an SDF is rarely every perfectly flat, even if the underlying polyhedron model from which it is constructed *is* flat. As a result, it is generally not possible for *e.g.* an edge of the spacecraft to be fully in contact with the SDF surface, without allowing penetration of that surface. Nevertheless, it is possible for any point along the considered edge to be in contact with the surface. As a result, the number of vertices distributed across the shell of a considered probe affects that probe’s ability to ‘resolve’ surface features of a given size. This is more clearly illustrated in Fig. 2.25, which depicts the surface impact of two different probes: one low-resolution probe with vertices only at its corners, and one high-resolution probe with additional vertices distributed across its faces. It can be seen that surface features smaller than the probe are able to penetrate the shell of the low-resolution probe quite significantly before contact is detected. In contrast, the high-resolution

probe interacts much more realistically with the surface feature. Furthermore, the low-resolution probe must always impact on its corners, while the high-resolution probe can impact in the middle of an edge or facet. Although such impacts do not occur for smooth surfaces that undulate at spatial periods several times greater than the probe size, they may occur often for surfaces with small surface features.

The reason for this is that with the SDF, because of the shape it has, you cannot have a hard contact with an entire edge or facet. It's because the SDF surface is almost never perfectly flat. Furthermore, it's easy to know what geometry between a triangular facet and the edge of a lander actually touches along the edge. This isn't true for the SDF, because the edge is a continuous object. But the SDF varies, so the intersection between e.g. an SDF and a plane is not obvious.

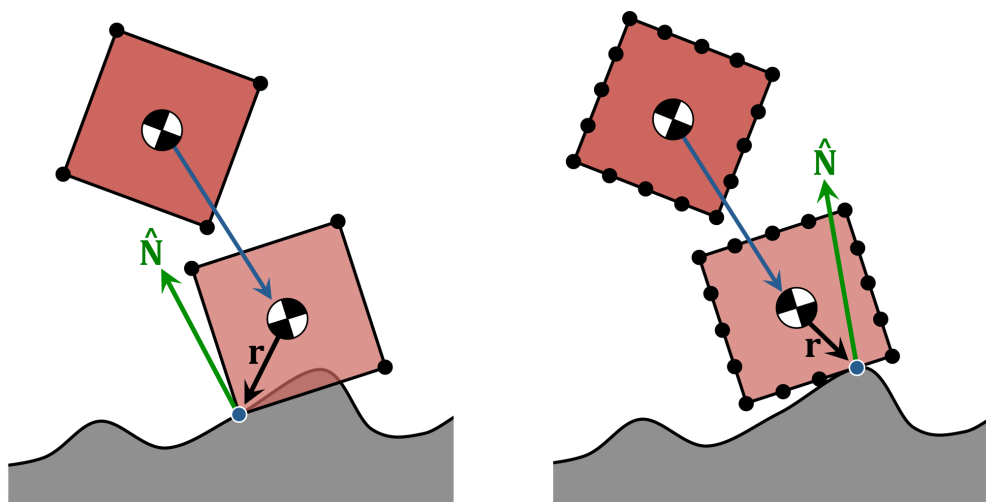


Figure 2.25: (left) Low- and (right) high-resolution probe models resolve surface impacts differently.

It is clear from the above illustration that the geometry between the surface normal $\hat{\mathbf{N}}$ and the collision vector \mathbf{r} will differ between low- and high-resolution probe shapes when small surface features are included. For the particular example shown, note that the normal force creates a clockwise (negative) torque on the low-resolution model, but a counterclockwise (positive) torque on the high-resolution model. It is stressed again that the consideration of probe resolution is only relevant when the target surface contains features smaller than the probe itself.

In order to properly resolve collisions when this occurs, the probe shapes are tessellated using a midpoint triangulation. Fig. 2.26 illustrates the generation of additional vertices across the shell of the octahedron as repeated tessellations are performed. To quantify the tessellation level required to resolve surface features of a given size, the mean distance between adjacent vertices of the five platonic solids with $m = 5 \text{ kg}$ and $\rho = 500 \text{ kg/m}^3$ is computed. These are plotted in Fig. 2.27 for increasing tessellation levels. This shows that the vertex resolution decreases exponentially as with respect to the tessellation level.




Figure 2.26: Progressive tessellation of an octahedron.




Figure 2.27: Vertex resolution as a function of the tessellation level for the platonic solids.

The trends of this figure suggest that the tetrahedron requires 2 additional tessellations in order to achieve the same vertex resolution as the dodecahedron. This result makes intuitive sense: near-spherical shapes such as the dodecahedron have a lower area-to-mass ratio than less spherical shapes such as the tetrahedron. If a shape has lower surface area, it also requires fewer vertices to achieve some inter-vertex distance than a shape with higher surface area.

Even when using a tessellated probe shape, one final issue must be addressed in order to properly resolve impacts on a noisy surface; Fig. 2.28 is used to illustrate this. Consider the left side of the figure, where a rectangular probe impacts a tiny surface feature. Although the probe is tessellated to allow for the resolution of small features, the considered feature is still small enough to partially penetrate into the probe before impact is detected. As a result, the use of the surface normal $\hat{\mathbf{N}}_0$ is incorrect: if the vertex resolution were infinitely small such that no penetration occurs, the surface normal at impact would be equal to the normal $\hat{\mathbf{N}}_1$ of the impacted edge. One final adjustment to our collision resolution is therefore made: whenever an impact is detected, the vector $\hat{\mathbf{N}}$, that defines the normal and tangential directions in which respectively restitution and friction act, is selected based on the ‘type’ of probe vertex that is impacted:

$$\text{Impact on a } \begin{cases} \text{corner vertex:} & \text{use the surface normal } \hat{\mathbf{N}}_0 = \nabla d(\mathbf{x}) \\ \text{edge vertex:} & \text{use the modified surface normal } \mathbf{N}_1 = \hat{\mathbf{N}}_0 - (\hat{\mathbf{N}}_0 \cdot \hat{\mathbf{e}}) \hat{\mathbf{e}} \\ \text{facet vertex:} & \text{use the } \textit{negative} \text{ facet normal } -\hat{\mathbf{N}}_{\mathcal{F}} \end{cases} \quad (2.33)$$

Note that the modified surface normal \mathbf{N}_1 in Eq. 2.33 must be normalized before use. On the right side of Fig. 2.28, the three vertex types are illustrated on a tessellated cube.




Figure 2.28: Modification of the impact normal.

2.3.3 Procedurally seeded rocks

The two previous sections have detailed different ways of including statistical surface features onto the small-body surface. The first, the polyhedral rock model, can be easily tuned to match observations of some target body, as made by an orbiting spacecraft. This rock model is *explicit*: the individual rocks consist of vertices and facets; collision detection is expensive for such models. This is especially true when the rocks are very densely distributed across the surface, or if very small rocks (of which there may be very many) are considered. Furthermore, this type of explicit model is incompatible with the implicit SDF shape model. This makes it poorly adaptable to fast, parallelized simulations. The second model, using fractional Brownian motion, is a fully procedural technique that generates surface features at very low computational cost. However, due to the ‘random’ nature of the model, it is not obvious how the fBm parameters may be tuned in order to match surface observations. Furthermore, fBm affects the entire small-body surface, whereas in reality the surface may be flat with dispersed rocks placed on top of that flat surface.

Having demonstrated the use of both techniques, it is natural to wonder: is there some way in which the polyhedral rock and fractional Brownian motion models can be combined, in order to produce a fully procedural, yet tuneable rock model? In other words, is it possible to convert the polyhedral rock model into a distance field-like procedure?

Seeded textures The fractional Brownian motion features are constructed from an initial ‘texture’ consisting of uniform random numbers. By modulating these numbers at different amplitudes and frequencies, and simply summing the different octaves of noise to the SDF model, small surface features can be generated. Instead of using random numbers to construct the noise texture, let us consider applying a more structured approach, and make use of the distance field around a set of rocks. To understand this process, consider first a two-dimensional example, similar to that shown previously in Fig. 2.22 for the fBm noise. This example is shown in Fig. 2.29, which visualizes the height field corresponding to a population of spherical rocks on a flat plane. This texture is also shown from a top view on the left side of Fig. 2.30.

The texture is constructed with periodic boundaries, such that it can be tiled indefinitely to populate an infinite plane with rocks, as shown on the right side of Fig. 2.30. By adjusting the size and spatial density of rocks in the small initial texture, the rock population can be tuned to match some particular distribution. This is a promising results, though it is obvious that the periodic tiling does not generate a ‘realistic’ population of rocks; a real rock field will be much more randomized.




Figure 2.29: Height field of a spherical rock field.




Figure 2.30: Periodic tiling of the spherical rock field.

To mitigate this issue, the methodology by Quilez [49] to aperiodically tile a two-dimensional texture is followed; the method easily extends to the three-dimensional case as well. For the moment, though, let us continue to investigate the two-dimensional example. In order to aperiodically tile the rock texture, a second texture that contains uniform random numbers between 0 and 8 is first sampled. These random numbers are used to establish some *texture index*, i , across the considered surface, as illustrated on the left side of Fig. 2.31. This index will have some integer part, $k = \text{floor}(i)$ and some fractional part, $q = i - k$.

Eight distinct tilings of the noise texture are now defined. Although all eight tilings are periodic, they are performed with some pseudorandom offset. That is, the noise texture sampling at position \mathbf{x} is performed at position \mathbf{x}_k in the respective eight tilings:

$$\mathbf{x}_k = [\mathbf{C}(k)] (\mathbf{x} + \mathbf{o}_k) \quad (2.34)$$

in which the rotation matrix $[\mathbf{C}(k)]$ is defined as:

$$[\mathbf{C}(k)] = \begin{bmatrix} \cos(k) & \sin(k) \\ -\sin(k) & \cos(k) \end{bmatrix} \quad (2.35)$$

and in which the offset vector \mathbf{o}_k is defined as:

$$\mathbf{o}_k = \begin{bmatrix} \sin(3k) \\ \sin(7k) \end{bmatrix} \quad (2.36)$$

In these expressions, the sine function effectively acts as a *hashing* that generates a random offset and orientation for each of the eight periodic tilings of the rock texture. In other words, they produce eight versions of the right side of Fig. 2.30, each of which has a random offset and rotation.

Using these eight *periodic* tilings, an *aperiodic* tiling can be achieved by sampling each of the eight versions of the texture in the respective regions defined by the integer part, k , of the texture index i . By using the fractional part, q , of the index, it is possible to smoothly blend between the eight periodic tilings, in order to obtain the aperiodic tiling shown on the right side of Fig. 2.31 and in Fig. 2.32. For more detail on the blending algorithm, the reader is referred to [49].

Although the tiling will still repeat over the scale of the random number texture from which the tiling indices i are obtained, the sampling frequency of this texture is generally low enough such that that spatial period of repetition is very large. The repetition will therefore not be noticeable over the scale of probe deployment simulations.




Figure 2.31: (left) Texture index and (right) corresponding aperiodic tiling of spherical rocks.




Figure 2.32: Three-dimensional view of the aperiodic tiling of spherical rocks.

Extension to 3D The extension of the two-dimensional to a three-dimensional rock population is achieved by following the same approach of the fBm, namely, by replacing the *height* field with a *distance* field. In order to do so, the distance function to some complexly-shaped rock is required. Although it is possible to construct the distance field around a set of polyhedral rocks, it is simpler to use the *generalized distance functions* by [4], which are given as:

$$d_{psr}(\mathbf{x}) = \left(\sum_{i=a}^b |\mathbf{n}_i \cdot \mathbf{v}|^p \right)^{1/p} - R \quad \text{with} \quad \mathbf{v} = [\mathbf{C}](\mathbf{x} - \mathbf{x}_0) \quad (2.37)$$

In this expression, \mathbf{x}_0 is the center of the considered rock, $[\mathbf{C}]$ expresses its attitude as a direction cosine matrix, and R is its radius. Furthermore, \mathbf{n}_i are a set of unit vectors that can be found in [4]. The shape of the resulting rock is determined by the range $[a, b]$ of vectors n_i that are considered in the series of Eq. 2.37. A few sample shapes are shown in Fig. 2.33. Finally, the value of the norm p determines the ‘sharpness’ of the considered shape, as illustrated in Fig. 2.34.



Figure 2.33: Realization of generalized distance fields of various shapes with $p = 100$.




Figure 2.34: Distance field of an icosahedron ($a = 4$ and $b = 13$), for different values of the norm p .

It is possible to construct a distance field around a set of rocks using Eq. 2.37. On a two-dimensional texture with surface area A , Eq. 2.22 can be directly used to determine the number of rocks across that texture. For a three-dimensional texture, it becomes necessary to distribute the rocks three-dimensionally. For this, the average diameter \bar{d} of the rocks of the considered population is first computed, using Eq. 2.23. Assuming that the K rocks on the two-dimensional area A fill $\bar{d}L^2$ of volume, an equivalently-filling population within a cube of volume L^3 must then contain $K_3 = K_2 \frac{L}{\bar{d}}$ rocks. The K_3 rocks are then each given a diameter d , using Eq. 2.23, and placed with random position and attitude in the three-dimensional texture. The generalized distance functions are then used to compute the corresponding distance field, which is again made periodic. As an example, Fig. 2.35 shows a sample distance field around a set of rocks. Only the points with negative distance value are shown; the points are also colored according to their z -coordinate for increased perspective. Side views along the three planes can also be found in Fig. 2.36.




Figure 2.35: Three-dimensional view of a distance field used to generate procedurally seeded rocks.




Figure 2.36: Two-dimensional view of a distance field used to generate procedurally seeded rocks.

In order to include these *procedurally seeded rocks* (PSRs) onto a signed distance field shape model, it is necessary to restrict the corresponding rock texture to contain only negative distance values; all positive distance values are set to zero. To illustrate why this is necessary, considered Fig. 2.37. The left side of the figure illustrates PSRs generated with an unmodified distance texture. It can be seen that such a texture generates a ‘continuous’ rock population, in which rocks always touch one another, and under which the original surface is no longer visible. Reviewing the right side of the figure, which uses a strictly negative distance texture, the desired behavior is seen, with dispersed rocks that leave the original surface visible. This approach is preferable for obvious reasons, but does have one downside: it distorts the shape of the considered rocks.




Figure 2.37: PSRs generated using (left) an unmodified and (right) a strictly negative distance field.

To understand why this distortion occurs, consider the simple example of Fig. 2.38. This shows the height field above a simple line (shown in black), and around a circle (shown in red), in the two leftmost subplots. The remaining subplots shown, in red, the circle placed at three different altitudes, and in black, the resulting surface when its height value is added to that of the line. This illustrates how the use of a strictly-negative distance field distorts the shape of the considered rocks. Nevertheless, it is considered to be an efficient method for generating tuneable rock distributions. As an illustration, Fig. 2.39 shows the SDF surface of asteroid Itokawa, with procedurally seeded rocks generated on top. Four models are shown, with different minimum considered rock sizes d_{min} . As expected, more rocks are visible when the minimum rock size is decreased. A detailed view of the densest rock field is also shown in Fig. 2.40 for additional visualization.




Figure 2.38: Illustration of the distortion of rock shapes by a strictly negative distance field.

On a final note, remark that the blending algorithm given in Eqs. 2.34 through 2.36 was two-dimensional. It is easily extended to the three-dimensional distance field by instead using the following rotation matrix (chosen arbitrarily as a combined z - and x -rotation) and offset vector:

$$[\mathbf{C}(k)] = [\mathbf{C}_x(k)][\mathbf{C}_z(k)] = \begin{bmatrix} \cos k & \cos k \sin k & \sin^2 k \\ -\sin k & \cos k \cos k & \cos k \sin k \\ 0 & -\sin k & \cos k \end{bmatrix} \quad (2.38)$$

and in which the offset vector \mathbf{o}_k is defined as:

$$\mathbf{o}_k = \begin{bmatrix} \sin 3k \\ \sin 5k \\ \sin 7k \end{bmatrix} \quad (2.39)$$

These rotation matrix and offset vector definitions were applied to generate the aperiodic three-dimensional tiling of procedurally seeded rocks shown in Figs. 2.39 and 2.40. This model can be used to populate the small-body surface, represented with a signed distance field, with procedurally-generated rocks, in accordance to some observed surface distribution.




Figure 2.39: Procedurally seeded rocks with varying minimum rock size d_{\min} .




Figure 2.40: Detailed view of procedurally seeded rocks with $d_{\min} = 0.10$ cm.

2.4 Gravity

Given that all mass has a gravitational field, it holds that irregular distributions of mass have irregular gravitational fields. Although the spherical and ellipsoidal harmonic models have useful applications to orbits about a small body, these models fail when evaluated within the circumscribing (or Brillouin) sphere around a target body. This renders them unusable for the simulation of small-body surface exploration. Instead, the constant-density polyhedron developed by Werner and Scheeres [140] is used, which converges both inside and outside of a small-body shape. It is stressed that this gravity field model is applied separately from the shape modeling, *i.e.*, it is possible to model a body's shape with the SDF while modeling its gravity with a polyhedron. The equations that evaluate the polyhedron gravity are first reviewed, followed by a quantification of the effects of the corresponding model resolution. Finally, the voxelization of the polyhedron gravity model is investigated, in order to reduce the run-time cost of gravity field evaluations.

2.4.1 Constant-density polyhedron

The constant-density polyhedron gravity model yields the gravitational field around some polyhedron model with n_V vertices, n_E edges, and n_F facets. Prior to setting the model up for evaluation, its consistency should be checked using Euler's criterion for polyhedra as given by Eq. 2.11. If the model passes this verification check, its edge dyads \mathbf{E}_{ij} and facet dyads \mathbf{F}_{ijk} can be computed as [140]:

$$\begin{aligned}\mathbf{E}_{ij} &= \hat{\mathbf{N}}_{ijk}\hat{\mathbf{M}}_l^T + \hat{\mathbf{N}}_{ijl}\hat{\mathbf{M}}_k^T \\ \mathbf{F}_{ijk} &= \hat{\mathbf{N}}_{ijk}\hat{\mathbf{N}}_{ijk}^T\end{aligned}\tag{2.40}$$

in which $\hat{\mathbf{N}}_{ijk}$ is the *outward-pointing* normal of the facet connecting vertices i , j , and k . These normals can be easily computed with the open-source `MeshLab` software [22]. The $\hat{\mathbf{M}}_k$ vector is the unit vector perpendicular to the edge \mathbf{e}_{ij} and through the vertex \mathbf{P}_k . These vectors are illustrated for two sample facets that share an edge, in Fig. 2.41. Given that the dyads \mathbf{E}_{ij} and \mathbf{F}_{ijk} are invariant (for a fixed shape), they can be pre-computed to reduce the computational overhead.




Figure 2.41: Vectors used in the evaluation of the constant-density polyhedron model.

The evaluation of the gravitational field at some position \mathbf{x} first requires computation of all per-edge factors L_{ij} and per-facet factors ω_{ijk} [140]:

$$L_{ij} = \ln \left(\frac{r_i + r_j + e_{ij}}{r_i + r_j - e_{ij}} \right) \quad (2.41)$$

$$\omega_{ijk} = 2 \arctan \left(\frac{\mathbf{r}_i \cdot (\mathbf{r}_j \times \mathbf{r}_k)}{r_i r_j r_k + r_i (\mathbf{r}_j \cdot \mathbf{r}_k) + r_j (\mathbf{r}_k \cdot \mathbf{r}_i) + r_k (\mathbf{r}_i \cdot \mathbf{r}_j)} \right)$$

in which \mathbf{r}_i is the vector from vertex \mathbf{P}_i to \mathbf{x} , *i.e.*, $\mathbf{r}_i = \mathbf{x} - \mathbf{P}_i$, and $r_i = \|\mathbf{r}_i\|$. Finally, $e_{ij} = \|\mathbf{e}_{ij}\|$ is the length of the edge connecting vertices i and j . Using this, the gravitational potential U , gravitational attraction \mathbf{g} , and gravity gradient matrix $\mathbf{\Gamma}$ can finally be evaluated at \mathbf{x} as [140]:

$$U(\mathbf{x}) = \frac{1}{2} G \rho \sum_{\forall n_E} \mathbf{r}_{ij} \cdot \mathbf{E}_{ij} \cdot \mathbf{r}_{ij} \cdot L_{ij} - \frac{1}{2} G \rho \sum_{\forall n_F} \mathbf{r}_{ijk} \cdot \mathbf{F}_{ijk} \cdot \mathbf{r}_{ijk} \cdot \omega_{ijk}$$

$$\mathbf{g}(\mathbf{x}) = -G \rho \sum_{\forall n_E} \mathbf{E}_{ij} \cdot \mathbf{r}_{ij} \cdot L_{ij} + G \rho \sum_{\forall n_F} \mathbf{F}_{ijk} \cdot \mathbf{r}_{ijk} \cdot \omega_{ijk} \quad (2.42)$$

$$\mathbf{\Gamma}(\mathbf{x}) = G \rho \sum_{\forall n_E} \mathbf{E}_{ij} \cdot L_{ij} - G \rho \sum_{\forall n_F} \mathbf{F}_{ijk} \cdot \omega_{ijk}$$

in which \mathbf{r}_{ij} and \mathbf{r}_{ijk} are, respectively, the vectors from \mathbf{x} to any vertex of the edge connecting vertices i and j , and to any vertex of the facet connecting vertices i , j , and k . Finally, G is the gravitational constant, while ρ is the density of the considered body. As implied by the name *constant-density polyhedron model*, the latter is assumed to be constant throughout the body.

When using high-resolution polyhedra such as that shown on the left side of Fig. 2.3, the corresponding gravity field evaluations are computationally expensive and therefore prohibitive to the speed of spacecraft simulations. In order to reduce this computational cost, two techniques are used.

2.4.2 Linearization

The first of these techniques is the linearization of the gravity field, as introduced by Tardivel et al. [118]. When the spacecraft velocity or integration step size is small (as is the case during continued spacecraft-target contact), the spacecraft moves only a small distance between successive integration steps. As a result, the gravitational field may be approximated at some \mathbf{x} using a previous evaluation at some \mathbf{x}_0 , provided that the distance $\Delta r = \|\mathbf{x} - \mathbf{x}_0\|$ is small. When this is true, the gravitational potential and attraction at \mathbf{x} are approximately equal to:

$$\begin{aligned} U(\mathbf{x}) &\simeq U(\mathbf{x}_0) + (\mathbf{x} - \mathbf{x}_0) \cdot \mathbf{g}(\mathbf{x}_0) + (\mathbf{x} - \mathbf{x}_0) \cdot \mathbf{\Gamma}(\mathbf{x}_0) \cdot (\mathbf{x} - \mathbf{x}_0) \\ \mathbf{g}(\mathbf{x}) &\simeq \mathbf{g}(\mathbf{x}_0) + (\mathbf{x} - \mathbf{x}_0) \cdot \mathbf{\Gamma}(\mathbf{x}_0) \end{aligned} \tag{2.43}$$

2.4.3 Model resolution

Second, it holds that small features on the body surface (which have negligible mass) have a negligible effect on the global gravitational field. A low-resolution polyhedron shape can therefore be used to approximate the ‘true’ gravitational field of a target body. The selection of the resolution of the considered polyhedron model is effectively a trade-off between computation cost and accuracy. In order to provide order-of-magnitude estimates of the inherent error of various model resolutions, the gravity fields of the three considered bodies are investigated. The open-source **MeshLab** software is used to reduce the high-resolution shape models of three sample bodies into lower-resolution models with a smaller number of facets [22]. This is done using a boundary-preserving, topology-preserving, normal-preserving quadratic edge collapse method. In Figs. 2.42 through 2.44, the shape models with $n_{\mathcal{F}} = [1.25, 10, 80] \times 10^3$ of the considered target bodies are shown, namely, asteroid Itokawa, comet 67P/C-G, and Phobos.




Figure 2.42: Itokawa shape models.




Figure 2.43: 67P/C-G shape models.




Figure 2.44: Phobos shape models.

The error resulting from the approximation of the small-body gravity field with a low-resolution model is now quantified. For this, the gravitational acceleration \mathbf{g}_0 yielded by the highest-resolution model is treated as the ‘true’ acceleration, while the ‘approximated’ \mathbf{g} is evaluated with one of the lower-resolution models. Both accelerations are evaluated across the surface of the considered bodies; the relative error of each low-resolution model is defined with respect to the ‘truth’ provided by the highest-resolution model as:

$$\epsilon = \frac{\|\mathbf{g} - \mathbf{g}_0\|}{\|\mathbf{g}_0\|} \quad (2.44)$$

Fig. 2.45 plots the cumulative density function of the relative gravity error distribution of each low-resolution model. This provides a straightforward visual comparison of the performance of the various models. In the figure, the 1% relative error is also marked with a dashed line. This value is used to judge what model resolution provides a sufficiently good approximation of a body's true gravity field. The 1% threshold is selected as an acceptable error, since even high-resolution polyhedron gravity models incur an error from the constant-density assumption. It is argued that local density variations are likely to cause surface gravity variations of at least 1% and thus accept that same error in the low-resolution approximation to allow us to move forward with analysis. For a more detailed discussion of why this error level may be considered acceptable, the reader is referred to [113].




Figure 2.45: Gravity error of various polyhedron model resolutions.

This figure shows notable differences between the three considered bodies. The size differences between Itokawa, 67P/C-G, and Phobos are stressed. These bodies have volumetric radii of respectively 1.6×10^2 , 1.6×10^3 , and 1.1×10^4 m. For a given number of facets, the Phobos model performs best; even the lowest-resolution 1,250-facet model has a relative error just barely exceeding 1%. This excellent performance is attributed to Phobos's near-spherical shape, which can be captured with relatively few facets. Asteroid Itokawa performs worse and requires a model with 20,000 facets to achieve maximum relative surface gravity error of 1%, due to its more irregular and less spherical shape. Finally, comet 67P/C-G performs the poorest and requires a 50,000-facet model to achieve the threshold of 1% maximum relative error, as it has a strongly irregular shape. If the error criterion is slightly modified by requiring only 95% of facets to achieve a relative error smaller than 1%, the required model resolutions do drop significantly. In this case, the 5,000-facet Itokawa and 20,000-facet 67P/C-G models perform sufficiently. As mentioned, the difference between these models can be explained from the complexity of their shapes. A useful characteristic in this is the *sphericity* ψ of a shape, which was originally defined by [138] and expresses the ratio of the surface area of a sphere with equal volume as the considered body to the surface area of the considered body. It can be computed as:

$$\psi = \frac{\pi^{1/3}(6V)^{2/3}}{A} \quad (2.45)$$

Evaluating this expression for the three considered bodies, Phobos is found to have a high sphericity of $\psi = 0.95$, Itokawa has intermediate $\psi = 0.81$, and 67P/C-G has low $\psi = 0.68$. Note that these values were computed with surface models of similar resolution, since the surface area of a shape with fractal-like properties (such as a small body) is dependent on the resolution of the considered shape model, in analogy with the coastline paradox (see [63] for an explanation of this phenomenon). As a final summary, the maximum relative surface gravity error is plotted as a function of the considered polyhedron model resolution, for the three considered bodies, in Fig. 2.46. This figure repeats the result that, for a given model resolution, Phobos has the lowest, Itokawa has intermediate, and 67P/C-G has the highest relative error.




Figure 2.46: Maximum relative gravity error for various polyhedron model resolutions.

2.4.4 Voxelization

Although the use of reduced-resolution polyhedron models can significantly limit the computational cost of gravity field evaluations, this cost may still be relatively high when performing evaluations ‘on the fly’ during simulation. Instead, it is possible to pre-compute the gravity field at the mesh points of a voxel grid and perform cheap run-time interpolations to obtain the gravitational acceleration in a simulation. Three voxelization schemes are considered:

- (1) Voxelization of the gravitational potential $U(\mathbf{x})$; the gravitational acceleration $\mathbf{g}(\mathbf{x})$ is found as the gradient $\nabla U(\mathbf{x})$ of this field. This requires storage of a single value at each mesh point.
- (2) Voxelization of the gravitational acceleration $\mathbf{g}(\mathbf{x})$. This requires storage of three values at each mesh point.
- (3) Voxelization of the gravitational perturbation $\Delta\mathbf{g}(\mathbf{x}) = \mathbf{g}(\mathbf{x}) - \frac{\mu}{\|\mathbf{x}\|^3}\mathbf{x}$. This also requires storage of three values at each mesh point.

Although the first field (voxelized potential) requires the least storage space per mesh point, the derivative operation required to obtain the acceleration $\mathbf{g}(\mathbf{x})$ incurs an error. The voxelized potential field is therefore intuitively expected to require a higher voxel resolution (*i.e.*, smaller Δx) in order to achieve the same accuracy as the second field (voxelized acceleration). This results

in the obvious question: which of the two fields requires the least memory storage to achieve the same gravity field accuracy, the voxelized potential or acceleration? This question will be answered below.

Finally, the voxelized perturbation is considered as a third field. When far away from even an irregular body, its gravitational field can be reasonably approximated by a central gravity field. When closer to the body or on its surface, the gravity field is often strongly perturbed from this central term. We therefore investigate the accuracy of a voxelized gravitational field where only the difference from the respective central gravity term is interpolated. This interpolation is expected to outperform the voxelized acceleration for a given Δx when far away from the considered body, but is of uncertain performance closer to or on the body surface, where the gravitational acceleration can be very non-central.

Acceleration vs. perturbation Given that small-body gravity fields tend to be more irregular on the surface than far away from a body, the inherent error of voxelization is quantified by evaluating the three fields on the small-body surface. However, in order to help clarify the differences between the $\mathbf{g}(\mathbf{x})$ and $\Delta\mathbf{g}(\mathbf{x})$ voxelizations, a model test is first performed in the greater neighborhood around a body. More specifically, both the $\mathbf{g}(\mathbf{x})$ and $\Delta\mathbf{g}(\mathbf{x})$ models are evaluated in the XY-plane surrounding Phobos. The two interpolated gravitational accelerations are then compared with those computed with the original polyhedron gravity model. Fig. 2.47 plots the relative gravity error of both voxelizations, as well as the difference between the two. In the latter, yellow indicates that the $\Delta\mathbf{g}(\mathbf{x})$ voxelization performs better, while green indicates that the $\mathbf{g}(\mathbf{x})$ voxelization performs better. Note that this example uses a relatively large voxel size of $\Delta x = 1,000$ m for visual reasons.

The results show that the $\Delta\mathbf{g}(\mathbf{x})$ voxelization performs best when away from Phobos's surface. On the surface, the $\mathbf{g}(\mathbf{x})$ voxelization performs best. When decreasing the voxel size to a more realistic value, both models are found to perform equally well when evaluated on the surface, while the $\Delta\mathbf{g}(\mathbf{x})$ voxelization continues to perform best when evaluated away from the surface. Since the two schemes have matching memory storage requirements for matching Δx , these results

suggest that it is more beneficial to use the $\Delta\mathbf{g}(\mathbf{x})$ voxelization than the $\mathbf{g}(\mathbf{x})$ voxelization. The model performance away from the surface will then greatly exceed the maximum surface error with respect to which the voxel resolution is selected.




Figure 2.47: Relative gravity error of voxelized acceleration and perturbation models of Phobos.

Voxel resolution To investigate the effects of the gravity voxel size Δx on the inherent interpolation error, the value is non-dimensionalized using the effective radius R of the considered

target bodies as $\Delta x = \eta \cdot R$. The gravitational acceleration on the surface of the target bodies is then evaluated using different η values and compare the \mathbf{g} values with those computed with the original polyhedron gravity model. As discussed above, the surface accelerations of the $\mathbf{g}(\mathbf{x})$ and $\Delta\mathbf{g}(\mathbf{x})$ voxelizations are equivalent; Figs. 2.48 through 2.50 there compare only the $U(\mathbf{x})$ and $\mathbf{g}(\mathbf{x})$ voxelizations. On the top of these figures, the cumulative distribution function of the relative error on the facets of the considered body is again plotted. On the bottom of each figure, the maximum encountered error is plotted as a function of the η parameter. As expected, the relative error decreases when the voxel size $\Delta x = \eta \cdot R$ is decreased. In agreement with the trends previously discussed, Phobos is found to have the lowest relative error for a given η . Itokawa has a larger error and comet 67P/C-G has the largest error. When plotted on a log-log scale, a strongly linear trend between the maximum relative error and the voxel size of an interpolation is found.

Further reviewing these results, the $U(\mathbf{x})$ voxelization with some Δx is found to have approximately the same relative error as the $\mathbf{g}(\mathbf{x})$ with $2\Delta x$. In other words, the voxelized potential requires twice the voxel resolution as the voxelized acceleration in order to achieve the same interpolation accuracy. Given that the $U(\mathbf{x})$ voxelization stores only a single value at each mesh point, while the $\mathbf{g}(\mathbf{x})$ voxelization stores three values at each mesh point, it is natural to wonder which scheme minimizes the required memory size.

For this, n is defined as the number of voxels along one dimension needed to achieve some error criterion when using $\mathbf{g}(\mathbf{x})$ voxelization. Along three dimensions, n^3 voxels are thus required when using $\mathbf{g}(\mathbf{x})$ voxelization. Per our observations from Figs. 2.48 through 2.50, the $U(\mathbf{x})$ voxelization requires $(2n)^3 = 8n^3$ voxels to achieve the same accuracy. Given that the $U(\mathbf{x})$ voxelization stores only a single value at each mesh point, this scheme therefore requires $N_U = 8n^3$ stored values. In contrast, since the $\mathbf{g}(\mathbf{x})$ voxelization stores three values at each mesh point, it requires a total of $N_g = 3n^3$ stored values. However, in order to track the energy of a spacecraft, knowledge of the potential U is also necessary. Therefore, this voxelization must store *four* values at each mesh point, leading to a total of $N_g = 4n^3$. This shows that it is optimal, from a memory storage perspective, to apply the $\mathbf{g}(\mathbf{x})$ voxelization; it requires precisely half the memory of the $U(\mathbf{x})$ voxelization.




Figure 2.48: Itokawa relative surface gravity error for varying voxel grid size.




Figure 2.49: 67P/C-G relative surface gravity error for varying voxel grid size.




Figure 2.50: Phobos relative surface gravity error for varying voxel grid size.

2.5 Contact

The previous sections have provided the equations of motion governing a deployed spacecraft, detailed a technique for detecting the contact between such a spacecraft and an implicitly-defined surface, and provided efficient methods for modeling the small-body gravitational field. As a final module required to simulate probe motion in this environment, a model for the forces and torques present during contact must therefore be included. Before proceeding with a derivation of our applied contact model, some challenges in the evaluation of contact between a spacecraft and target body with arbitrary shapes are highlighted. This is illustrated using Fig. 2.51.

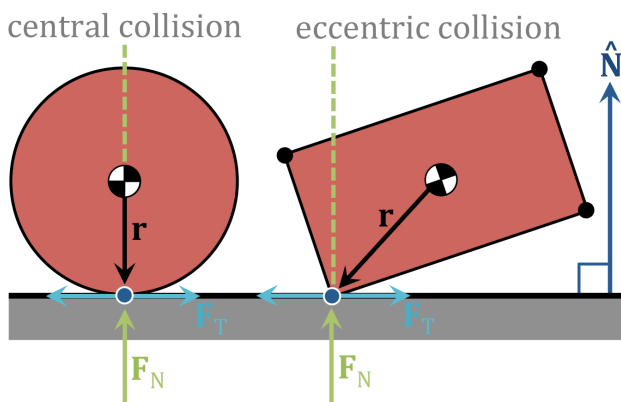


Figure 2.51: Alignment of the collision vector in central and eccentric collisions.

Consider first the spherical probe shown on the left side the figure. When this probe impacts a surface, the *collision vector* \mathbf{r} that points from the probe center of mass to the contact point must always be parallel to the surface normal. This is true regardless of the topography of the underlying surface; it holds even when the spherical probe impacts a sharp rock. This kind of collision, in which $\mathbf{r} \times \hat{\mathbf{N}} = \mathbf{0}$ is known as a *central collision*. Central collisions are relatively easy to resolve, because the normal forces \mathbf{F}_N and tangential forces \mathbf{F}_T typically modeled in impact mechanics are *uncoupled* in such collisions. Although the classical Coulomb friction force is dependent on the applied normal force, the uncoupling of the two forces imply that they can be computed separately. This allows for algebraic collision laws, such as those applied by Tardivel et al. [118], which allow for a fast and straightforward resolution of spacecraft-target impacts. For complex probe shapes

such as that shown on the right side of Fig. 2.51, the situation is far more complicated: \mathbf{r} and $\hat{\mathbf{N}}$ are *not* parallel (*i.e.* $\mathbf{r} \times \hat{\mathbf{N}} \neq \mathbf{0}$), such that any normal force imparts a torque on the probe, which changes the tangential velocity of the contact point. The reverse holds for forces applied in the tangential direction: they change the normal velocity of the contact point. This creates a coupling between normal and tangential forces that renders the collision resolution non-trivial. Precisely how this coupling appears in the mathematics of the collision resolution will be discussed shortly. We suspect that this force-coupling is one of the reasons why we have not found attempts by previous authors to fully account for complex probe shapes in simulations to small-body environments.

Many different rigid- and soft-body collision models exist. Ahmad, Ismail, and Mat provide an excellent overview of different models, with a strong focus on the definition of the commonly-used *coefficient of restitution* [3]. The first coefficient of restitution model was developed by Newton, with his *kinematic* definition relating the incoming and outgoing velocities of a pointmass under impact. Poisson used a more involved *kinetic* definition, relating to the applied impulse at impact. Most recently, Stronge developed an *energetic* definition of the coefficient of restitution, which is the most appropriate definition currently available, according to some authors [3].

In light of Philae’s interaction with both a soft regolith layer and an underlying hard surface layer on comet 67P/C-G, one may be inclined to include separate models for these two layers. A high fidelity modeling of the regolith layer may be achieved through discrete element modeling (DEM), in which individual regolith grains are accounted for, and both their mutual interactions as well as their interaction with a deployed probe simulated. For example, consider the hard-sphere model by Richardson et al. [92] or the soft-sphere model by Sánchez and Scheeres [96]. Although such models can achieve significant simulation fidelity, their integration into our surface exploration framework is problematic due to two distinctly different reasons:

- (1) Granular DEM models require the tracking of several thousand particles at any given time. Accounting for all particle-particle and particle-lander interactions in such a system carries a computational burden that *significantly* reduces the speed at which simulations can be

performed. For example, Thuillet et al. use the `pkdgrav` N-body gravity tree code to simulate the interaction between the MASCOT lander and a granular surface, see also Fig. 2.52. Simulation times of this magnitude prevent the type of broad study that is attempted with our work. Furthermore, it is not obvious how a highly localized DEM model can be integrated onto a global small-body shape model.




Figure 2.52: Snapshots of MASCOT impacting a granular surface, as simulated by Thuillet et al. [122]

- (2) Even when using hard-sphere DEM models that are governed only by a particle-particle restitution and friction coefficient, selecting the particular values of these coefficients remains a significant challenge. In particular, it is not obvious whether these coefficients should have equal values for particle-particle collisions at different velocities and particles size and for particle-probe collisions. If a hard surface layer is added underneath a soft regolith layer, this introduces an additional interface along which collisions occur. Furthermore, these DEM simulations require knowledge of the particle size distribution on the relevant target surface. Given that our current knowledge of these properties is extremely limited, it is difficult to justify the significant effort required to implement a DEM model into our simulation framework.

In light of these two limitations, we instead choose to abstract away the complex impact mechanics that occur between a surface exploration probe and the small-body surface, hard or soft. This is done by using a *hard* contact model with impulsive collisions and continued contact motion. The model is based on the work by Stronge [110]; it is developed for generic purposes, with certain simplifications applied in order to extract both an impulsive collision and a continued contact

model. The models are optimized for evaluation on parallelized GPU hardware. This ‘simplified’ model allows us to investigate how the net effects of different surface properties affect the motion of surface exploration probes. For example, if a thick regolith layer is expected to dissipate significant energy upon impact, the coefficient of restitution can be set to some particularly low value.

Finally, note that it should be possible to implicitly merge high-fidelity DEM models into our simplified contact model. With a separate DEM simulation framework similar to that of [122], it should be possible to simulate probe impacts at different velocities and attitudes. The net restitution and friction coefficients can be extracted from these simulations and empirically fitted to produce a statistical coefficient model. This model could then be implemented into our global simulation framework in order to (statistically) mimic the impact behavior observed in the DEM simulations. The main difficulty of such an approach is accounting for all relevant variables, such as angle of impact, initial velocities, initial attitude, etc., creating an unfeasibly large parameter space. This is considered beyond the scope of the current thesis, but recommended as a piece of future work.

2.5.1 Geometry

In our contact model, the spacecraft shape is defined as a set of vertices. At the beginning of a contact phase, the spacecraft touches the target surface at some *contact point* with position vector \mathbf{H} relative to the target frame, and collision vector \mathbf{r} relative to the spacecraft frame. In other words:

$$\mathbf{H} = \mathbf{x} + \mathbf{r} \tag{2.46}$$

This contact geometry is illustrated in a simple two-dimensional setup in Fig. 2.53. This is done purely for visual purposes; all derivations are performed for the full three-dimensional setup.

The SDF can be evaluated at the contact point to yield the corresponding surface normal $\hat{\mathbf{N}}$ using Eq. 2.19. The contact point has a net velocity \mathbf{v} relative to the target, resulting from the




Figure 2.53: Illustration of the contact geometry.

combination of the linear and angular spacecraft velocities, as:

$$\mathbf{v} = \dot{\mathbf{v}} + \boldsymbol{\omega}_{S/T} \times \mathbf{r} \quad (2.47)$$

The *contact frame* \mathcal{C} is now defined with origin at the contact point and orthonormal axes $\{\hat{\mathbf{u}}_1, \hat{\mathbf{u}}_2, \hat{\mathbf{N}}\}$. In other words, the third component spans the *normal* direction at the contact point, while the first and second components together span the *tangential* plane. This contact frame is included in Fig. 2.53. The two tangential unit vectors are defined as:

$$\hat{\mathbf{u}}_1 = \frac{\hat{\mathbf{N}} \times \hat{\mathbf{t}}_3}{\|\hat{\mathbf{N}} \times \hat{\mathbf{t}}_3\|} \quad \text{and} \quad \hat{\mathbf{u}}_2 = \hat{\mathbf{N}} \times \hat{\mathbf{u}}_1 \quad (2.48)$$

During contact, the spacecraft is subject to a contact force \mathbf{F}_c that acts at the contact point and creates an associated torque $\mathbf{L}_c = \mathbf{r} \times \mathbf{F}_c$. The force and torque change both the linear and angular velocity of the spacecraft, and thus also of the contact point. The change in the contact point velocity \mathbf{v} can be expressed by taking the derivative of Eq. 2.47, as:

$$\dot{\mathbf{v}} = \dot{\mathbf{v}} + \dot{\boldsymbol{\omega}}_{S/T} \times \mathbf{r} + \boldsymbol{\omega}_{S/T} \times \dot{\mathbf{r}} \quad (2.49)$$

The above expression consists of three terms, which are reviewed separately. The first term corresponds to the change in spacecraft linear velocity, which was given in Eq. 2.2 as:

$$\dot{\mathbf{v}} = \mathbf{g}(\mathbf{x}) - [\tilde{\boldsymbol{\Omega}}_{T/N}][\tilde{\boldsymbol{\Omega}}_{T/N}]\mathbf{x} - 2[\tilde{\boldsymbol{\Omega}}_{T/N}]\mathbf{v} + \frac{1}{m}\mathbf{F}_c \quad (2.50)$$

The second term, corresponding to the change in spacecraft angular velocity, can be expanded using Eq. 2.5 as:

$$\begin{aligned}\dot{\boldsymbol{\omega}}_{S/\mathcal{T}} \times \mathbf{r} &= -[\tilde{\mathbf{r}}] \left(-[\mathbf{I}_S]^{-1} [\tilde{\boldsymbol{\omega}}_{S/\mathcal{N}}] [\mathbf{I}_S] \boldsymbol{\omega}_{S/\mathcal{N}} + [\mathbf{I}_S]^{-1} \mathbf{L}_c - [\tilde{\boldsymbol{\omega}}_{S/\mathcal{N}}] \boldsymbol{\omega}_{S/\mathcal{T}} \right) \\ &= [\tilde{\mathbf{r}}] [\mathbf{I}_S]^{-1} [\tilde{\boldsymbol{\omega}}_{S/\mathcal{N}}] [\mathbf{I}_S] \boldsymbol{\omega}_{S/\mathcal{N}} - [\tilde{\mathbf{r}}] [\mathbf{I}_S]^{-1} [\tilde{\mathbf{r}}] \mathbf{F}_c + [\tilde{\mathbf{r}}] [\tilde{\boldsymbol{\omega}}_{S/\mathcal{N}}] \boldsymbol{\omega}_{S/\mathcal{T}}\end{aligned}\quad (2.51)$$

Note that if the target is not in uniform rotation, this equation must include an additional term to account for that rotation (see Eq. 2.5). Finally, the third term relates to the change of the contact point position \mathbf{r} . Since this vector is fixed in the \mathcal{S} -frame, one may write that:

$$\boldsymbol{\omega}_{S/\mathcal{T}} \times \dot{\mathbf{r}} = [\tilde{\boldsymbol{\omega}}_{S/\mathcal{T}}] [\tilde{\boldsymbol{\omega}}_{S/\mathcal{T}}] \mathbf{r} \quad (2.52)$$

Substitution of Eqs. 2.50, 2.51, and 2.52 into Eq. 2.49 yields for the contact point velocity change:

$$\begin{aligned}\dot{\mathbf{v}} &= \mathbf{g}(\mathbf{x}_S) - [\tilde{\boldsymbol{\Omega}}_{\mathcal{T}/\mathcal{N}}] [\tilde{\boldsymbol{\Omega}}_{\mathcal{T}/\mathcal{N}}] \mathbf{x}_S - 2[\tilde{\boldsymbol{\Omega}}_{\mathcal{T}/\mathcal{N}}] \mathbf{v}_S + \left(\frac{1}{m} \mathbf{I}_3 - [\tilde{\mathbf{r}}] [\mathbf{I}_S]^{-1} [\tilde{\mathbf{r}}] \right) \mathbf{F}_c \\ &\quad + [\tilde{\mathbf{r}}] [\mathbf{I}_S]^{-1} [\tilde{\boldsymbol{\omega}}_{S/\mathcal{N}}] [\mathbf{I}_S] \boldsymbol{\omega}_{S/\mathcal{N}} + [\tilde{\mathbf{r}}] [\tilde{\boldsymbol{\omega}}_{S/\mathcal{N}}] \boldsymbol{\omega}_{S/\mathcal{T}} + [\tilde{\boldsymbol{\omega}}_{S/\mathcal{T}}] [\tilde{\boldsymbol{\omega}}_{S/\mathcal{T}}] \mathbf{r}\end{aligned}\quad (2.53)$$

in which \mathbf{I} is the identity matrix. In this expression, the terms from Eqs. 2.50 and 2.51 that contain the contact force \mathbf{F}_c have been combined. Grouping some terms, this expression can be rewritten as:

$$\dot{\mathbf{v}} = \mathbf{g}(\mathbf{x}) + \mathbf{C} + \mathbf{W} + [\mathbf{M}] \mathbf{F}_c \quad (2.54)$$

in which:

$$\begin{aligned}\mathbf{C} &= -[\tilde{\boldsymbol{\Omega}}_{\mathcal{T}/\mathcal{N}}] [\tilde{\boldsymbol{\Omega}}_{\mathcal{T}/\mathcal{N}}] \mathbf{x} - 2[\tilde{\boldsymbol{\Omega}}_{\mathcal{T}/\mathcal{N}}] \mathbf{v} \\ \mathbf{W} &= [\tilde{\mathbf{r}}] [\mathbf{I}_S]^{-1} [\tilde{\boldsymbol{\omega}}_{S/\mathcal{N}}] [\mathbf{I}_S] \boldsymbol{\omega}_{S/\mathcal{N}} + [\tilde{\mathbf{r}}] [\tilde{\boldsymbol{\omega}}_{S/\mathcal{N}}] \boldsymbol{\omega}_{S/\mathcal{T}} + [\tilde{\boldsymbol{\omega}}_{S/\mathcal{T}}] [\tilde{\boldsymbol{\omega}}_{S/\mathcal{T}}] \mathbf{r} \\ [\mathbf{M}] &= \frac{1}{m} \mathbf{I}_3 - [\tilde{\mathbf{r}}] [\mathbf{I}_S]^{-1} [\tilde{\mathbf{r}}]\end{aligned}\quad (2.55)$$

Eq. 2.54, though simple in form, lies at the basis of our contact model development. The \mathbf{C} term combines accelerations resulting from the non-inertiality of the target frame. Similarly, the \mathbf{W} term accounts for the rigid body dynamics at play within the spacecraft frame. More important is the $[\mathbf{M}]$ matrix, which Stronge denotes as the “*inverse inertia matrix* [110].” This name is somewhat confusing, since $[\mathbf{M}]$ is *not* simply the inverse of the inertia matrix $[\mathbf{I}]$. Instead, it is related to the

inverse of the *total* inertia matrix that combines the mass m and inertia matrix $[\mathbf{I}]$ of the considered body.

The $[\mathbf{M}]$ matrix expresses how some force, applied at the contact point, changes the velocity of that contact point. It takes into account both the changes in the spacecraft linear velocity (from the force) and the spacecraft angular velocity (from the corresponding torque). This matrix is used by Stronge in his derivation of a collision model, though he only provides an element-wise definition and does not make mention of our more intuitive understanding of the physical meaning of $[\mathbf{M}]$ [110]. In fact, our expression for the inverse inertia matrix can be generalized, *i.e.*, one may define an $[\mathbf{M}_{ij}]$ that expresses how a force applied at some point \mathbf{r}_j on the spacecraft changes the velocity of some other point \mathbf{r}_i , as:

$$[\mathbf{M}_{ij}] = \frac{1}{m} \mathbb{I}_3 - [\tilde{\mathbf{r}}_i][\mathbf{I}_S]^{-1}[\tilde{\mathbf{r}}_j] \quad (2.56)$$

The inverse inertia matrix used in Eq. 2.54 is a ‘special case’ in which $\mathbf{r}_i = \mathbf{r}_j = \mathbf{r}$.

2.5.2 Collisions

We distinguish between two phases of contact: *contact motion* during which the spacecraft remains in continued contact with the target surface and ‘rolls’ across its surface, and impulsive *collisions* during which the spacecraft velocities are changed in an infinitesimally short impact on the surface and after which the spacecraft is returned to ballistic flight. The majority of the energy dissipation required for a spacecraft to come to rest occurs during collisions; they are therefore discussed first.

Our model includes a restitutive normal force and a tangential Coulomb friction force. As previously mentioned, a central collision occurs when \mathbf{r} and $\hat{\mathbf{N}}$ are parallel. Mathematically, this implies that the inverse inertia matrix $[\mathbf{M}]$ in Eq. 2.54 is diagonal when expressed in the contact frame, such that the restitution and friction forces are uncoupled. For three-dimensional collisions that are generally eccentric (or non-central), $[\mathbf{M}]$ is *not* diagonal, such that the two forces *are* coupled. In order to evaluate such a collision, recognize that the net contact force \mathbf{F}_c acting at \mathbf{H}

during a collision applies an impulse \mathbf{P} on the spacecraft, for which it can be written that:

$$d\mathbf{P} = \mathbf{F}_c \cdot dt \quad (2.57)$$

Assuming that the collision takes place in an infinitesimally short amount of time, the spacecraft position and attitude remain fixed over the course of the collision. Furthermore, recognize that the normal component $p = \mathbf{P}^T \hat{\mathbf{N}}$ of the applied contact force must monotonically increase during a collision, because the normal force \mathbf{F}_N is compressive. In other words, the normal force can only ‘push’ the spacecraft away from the surface, it can never ‘pull’ it onto the surface. Because of this monotonic increase, the applied normal impulse p can be used as an independent variable with respect to which the velocity changes during collision can be expressed. Applying these simplifications to the spacecraft equations of motion, one can write for its velocity change during collision that:

$$\begin{aligned} \frac{d\mathbf{v}}{dp} &= \frac{1}{m} \frac{d\mathbf{P}}{dp} \\ \frac{d\boldsymbol{\omega}_{S/\mathcal{T}}}{dp} &= [\mathbf{I}_S]^{-1} [\hat{\mathbf{r}}] \frac{d\mathbf{P}}{dp} \end{aligned} \quad (2.58)$$

Correspondingly, the contact point velocity change is reduced from Eq. 2.54 to:

$$\frac{d\mathbf{v}}{dp} = [\mathbf{M}] \frac{d\mathbf{P}}{dp} \quad (2.59)$$

During a collision, two deformable objects are driven together under an initially negative relative normal velocity. The objects will deform, reducing the relative normal velocity and storing the corresponding energy as internal mechanical energy. After this ‘compression’ has stopped, part of the stored energy is returned to the bodies during ‘restitution,’ with the collision eventually ending with the two bodies having some positive relative normal velocity.

Each collision is evaluated in analogy with the illustration above, and consists of two phases: a *compressive* phase during which the normal contact point velocity increases from some negative value to zero, followed by a *restitutive* phase during which it is returned to some positive value. In order to determine when the collision ends, the variable ξ is tracked, defined as the integral of the $v_3(p)$ curve. In other words, ξ represents the work done on the spacecraft by the normal impulse.

Over the course of the compression phase, the normal force performs a total $\xi_c < 0$ of work. The collision is then defined to end once the restitution phase has performed a total $\xi_r \geq 0$ of work, with ξ_r defined relative to ξ_c as [110]:

$$e^2 = -\frac{\xi_r}{\xi_c} \iff \xi_r = -e^2 \xi_c \quad (2.60)$$

in which $0 \leq e \leq 1$ is the *coefficient of restitution*. Note that Eq. 2.60 provides an *energetic* definition of the coefficient, as opposed to the classical *kinetic* definition of $e = -v_3^+/v_3^-$, in which v_3^+ and v_3^- are the normal contact point velocities at the end and start of the collision, respectively. It is important to note that the two definitions (kinetic and energetic e) are equivalent for central collisions. This is *not* true for eccentric collisions. In fact, the use of a kinetically-defined coefficient of restitution may lead to energy increases in the resolution of eccentric collisions. Again, it is stressed that for arbitrary eccentric collisions, it holds that:

$$-\frac{\xi_r}{\xi_c} \neq \left(\frac{v_3^+}{v_3^-} \right)^2 \quad (2.61)$$

The energetic definition of the coefficient of restitution, as introduced by Stronge [110], is thought by some to be “*the most consistent and applicable*” [3]. Nevertheless, note the lack of experimental measurements of the energetic coefficient of restitution of complex shapes in low-gravity environments. This is an area of basic research that should be performed as part of the validation of the applied impact model, which is a piece of suggested future work that will be discussed further on.

Slip Eq. 2.59 relates the change in contact point velocity to the total impulse applied at the contact point. As mentioned, all variables are expressed in the contact frame \mathcal{C} . In other words, the contact point velocity \mathbf{v} has the following components:

$$\mathbf{v} = \begin{bmatrix} v_1 & v_2 & v_3 \end{bmatrix}^T \quad (2.62)$$

in which $v_3 = \mathbf{v} \cdot \hat{\mathbf{N}}$ is the *normal* contact point velocity, and in which $v_1 = \mathbf{v} \cdot \hat{\mathbf{u}}_1$ and $v_2 = \mathbf{v} \cdot \hat{\mathbf{u}}_2$ are the two components of the *tangential* contact point velocity. The latter two can also be expressed

using the *sliding velocity* s and *direction* ϕ relative to the $\hat{\mathbf{u}}_1$ axis, as:

$$\begin{cases} s^2 = v_1^2 + v_2^2 \\ \tan \phi = \frac{v_2}{v_1} \end{cases} \quad (2.63)$$

Equivalently, it may be written that:

$$\begin{cases} v_1 = s \cos \phi \\ v_2 = s \sin \phi \end{cases} \quad (2.64)$$

As mentioned, our contact model includes both a restitution and Coulomb friction impulse. The normal impulse acts to reduce v_3 from some negative value at the start of a collision, to some positive v_3 at the end of the collision. The Coulomb friction impulse acts against v_1 and v_2 in an attempt to reduce $s \rightarrow 0$ and to maintain $s = 0$ once it is reached. When $s > 0$, the Coulomb friction impulse acts against the tangential contact point velocity with magnitude $d\mathbf{P}_F = f \cdot d\mathbf{P}_N$, in which f is the *coefficient of friction* and $d\mathbf{P}_N$ is the magnitude of the normal impulse. Using this, the combined impulse from restitution and friction can be written as [110]:

$$\left(\frac{d\mathbf{P}}{dp} \right)_{slip} = \begin{bmatrix} -f \cos \phi \\ -f \sin \phi \\ 1 \end{bmatrix} \quad (2.65)$$

whenever $s > 0$. When this is true, the contact point is said to be in *slip*. By substituting Eq. 2.65 into Eqs. 2.58 and 2.59, one can integrate \mathbf{v} , \mathbf{v} , and $\boldsymbol{\omega}_{S/\mathcal{T}}$ as the restitution and friction impulses are applied.

Stick Once the sliding velocity s reaches zero, friction is applied at some reduced magnitude $d\mathbf{P}_F = f^* \cdot d\mathbf{P}_N$ with $f^* \leq f$ and in direction ϕ^* , in order to maintain $s = 0$. In other words, friction is applied at some critical f^* and ϕ^* such that $\dot{s} = 0$. The latter implies that $\dot{v}_1 = \dot{v}_2 = 0$

during stick. To compute the corresponding f^* and ϕ^* , start from Eq. 2.59:

$$\begin{aligned} \left(\frac{d\mathbf{v}}{dp}\right)_{stick} &= \mathbf{M} \begin{bmatrix} -f^* \cos \phi^* \\ -f^* \sin \phi^* \\ 1 \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \\ \dot{v}_3 \end{bmatrix} \\ \Leftrightarrow \begin{bmatrix} -f^* \cos \phi^* \\ -f^* \sin \phi^* \\ \dot{v}_3 \end{bmatrix} &= - \begin{bmatrix} M_{11} & M_{12} & 0 \\ M_{21} & M_{22} & 0 \\ M_{31} & M_{32} & 1 \end{bmatrix}^{-1} \begin{bmatrix} M_{13} \\ M_{23} \\ M_{33} \end{bmatrix} \end{aligned} \quad (2.66)$$

Writing out the inverse, it holds that:

$$\begin{bmatrix} -f^* \cos \phi^* \\ -f^* \sin \phi^* \\ \dot{v}_3 \end{bmatrix} = M^* \begin{bmatrix} M_{22} & -M_{12} & 0 \\ -M_{21} & M_{11} & 0 \\ M_{21}M_{32} - M_{22}M_{31} & M_{12}M_{31} - M_{11}M_{32} & 1 \end{bmatrix} \begin{bmatrix} M_{13} \\ M_{23} \\ M_{33} \end{bmatrix} \quad (2.67)$$

in which:

$$M^* = \frac{-1}{M_{11}M_{22} - M_{12}M_{21}} \quad (2.68)$$

Dividing the second component of Eq. 2.67 by the first component, the critical ϕ^* is found as:

$$\begin{aligned} \tan \phi^* &= \frac{\sin \phi^*}{\cos \phi^*} = \frac{M_{11}M_{23} - M_{21}M_{13}}{M_{22}M_{13} - M_{12}M_{23}} \\ \Leftrightarrow \phi^* &= \arctan \left(\frac{M_{11}M_{23} - M_{21}M_{13}}{M_{22}M_{13} - M_{12}M_{23}} \right) \end{aligned} \quad (2.69)$$

Similarly, by summing the squares of the first and second components of Eq. 2.67, the critical coefficient of friction f^* is found as:

$$\begin{aligned} (f^*)^2 &= \left(\frac{M_{12}M_{23} - M_{22}M_{13}}{M_{11}M_{22} - M_{12}M_{21}} \right)^2 + \left(\frac{M_{21}M_{13} - M_{11}M_{23}}{M_{11}M_{22} - M_{12}M_{21}} \right)^2 \\ \Leftrightarrow f^* &= \sqrt{\frac{(M_{12}M_{23} - M_{22}M_{13})^2 + (M_{21}M_{13} - M_{11}M_{23})^2}{(M_{11}M_{22} - M_{12}M_{21})^2}} \end{aligned} \quad (2.70)$$

In summary, the friction impulse during stick is expressed as [110]:

$$\left(\frac{d\mathbf{P}}{dp}\right)_{stick} = \begin{bmatrix} -f^* \cos \phi^* \\ -f^* \sin \phi^* \\ 1 \end{bmatrix} \quad (2.71)$$

with the critical friction coefficient and direction given by Eqs. 2.69 and 2.70.

Slip reversal If the coefficient of friction is sufficiently high, friction will drive $s \rightarrow 0$ during a collision. Once $s = 0$ is achieved, two situations are possible, depending on the value of the critical coefficient of friction f^* . In the large majority of cases, it holds that $f^* \leq f$, such that friction is sufficiently strong to maintain stick at $s = 0$. If instead $f^* < f$, friction is *not* sufficiently strong to maintain $s = 0$ and slip resumes in a final phase of *slip reversal*.

It was shown by [110] that the slip direction $\hat{\phi}$ during slip reversal is constant. The direction of slip reversal $\hat{\phi}$ can be found by solving for $d\phi = 0$ with $ds > 0$ [110]. In order to do so, expressions for ds and $d\phi$ must first be obtained. From Eq. 2.63, one may write that:

$$\begin{aligned} s^2 &= v_1^2 + v_2^2 \\ \iff 2s \frac{ds}{dp} &= 2v_1 \frac{dv_1}{dp} + 2v_2 \frac{dv_2}{dp} \\ \iff \frac{ds}{dp} &= \cos \phi \frac{dv_1}{dp} + \sin \phi \frac{dv_2}{dp} \end{aligned} \quad (2.72)$$

and:

$$\begin{aligned} \tan \phi &= \frac{v_2}{v_1} \\ \iff \frac{1}{\cos^2 \phi} \frac{d\phi}{dp} &= \frac{1}{v_1} \frac{dv_2}{dp} - \frac{v_2}{v_1} \frac{dv_1}{dp} \\ \iff \frac{d\phi}{dp} &= \frac{\cos \phi}{s} \frac{dv_2}{dp} - \frac{\sin \phi}{s} \frac{dv_1}{dp} \end{aligned} \quad (2.73)$$

We wish to solve for the $\hat{\phi}$ that satisfies $d\phi = 0$. However, notice the appearance of s in the denominator of Eq. 2.73. Since this expression is being evaluated at the instant when stick is reached, *i.e.*, when $s = 0$, Eq. 2.73 becomes singular. To avoid this, we multiply both sides by s and solve for $sd\phi = 0$, *i.e.*:

$$\begin{aligned} h = sd\phi &= \cos(\hat{\phi})dv_2 - \sin(\hat{\phi})dv_1 = 0 \\ \iff \frac{1}{2}f(M_{11} - M_{22}) \sin(2\hat{\phi}) - fM_{12} \cos(2\hat{\phi}) + M_{23} \cos(\hat{\phi}) - M_{13} \sin(\hat{\phi}) &= 0 \end{aligned} \quad (2.74)$$

in which the root $\hat{\phi}$ must satisfy $ds > 0$ for slip to resume. Thus, it must satisfy that:

$$\begin{aligned} g = ds &= \cos \phi dv_1 + \sin \phi dv_2 > 0 \\ \iff M_{13} \cos \phi + M_{23} \sin \phi - fM_{11} \cos^2 \phi - 2fM_{12} \cos \phi \sin \phi - fM_{22} \sin^2 \phi > 0 \end{aligned} \quad (2.75)$$

It can be shown that Eq. 2.74 has three roots, of which only one satisfies Eq. 2.75 [110]. Classical iterative root-finding methods work well when using an initial guess of $\hat{\phi}_0 = \phi^*$; a third-order Householder method is applied to accelerate the convergence rate on the proper root. In this, $\hat{\phi}$ is iteratively solved for as:

$$\hat{\phi}_{i+1} = \hat{\phi}_i - \frac{6hh''^2 - 3h^2h''}{6h'^3 - 6hh'h'' + h^2h'''} \quad (2.76)$$

in which the h^i partials can be found from Eq. 2.74 as:

$$\begin{aligned} h' &= \frac{\partial h}{\partial \phi} = -M_{13} \cos \phi - M_{23} \sin \phi + f(M_{11} - M_{22}) \cos(2\phi) + 2fM_{12} \sin(2\phi) \\ h'' &= \frac{\partial^2 h}{\partial \phi^2} = M_{13} \sin \phi - M_{23} \cos \phi - 2f(M_{11} - M_{22}) \sin(2\phi) + 4fM_{12} \cos(2\phi) \\ h''' &= \frac{\partial^3 h}{\partial \phi^3} = M_{13} \cos \phi - M_{23} \sin \phi - 4f(M_{11} - M_{22}) \cos(2\phi) - 8fM_{12} \sin(2\phi) \end{aligned} \quad (2.77)$$

In practice, three steps of this iteration are found to be sufficient to achieve convergence. In summary, if slip reversal is detected, the spacecraft state can be integrated for the remainder of the collision with:

$$\left(\frac{d\mathbf{P}}{dp} \right)_{rev} = \begin{bmatrix} -f \cos \hat{\phi} \\ -f \sin \hat{\phi} \\ 1 \end{bmatrix} \quad (2.78)$$

in which the direction of slip reversal $\hat{\phi}$ can be iteratively found using Eq. 2.76.

To understand what physically happens during a collision with slip reversal, consider the three sample velocity hodographs shown in Fig. 2.54. In all three subplots, the horizontal axes express the two tangential contact point velocities, v_1 and v_2 . All collisions begin with some initial sliding velocity $s_i > 0$ and end with some final sliding velocity s_f . In these hodographs, the friction impulse always points at the origin, since friction aims to reduce the sliding velocity to zero. The normal impulse points in some other, fixed direction, in this case towards the positive v_1 axis. If the friction coefficient f is sufficiently large (as shown on the left side of the figure), friction dominates the collision and draws the sliding velocity $s \rightarrow 0$; it achieves stick. If f is small (as shown in the center of the figure), the normal impulse dominates the collision; the contact point is unable to achieve stick, but instead continues sliding throughout the entire collision. In the rare case where

the normal impulse happens to point at the origin of the hodograph (as shown on the right side of the figure), both of the impulses draw the sliding velocity to the $s = 0$ point (the origin). However, assuming that f is small enough, the normal impulse will ‘overpower’ the friction impulse and simply continue pushing the s curve, away from the $s = 0$ point. This latter example is precisely how slip reversal occurs, though the hodographs have been simplified for visualization purposes.




Figure 2.54: Illustration of velocity hodographs for different slip/stick behaviors.

Numerical implementation In light of the collision dynamics detailed above, one may be inclined to integrate the following *collision state* \mathbf{X}_c :

$$\mathbf{X}_c = \begin{bmatrix} \mathbf{v} \\ \mathbf{P} \\ \mathbf{v} \\ \omega_{S/\mathcal{T}} \\ \xi \end{bmatrix} \quad \text{with} \quad \frac{d\mathbf{X}_c}{dp} = \begin{bmatrix} [\mathbf{M}] \frac{d\mathbf{P}}{dp} \\ \text{Eq. 2.65 or Eq. 2.71 or Eq. 2.78} \\ \frac{1}{m} \frac{d\mathbf{P}}{dp} \\ [\mathbf{I}_S]^{-1} [\tilde{\mathbf{r}}] \frac{d\mathbf{P}}{dp} \\ v_3 \end{bmatrix} \quad (2.79)$$

in which ξ is used to determine when the collision ends. However, a naïve implementation of these equations will be numerically unstable. More specifically, the collision integration will begin in slip, and the onset of stick must be converged on using Eq. 2.65 until $s = 0$, after which Eq. 2.71 must be used. As both v_1 and v_2 approach zero as $s \rightarrow 0$, the direction of slip ϕ will become ill-defined, resulting in oscillations that make it difficult to properly converge on $s = 0$. To avoid this problem, a change of variable is first performed, by replacing v_1 and v_2 with the equivalent s and ϕ in the

collision state, such that instead:

$$\mathbf{X}_c = \begin{bmatrix} s & \phi & v_3 & \mathbf{P} & \mathbf{v} & \boldsymbol{\omega}_{S/\mathcal{T}} & \xi \end{bmatrix}^T \quad (2.80)$$

Note that we do continue to use the normal contact point velocity v_3 . The differential equations expressing the change in s and ϕ were given in Eqs. 2.72 and 2.73. By itself, this change of variables does not alleviate the challenge of converging on $s = 0$. For this, the approach by Bhatt and Koechling is followed, by introducing the *stretching variable* τ , which is defined as [10]:

$$d\tau = \frac{f}{s} dp \quad \iff \quad \frac{dp}{d\tau} = \frac{s}{f} \quad (2.81)$$

The coefficient of friction f is included in this expression in order to pseudo-normalize the ‘velocity flow,’ as previously illustrated in Fig. 2.54, such that impacts with different f values experience similar ranges of τ . Bhatt and Koechling use this definition in order to perform detailed studies of the velocity flow that occurs during eccentric rigid-body collisions, but they note that τ can also be used “to guide the numerical integration as the point of sticking is reached, making the integration of the equations defining the impact process more robust” [10]. This is precisely how the stretching variable is applied. During slip, the collision equations of motion are integrated with respect to τ rather than with respect to p . The reason for doing so is clear from Eq. 2.81: when s approaches zero, $dp/d\tau$ also approaches zero. As a result, even though $s = 0$ is achieved at some finite p that is difficult to numerically converge on, it is reached at infinite τ . The use of τ as integration variable therefore avoids having to actively converge on $s = 0$. Instead, we continue to integrate along τ towards infinity, until the sliding velocity reaches below some small tolerance $s < \epsilon$ at some finite τ . When this occurs, we switch to integrating sticking motion using p instead, since $dp/d\tau = 0$ during stick. Given the equations of motion for the contact state with respect to dp , the equivalent expression with respect to $d\tau$ are easily found using the chain rule. For any variable k , it holds that:

$$\frac{dk}{d\tau} = \frac{dk}{dp} \cdot \frac{dp}{d\tau} = \frac{dk}{dp} \cdot \frac{s}{f} \quad (2.82)$$

In order to minimize the amount of switch statements² required for a numerical implementation

² The reason for preferring this minimization is discussed in Section 2.7

of collision evaluation that distinguishes between slip and stick, the variable η is therefore used, which is defined as:

$$\eta = \begin{cases} s/f & \text{during slip} \Rightarrow \text{integration wrt. } \tau \\ 1 & \text{during stick and slip reversal} \Rightarrow \text{integration wrt. } p \end{cases} \quad (2.83)$$

The collision state is then propagated with respect to the integration variable σ as:

$$\dot{\mathbf{X}}_c = \eta \cdot \begin{bmatrix} \frac{ds}{d\sigma} \\ \frac{d\phi}{d\sigma} \\ \frac{dv_3}{d\sigma} \\ \frac{d\mathbf{P}}{d\sigma} \\ \frac{d\mathbf{v}}{d\sigma} \\ \frac{d\boldsymbol{\omega}}{d\sigma} \\ \frac{d\xi}{d\sigma} \end{bmatrix} = \eta \cdot \begin{bmatrix} \cos \phi \frac{dv_1}{dp} + \sin \phi \frac{dv_2}{dp} \\ \frac{\cos \phi}{s} \frac{dv_2}{dp} - \frac{\sin \phi}{s} \frac{dv_1}{dp} \\ \frac{dv_3}{dp} \\ \text{Eq. 2.65 or Eq. 2.71 or Eq. 2.78} \\ \frac{1}{m} \frac{d\mathbf{P}}{dp} \\ [\mathbf{I}_S]^{-1} [\tilde{\mathbf{r}}] \frac{d\mathbf{P}}{dp} \\ v_3 \end{bmatrix} \quad \text{with} \quad \frac{d\mathbf{v}}{dp} = [\mathbf{M}] \frac{d\mathbf{P}}{dp} \quad (2.84)$$

The integration variable σ takes on two variants: during stick, it holds that $\sigma = p$. During slip and slip reversal, it holds that $\sigma = \tau$. Using this implementation, the only type of ‘switch’ statement to be implemented in a numerical propagator are a simple `if/else` switch on η , using Eq. 2.83), and the three cases for $d\mathbf{P}/dp$ (as given by Eqs. 2.65, 2.71, and 2.78). On a final note, remark that the f^* , ϕ^* , and $\hat{\phi}$ values are invariant over the course of any given collision. They must therefore only be computed once, at the start of each collision.

Rolling resistance The contact model presented above includes both restitution and Coulomb friction. This model works well for shapes that naturally resist rolling motion, such as a cube or tetrahedron. The resistance or ability to roll can be empirically related to the angular defect at the vertices of a considered shape. This will be shown in Section 4.2 when the motion of different shapes is compared. When a shape *does* have the ability to roll, such as a sphere or disk, the contact model above is deficient. This is because both restitution and Coulomb friction cannot dissipate any more energy once a shape has settled into ‘rolling’ on a flat surface. Intuitively, we know that a ball rolling on a flat surface does eventually come to a halt; the phenomenon responsible for

the corresponding energy loss is known in literature as rolling friction or rolling resistance during continued contact motion [130, 26]. Although the implementation of rolling resistance into our simulation framework has been verified, it is not nominally it during *collisions*. Nevertheless, the implementation is detailed below.

In order to account for rolling resistance during collisions, the collision model shown above is augmented by including a torque $d\mathbf{L}_{RR,roll}$ that opposes the angular velocity of the spacecraft:

$$\frac{d\mathbf{L}_{RR,roll}}{dp} = -[\mathbf{C}_{RR}]R \frac{\boldsymbol{\omega}_{S/\mathcal{T}}}{\omega_{S/\mathcal{T}}} \quad (2.85)$$

in which $[\mathbf{C}_{RR}]$ is a matrix containing the *coefficients of rolling resistance*. All of [118, 130, 26] consider the rolling resistance of a sphere, which has the same resistance to rolling about all three of its axes. In that case, $[\mathbf{C}_{RR}] = C_{RR} \cdot \mathbf{I}_3$ with C_{RR} analogous to the coefficient of friction. However, certain shapes are able to roll in one direction while resisting rolling in other directions, such as a cylinder. In this case, rolling resistance should only be applied along the axis of rolling, such that the diagonal³ entries of $[\mathbf{C}_{RR}]$ are no longer identical. For example, a cylinder with the z axis as axis of revolution could have $[\mathbf{C}_{RR}] = \text{diag}(0, 0, 1)$. Indeed, $C_{RR} = 0$ in the directions in which natural rolling behavior does not occur. Finally, note that in Eq. 2.85, R is the radius of the spacecraft in the direction of rolling.

In analogy to the Coulomb friction force, rolling resistance acts with some fixed magnitude as it drives $\omega_{S/\mathcal{T}} \rightarrow 0$. Once $\omega_{S/\mathcal{T}} = 0$, it acts with some reduced magnitude to keep this angular velocity at zero. The torque $d\mathbf{L}_{RR,lock}$ required to maintain $\omega_{S/\mathcal{T}} = 0$, a state which we denote as *lock*, can be computed from Eq. 2.58 as:

$$\begin{aligned} \frac{d\boldsymbol{\omega}_{S/\mathcal{T}}}{dp} &= [\mathbf{I}_S]^{-1}[\tilde{\mathbf{r}}] \frac{d\mathbf{P}}{dp} + [\mathbf{I}_S]^{-1} \frac{d\mathbf{L}_{RR,lock}}{dp} = \mathbf{0} \\ \iff \frac{d\mathbf{L}_{RR,lock}}{dp} &= -[\tilde{\mathbf{r}}] \frac{d\mathbf{P}}{dp} \end{aligned} \quad (2.86)$$

Although another stretching variable could be used to handle the $\omega_{S/\mathcal{T}} \rightarrow 0$ approach, the simultaneous use of two stretching variables is not trivial. Instead, we make use of the regularization

³ The directions of rolling are assumed to be aligned with the axes of the spacecraft frame, such that $[\mathbf{C}_{RR}]$ is always diagonal when expressed in the spacecraft frame.

variable λ_{RR} . This variable is defined as:

$$\lambda_{RR} = \frac{\min(\omega_{S/\mathcal{T}}, \omega^*)}{\omega^*} \quad (2.87)$$

It is equal to 1 when $\omega_{S/\mathcal{T}} \geq \omega^*$ and linearly reduces to zero as $\omega_{S/\mathcal{T}} \rightarrow 0$ while $\omega_{S/\mathcal{T}} < \omega^*$. Using this variable, rolling resistance is regularized as:

$$\frac{d\mathbf{L}_{RR}}{dp} = \lambda_{RR} \frac{d\mathbf{L}_{RR,roll}}{dp} + (1 - \lambda_{RR}) \frac{d\mathbf{L}_{RR,lock}}{dp} \quad (2.88)$$

which avoids numerical oscillations as the angular velocity approaches zero. When the torque of Eq. 2.88 is applied by itself, the contact point velocity \mathbf{v} is changed. This introduces a coupling between rolling resistance and respectively restitution and Coulomb friction, such that the respective coefficients can no longer be independently defined. This issue is highlighted in [130] and [118]. We follow their approach and also include a rolling resistance *force*, applied at the spacecraft center of mass, such that the net effect on the contact point velocity is nullified. To obtain an expression for this force, we start from the equation for the contact point velocity change resulting from some force and torque:

$$\frac{d\mathbf{v}_{RR}}{dp} = \frac{1}{m} \frac{d\mathbf{P}_{RR}}{dp} - [\tilde{\mathbf{r}}][\mathbf{I}_S]^{-1} \frac{d\mathbf{L}_{RR}}{dp} \quad (2.89)$$

Equating this to zero and solving for the corresponding linear impulse:

$$\frac{d\mathbf{v}_{RR}}{dp} = \mathbf{0} \iff \frac{d\mathbf{P}_{RR}}{dp} = m[\tilde{\mathbf{r}}][\mathbf{I}_S]^{-1} \frac{d\mathbf{L}_{RR}}{dp} \quad (2.90)$$

In summary, the inclusion of rolling resistance updates the collision equations of motion to:

$$\dot{\mathbf{X}}_c = \eta \cdot \begin{bmatrix} \cos \phi \frac{dv_1}{dp} + \sin \phi \frac{dv_2}{dp} \\ \frac{\cos \phi}{s} \frac{dv_2}{dp} - \frac{\sin \phi}{s} \frac{dv_1}{dp} \\ \frac{dv_3}{dp} \\ \text{Eq. 2.65 or Eq. 2.71 or Eq. 2.78} \\ \frac{1}{m} \frac{d\mathbf{P}}{dp} \\ [\mathbf{I}_S]^{-1} [\tilde{\mathbf{r}}] \frac{d\mathbf{P}}{dp} \\ v_3 \end{bmatrix}_{N+F} + \eta \cdot \begin{bmatrix} 0 \\ 0 \\ 0 \\ \frac{d\mathbf{P}_{RR}}{dp} \\ \frac{1}{m} \frac{d\mathbf{P}_{RR}}{dp} \\ [\mathbf{I}_S]^{-1} \frac{d\mathbf{L}_{RR}}{dp} \\ 0 \end{bmatrix}_{RR} \quad (2.91)$$

2.5.3 Contact motion

The motion of a rectangular payload deployed to the surface of some body is illustrated in Fig. 2.55. The probe is released above the surface at some t_0 , ballistically descends towards the surface, and first impacts it at t_1 . A number of successive impacts occur, some of which are *single* bounces and some of which are *multiple* bounces in which the probe performs touches the surface several times before ‘flying’ up into another ballistic arc.




Figure 2.55: Illustration of the different phases in the rover deployment to a small body.

During each collision, the energy of the spacecraft is reduced, until the normal contact point velocity reaches below some small threshold value at the start of a collision. When this is detected, two situations are possible:

- (1) If the *external* normal acceleration of the contact point is negative, this implies that the surface must continuously apply a normal force in order to enforce the non-penetration constraint. One final collision with $e = 0$ is then performed, such that the normal contact point velocity $v_3 = 0$ at the end of the compression phase. *Contact motion* then begins, in which the spacecraft is in continuous contact with the small-body surface.
- (2) If the external normal acceleration is positive, this implies that the accelerations on the spacecraft are such that the contact point will move upwards from the surface, of its own accord. Contact motion *does not* start, and the spacecraft is returned to another ballistic arc instead. This typically only happens when the coefficient of restitution is very low, at roughly $e < 0.05$.

The contact motion phase is different from a collision in that it is not evaluated impulsively. In other words, time continues to run during contact motion, while it was paused during the evaluation of a collision. As an example of contact motion, consider the ‘rolling’ of the rectangle on one of its vertices on the right side of Fig. 2.55. Although contact motion is a negligible phase for objects that oppose rolling (such as a cube or tetrahedron), it can be much more significant for objects that do roll (such as a sphere or disk). This will be shown through simulation in Section 4.2.

During contact motion, the spacecraft is subject to the same surface interactions as during collision phases. First, a normal force is applied in order to prevent the contact point from penetrating through the small-body surface. Second, a classical Coulomb friction force resists the tangential contact point velocity. Finally, rolling resistance acts against the spacecraft angular velocity. Starting from the expression for the contact point velocity change, as given by Eq. 2.54, the equations of motion for contact motion are now derived. Similar to the approach for collisions, expressions for the normal and friction forces during respectively stick and slip are first obtained. Rolling resistance is then included into the model. As mentioned, the normal force enforces the non-penetration constraint during contact motion, *i.e.*, it enforces:

$$\dot{v}_N = \frac{d}{dt} (\mathbf{v} \cdot \hat{\mathbf{N}}) = \frac{d\mathbf{v}}{dt} \cdot \hat{\mathbf{N}} + \mathbf{v} \cdot \frac{d\hat{\mathbf{N}}}{dt} = 0 \quad (2.92)$$

This expression holds regardless of whether the contact point is in stick or slip.

Stick If the contact point is in stick, it holds that also its tangential velocity $\dot{v}_T = v_T = 0$. Combining this with the non-penetration constraint, it holds that during sticking contact motion, $\dot{\mathbf{v}} = \mathbf{0}$. From Eq. 2.54, the corresponding force to be applied at the contact point can therefore be obtained:

$$\begin{aligned} \dot{\mathbf{v}} &= \mathbf{g} + \mathbf{C} + \mathbf{W} + [\mathbf{M}]\mathbf{F}_{c,stick} = \mathbf{0} \\ \iff \mathbf{F}_{c,stick} &= -[\mathbf{M}]^{-1} (\mathbf{g} + \mathbf{C} + \mathbf{W}) \end{aligned} \quad (2.93)$$

This force can be broken up into its normal and tangential components with respect to $\hat{\mathbf{N}}$ to yield

the normal and friction forces, as:

$$\begin{cases} \mathbf{F}_{N,stick} = \left(\mathbf{F}_{c,stick}^T \hat{\mathbf{N}} \right) \hat{\mathbf{N}} = \hat{\mathbf{N}} \hat{\mathbf{N}}^T \mathbf{F}_{c,stick} \\ \mathbf{F}_{F,stick} = \mathbf{F}_{c,stick} - \mathbf{F}_{N,stick} = \left(\mathbb{I}_3 - \hat{\mathbf{N}} \hat{\mathbf{N}}^T \right) \mathbf{F}_{c,stick} = [\mathbf{T}] \mathbf{F}_{c,stick} \end{cases} \quad (2.94)$$

in which the tangential projection matrix $[\mathbf{T}]$ is defined as:

$$[\mathbf{T}] = \mathbb{I}_3 - \hat{\mathbf{N}} \hat{\mathbf{N}}^T \quad (2.95)$$

During sticking contact motion, both $\hat{\mathbf{N}}$ and $[\mathbf{T}]$ remain fixed.

Slip Although the *normal* contact point velocity remains zero during sliding contact motion, its *tangential* velocity is allowed to change. In order to derive expressions for the corresponding normal and friction forces, the normal and tangential contact point velocities are first expressed (vectorially) as:

$$\mathbf{v}_N = \hat{\mathbf{N}} \hat{\mathbf{N}}^T \mathbf{v} \quad \text{and} \quad \mathbf{v}_T = \left(\mathbb{I}_3 - \hat{\mathbf{N}} \hat{\mathbf{N}}^T \right) \mathbf{v} = [\mathbf{T}] \mathbf{v} \quad (2.96)$$

In contrast to sticking contact, the surface normal $\hat{\mathbf{N}}$ can change during sliding contact. As a result, both the normal and tangential directions change, which affects the magnitude of the contact forces. Our methodology of evaluating the change in surface normal will be addressed shortly.

When the sliding velocity v_T is non-zero, the friction force acts with magnitude $F_F = f \cdot F_N$ opposite the direction of slip. It can thus be written as:

$$\mathbf{F}_{F,slip} = -f F_{N,slip} \hat{\mathbf{v}}_T \quad (2.97)$$

To obtain the magnitude of the normal force during sliding contact, Eq. 2.54 is substituted into Eq. 2.92. This yields for the change in normal contact point velocity that:

$$\dot{v}_N = \left(\mathbf{g} + \mathbf{C} + \mathbf{W} + [\mathbf{M}] F_{N,slip} \hat{\mathbf{N}} + [\mathbf{M}] f F_{N,slip} \hat{\mathbf{v}}_T \right) \cdot \hat{\mathbf{N}} + \mathbf{v} \cdot \frac{d\hat{\mathbf{N}}}{dt} = 0 \quad (2.98)$$

Rewriting, this yields the following normal force magnitude:

$$F_{N,slip} = - \left[\left([\mathbf{M}] \hat{\mathbf{N}} + f [\mathbf{M}] \hat{\mathbf{v}}_T \right) \cdot \hat{\mathbf{N}} \right]^{-1} \left(\mathbf{v} \cdot \frac{d\hat{\mathbf{N}}}{dt} + (\mathbf{g} + \mathbf{C} + \mathbf{W}) \cdot \hat{\mathbf{N}} \right) \quad (2.99)$$

which yields a (vectorialized) normal force of $\mathbf{F}_{N,slip} = F_{N,slip} \hat{\mathbf{N}}$.

Changes to the normal Eq. 2.99 for the normal force during sliding contact contains the term $d\hat{\mathbf{N}}/dt$. Indeed, as mentioned, the surface normal $\hat{\mathbf{N}}$ changes during sliding contact. Although the SDF voxel size is typically such that changes to $\hat{\mathbf{N}}$ within any given voxel are small, they cannot be ignored when computing contact forces. The reader is referred back to Fig. 2.12 for an illustration of how the surface normal can change within a given voxel.

In order to derive an expression for the rate of change of the surface normal, it is repeated that the surface normal is obtained as the gradient of the signed distance field. However, as mentioned in Section 2.2, it is necessary to normalize this gradient, as the linear interpolant does not preserve the distance function property that $\|\nabla d(\mathbf{x})\| = 1$. The surface normal is thus computed as:

$$\hat{\mathbf{N}} = \frac{\nabla d(\mathbf{x})}{\|\nabla d(\mathbf{x})\|} = \frac{\mathbf{N}}{N} \quad (2.100)$$

in which $\nabla d(\mathbf{x})$ is obtained using Eq. 2.20. The rate of change of this expression is found using the chain rule as:

$$\frac{d\hat{\mathbf{N}}}{dt} = \frac{d\mathbf{N}}{dt} \frac{1}{N} - \frac{\mathbf{N}}{N^2} \frac{dN}{dt} \quad (2.101)$$

which consists of two terms. Applying the chain rule again, the first term of Eq. 2.101 is found as:

$$\frac{d\mathbf{N}}{dt} = \frac{d\mathbf{N}}{d\mathbf{x}} \frac{d\mathbf{x}}{dt} = \nabla \nabla d(\mathbf{x}) \mathbf{v} \quad (2.102)$$

To obtain the second term of Eq. 2.101, start from the definition of N^2 :

$$\begin{aligned} N^2 &= \mathbf{N} \cdot \mathbf{N} \\ \iff 2N \frac{dN}{dt} &= 2 \frac{d\mathbf{N}}{dt} \cdot \mathbf{N} \\ \iff \frac{dN}{dt} &= \hat{\mathbf{N}} \cdot \nabla \nabla d(\mathbf{x}) \mathbf{v} \end{aligned} \quad (2.103)$$

Substituting Eqs. 2.102 and 2.103 into Eq. 2.101, the rate of change of the surface normal is found as:

$$\frac{d\hat{\mathbf{N}}}{dt} = \frac{1}{N} \left(\mathbb{I}_3 - \hat{\mathbf{N}} \hat{\mathbf{N}}^T \right) \nabla \nabla d(\mathbf{x}) \mathbf{v} = \frac{1}{N} [\mathbf{T}] \nabla \nabla d(\mathbf{x}) \mathbf{v} \quad (2.104)$$

This can be substituted into Eq. 2.99 in order to obtain the normal force during sliding contact.

Regularization Analogous to the collision evaluation, it is necessary to detect and converge on the onset of stick. However, since the both surface normal $\hat{\mathbf{N}}$ and the tangential direction

change as the contact point slides, the sliding magnitude s and direction ϕ change in a non-trivial way. As a result, the approach using the stretching variable τ is not well applicable to continued contact motion. Instead, the net contact force from normal and friction interactions is regularized, analogous to what was done for rolling resistance in collisions. The regularization variable λ_F is defined as:

$$\lambda_F = \frac{\min(v_T, v_T^*)}{v_T^*} \quad (2.105)$$

This is then used to compute the regularized net contact force from normal and friction as:

$$\mathbf{F}_c = \lambda_F \mathbf{F}_{c,slip} + (1 - \lambda_F) \mathbf{F}_{c,stick} \quad (2.106)$$

The use of this regularization effectively avoids the need for a transition between sliding and sticking contact phases, but rather allows some very small v_T to exist in order to smooth the stick/slip behavior without requiring active convergence on the precise transition time.

Rolling resistance Finally, the contact model is extended by including a rolling resistance model. Again, this model is not nominally active during simulations, but has been implemented and verified. As done before, rolling resistance is accounted for through a rolling resistance torque that acts to reduce the angular velocity of the spacecraft. When the spacecraft is rolling, this torque is defined in analogy with the collision evaluation as:

$$\mathbf{L}_{RR,roll} = -[\mathbf{C}_{RR}]RF_N \frac{\boldsymbol{\omega}_{S/\mathcal{T}}}{\omega_{S/\mathcal{T}}} \quad (2.107)$$

When the spacecraft angular velocity is in ‘lock,’ the corresponding torque required to maintain $\omega_{S/\mathcal{T}} = 0$ can be derived from Eq. 2.5, as:

$$\mathbf{L}_{RR,lock} = [\tilde{\boldsymbol{\omega}}_{S/\mathcal{N}}][\mathbf{I}_S]\boldsymbol{\omega}_{S/\mathcal{N}} + [\mathbf{I}_S][\tilde{\boldsymbol{\omega}}_{S/\mathcal{N}}]\boldsymbol{\omega}_{S/\mathcal{T}} \quad (2.108)$$

This is regularized using the familiar scheme, as:

$$\mathbf{L}_{RR} = \lambda_{RR} \mathbf{L}_{RR,roll} + (1 - \lambda_{RR}) \mathbf{L}_{lock} \quad (2.109)$$

in which λ_{RR} is the rolling regularization variable, given before by Eq. 2.87. Finally, a corresponding force is applied at the spacecraft center of mass to ensure that rolling resistance does not affect the

contact point velocity, again in analogy with our rolling resistance model during collision evaluation. This can be derived from Eq. 2.54 as:

$$\mathbf{F}_{RR} = m[\tilde{\mathbf{r}}][\mathbf{I}_S]^{-1}\mathbf{L}_{RR} \quad (2.110)$$

2.6 Regolith

One of the major discoveries made by the Hayabusa spacecraft was that asteroid Itokawa is “covered with unconsolidated millimeter-sized and larger gravels” [72, 123]. This observation is important for the modeling of spacecraft-target interactions: when a spacecraft impacts a regolith-covered surface, it will interact with and ‘dig through’ the small regolith grains, dissipating energy in the process. This was witnessed most notably during the deployment of the Philae lander onboard the Rosetta mission. As previously mentioned, Philae’s active landing systems failed to anchor the lander to the surface of comet 67P/C-G, after which the craft bounced away and settled in a poorly-lit location on the surface. Imaging of the initial touchdown site before and after impact revealed that two smaller craters had appeared after the impact. This is illustrated in Fig. 2.56, taken from [11], in which the ‘B’ and ‘C’ crater features are respectively 10 and 20 cm deep with a depth of approximately 2 m. The corresponding excavated masses are respectively 60 and 120 kg. Detailed reconstruction of Philae’s descent trajectory and of the first impact indicated that “the comet surface is strongly damping: 50 to 80% of the lander’s kinetic energy before TD1 was dissipated in the comet soil.” [11] In other words, Philae experienced a *kinetic* coefficient of restitution of 0.2 to 0.5. The reconstruction also showed that, after penetrating the 20 cm regolith layer, the lander interacted with a hard surface layer underneath the regolith.

As discussed prior to our contact model development, we chose to abstract away the complex granular mechanics that occur between surface regolith and a deployed spacecraft, such as Philae’s ‘cratering’ illustrated above. Instead, effective coefficients of restitution and friction are used, that are easily tuned and allow for fast simulation of a large number of spacecraft trajectories. We also discussed how localized, high-fidelity DEM simulations could be used to empirically obtain




Figure 2.56: Surface morphology changes at Philae's first impact site [11].

expressions for the coefficients of restitution and friction as a function of impact conditions. Although this process is beyond the scope of the current thesis, there is one simple ‘modulation’ of the coefficient of restitution that may be applied, in light of Philae's bouncing behavior and some relevant experimental results.

Impact experiments Projectile impacts in Earth gravity have received considerable experimental attention, see for example the work by Uehara et al [127], Newhall and Durian [79], Ambroso et al. [5], Goldman and Umbanhowar [37] and Walsh et al. [139]. The amount of research in low-velocity, micro-gravity impacts is far smaller. Colwell and Taylor performed experiments on Space Shuttle and the ISS, during which they shot small spheres into various sands and JSR (a regolith simulant). They observed particularly low restitution coefficients of $e = 0 - 0.1$ at impact velocities of $V_0 = 10 - 100$ cm/s [24]. These results were repeated in successive experiments, in which they observed notably higher $e \sim 0.3$ at particularly low impact velocities of $V_0 \sim 5$ cm/s [23]. Murdoch et al. performed experiments with an Atwood machine in a drop tower to simulate a wide range of microgravity values $g = 0.2 - 1.0$ m/s². Dropping a 10 cm sphere into quartz sand at velocities of $V_0 = 2 - 40$ cm/s, they also did not observe any rebounds [78]. These observations run counter to the bouncing motion of Philae on comet 67P/C-G. It is possible that the Earth soils experienced more significant packing, which would affect the amount of energy dissipated during collisions with the regolith. In any case, Murdoch et al. note that the shape and angle at

which an object impacts regolith can strongly affect its rebounding behavior. Indeed, all of the aforementioned experiments were performed for (approximately) normal impacts.

To the best of our knowledge, Nishida et al. are the only authors to have investigated oblique, low-velocity impacts into granular materials. They performed experiments at $1g$ with small spheres between 6 and 20 mm in diameter and of different materials, impacting different granular materials. They varied the impact angle θ , as illustrated on the left side of Fig. 2.57. The impact angle θ is thus defined as:

$$\theta = \arcsin \left(\frac{-\mathbf{v} \cdot \hat{\mathbf{N}}}{\|\mathbf{v}\|} \right) \quad (2.111)$$




Figure 2.57: Definition of the impact angle for (left) a pointmass and (right) a complex probe.

Nishida et al. found significantly different behaviors at different impact angles [80]. More specifically, at low (tangential) θ , the spherical particles rebounded with some restitution coefficient $e > 0$. As the impact angle was increased towards more normal impacts, the corresponding restitution coefficient decreased. When θ was increased beyond some transition θ_E , the spheres no longer rebounded but instead penetrated the granular surface. Combining multiple sets of experiments with different spheres and granular materials, they obtained the following empirical relation for the critical θ_E [80]:

$$\theta_E = \left(\frac{\rho_t}{\rho_p} \right)^{1/2} \cdot \left(\frac{D_t}{D_p} \right)^{2/3} \cdot 210 \text{ deg} \quad (2.112)$$

in which ρ_t and ρ_p are the (bulk) densities of respectively the granular material and the projectile. D_t and D_p are the corresponding diameters.

Restitution modulation Eq. 2.112 can be used to mimic the effect of a granular layer

on the small-body surface on the resulting probe motion. Indeed, the coefficient of restitution may be modulated based on the effective impact angle of a probe as it collides with the small-body surface. In this, one may be inclined to simply set $e = 0$ when $\theta > \theta_E$ and to scale the restitution coefficient when $\theta < \theta_E$. However, this kind of modulation is counter to an important real-world result, namely, the dynamics observed in the bouncing motion of the Philae lander. Despite its near-normal first impact on comet 67P/C-G, Philae rebounded with an (energetic) coefficient of restitution e between 0.2 and 0.5 [13]. This occurred because the granular layer on the comet surface had a limited thickness of approximately 20 cm; Philae penetrated this soft layer and then impacted the hard surface layer underneath. In order to simulate such a surface environment, some lower limit e_0 on the coefficient of restitution is therefore applied, such that $e = e_0$ whenever the impact angle $\theta > \theta_0$. When the impact angle is reduced below its critical value, e is then increased up to some maximum e_1 in the limit of a tangential impact where $\theta = 0$ deg. This modulation is illustrated in Fig. 2.58 and can be expressed as:

$$e = e_1 + \frac{\min(\theta, \theta_E)}{\theta_E} (e_0 - e_1) \quad (2.113)$$




Figure 2.58: Modulation of the coefficient of restitution e as a function of impact angle θ .

Eq. 2.113 provides a simple model that modulates the restitution coefficient in order to account for the presence of surface regolith on top of a hard, rocky layer. Although far more advanced models based on, *e.g.*, the impact velocity or probe shape could be implemented, this would require extensive experimental work. Given that the model presented here can be used to

match the few relevant experiments, it is considered a good first attempt at accounting for regolith presence in global small-body simulations.

2.7 Numerical methods

The above sections have detailed our methods for modeling the small-body environment and the motion of a ballistically-deployed probe through this environment. Here, our numerical implementation used to actually simulate that motion is discussed. Some general characteristics and techniques applied in the numerical propagator are first presented. This is followed by a summary of how such an integrator can be classically implemented in a sequential manner that is executed on the CPU. Finally, details are provided on how the integration can be *parallelized* through implementation on the GPU, which allows for the simultaneous propagation of multiple simulations and significantly increases the speed at which a batch of simulations can be performed.

2.7.1 Trajectory propagation

During simulation of a probe's trajectory, a distinction between three 'phases' is made: *flying* motion in which the probe (ballistically) moves above the asteroid surface, *contact* motion in which the probe moves while in contact with the surface, and *collision* evaluation in which impulsive target-probe collisions are handled. The spacecraft state \mathbf{X} propagated during flying and contact phases was previously given in Eq. 2.1. This state is propagated using the equations of motion in Eq. 2.2. During collision evaluation, the collision state \mathbf{X}_c as given in Eq. 2.80 is propagated instead, using the equations of motion in Eq. 2.91.

These equations of motion are integrated numerically using a Runge-Kutta (RK) integrator, which is a form of Gauss quadrature that approximates the state \mathbf{X}_1 at some time $t_1 = t_0 + h$ based on an initial state \mathbf{X}_0 at t_0 , with h the applied step size [75]:

$$\mathbf{X}_1 = \mathbf{X}(t_0 + h) = \mathbf{X}_0 + \Delta\mathbf{X} \approx \mathbf{X}_0 + h \sum_{i=1}^p b_i \mathbf{k}_i \quad (2.114)$$

In this expression, b_i are coefficients of the applied quadrature rule; the reader is referred to [75]

for their values. Furthermore, the \mathbf{k}_i are evaluations of the equations of motion $\dot{\mathbf{X}}(\dots)$ at different times t and states \mathbf{X} , given generally as:

$$\mathbf{k}_i = \dot{\mathbf{X}}(t_0 + h \cdot c_i, \mathbf{X}_0 + h \cdot \sum_{j=1}^{i-1} a_{ij} \mathbf{k}_j) \quad (2.115)$$

in which a_{ij} and c_i are quadrature coefficients similar to the aforementioned b_i and that can also be found in [75]. Furthermore, it holds that $\mathbf{k}_0 = \mathbf{0}$. The upper limit p of the sum in Eq. 2.114 is called the *order* of the applied integration method. The use of higher-order methods allows for larger step sizes h while still achieving some accuracy, but does require additional evaluations of the equations of motion. Here, the RK5(4) scheme is applied, which uses a fifth-order method for propagation purposes. Additionally, the scheme uses a fourth-order method to estimate the accuracy of the integration. Based on the accuracy of an integration with step size h , it is possible to estimate the step size h^* necessary to achieve some particular tolerance ϵ . If \mathbf{X}_1 is the updated state computed with a high-order method and \mathbf{X}_1^* is that state computed with a lower-order method, the new step size h^* can be computed as [75]:

$$h^* = h \cdot \left(\frac{\epsilon}{\max(|\mathbf{X}_1 - \mathbf{X}_1^*|)} \right)^{\frac{1}{p+1}} \quad (2.116)$$

In our propagation, a *relative* tolerance of 10^{-10} is used. After each integration step, the spacecraft and target attitude quaternion are normalized. This is then followed by an event check that evaluates some even function to obtain the event vector $\mathbf{E}(\mathbf{X})$ of size $Q \times 1$, where Q is the number of tracked events. Each of the vector elements e_i quantifies some criterion about the spacecraft state \mathbf{X} . For the spacecraft-target collision event, this element is equal to the minimum distance $\min(\mathbf{d})$ between the target and the spacecraft, as computed from the SDF with Eq. 2.17.

Using this event vector, an event is detected when the sign of any of the event vector elements changes between successive time steps. When this occurs, the main integration loop is adjusted and event convergence begins, in which the simulation converges on the particular time t_e at which the event takes place. A flowchart of the event convergence algorithm is provided in Fig. 2.59. Effectively, this algorithm continues to reduce the time step h while evaluating the event function,

until $h \leq \Delta t_e$. Once this happens, the simulation has converged on the event, such that its effect may be evaluated.




Figure 2.59: Flowchart of event convergence algorithm.

2.7.2 Sequential simulation on the CPU

In a standard implementation of this integration algorithm, simulations are performed sequentially. In this, the computer handles a single simulation at any given time. The overall algorithm applied in this sequential simulation is illustrated in Fig. 2.60. The relevant shape, gravity, and spacecraft models necessary to enable simulation are first set up. The first simulation is then initialized with some set of initial conditions (position, velocities, and attitude). The simulation then consists of a single integration loop in which repeated RK5(4) integration steps are performed, followed by an event check and (if necessary) event handling. Once the spacecraft has settled on the target surface, the simulation outputs are saved. The algorithm then initializes the

second simulation, propagates it, and saves the corresponding outputs. This is repeated for all N simulations.

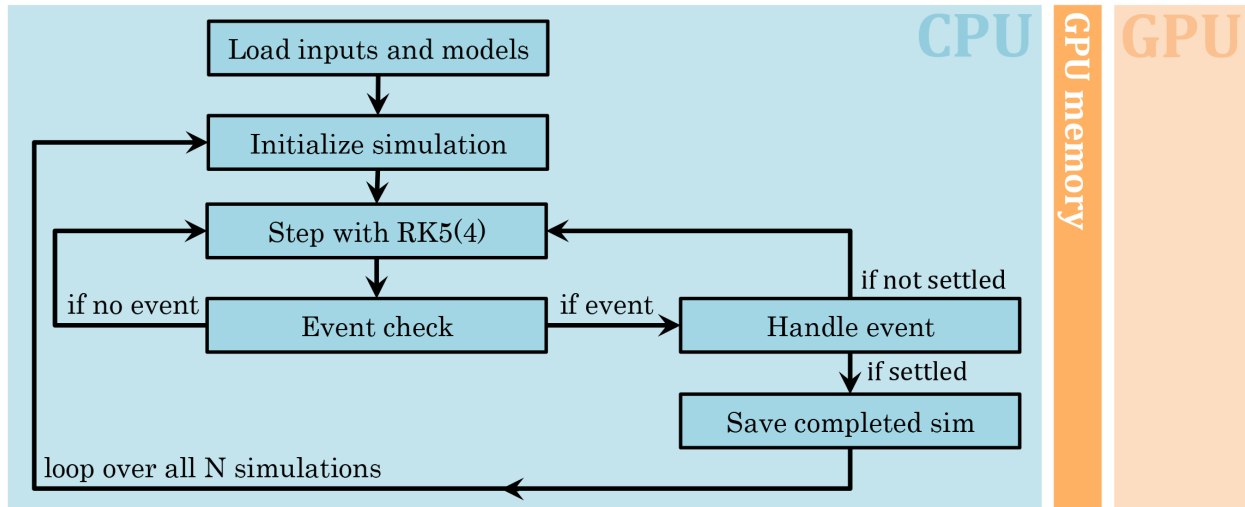


Figure 2.60: Flowchart of serial simulation algorithm.

2.7.3 Parallel simulation on the GPU

Although the sequential integration discussed above is relatively straightforward to implement, it is slow to evaluate a sizable amount of simulations. This is because the sequential propagation is limited to handling a single simulation at any given time. To increase the speed at which simulations can be performed, the simulation methods are adjusted to allow for parallelized implementation on a GPU, which may have several thousand computation cores and are thus able to simultaneously perform several thousand simulations. Although such an implementation has the potential of significantly increasing the simulation speed, parallelization requires efficient memory management. More specifically, it holds that the CPU (where models and simulations are initialized) and GPU (where simulations are actually performed) have separate memories. Any exchange of variables between the two memories must be commanded in code. Although the GPU is efficient at performing computations, memory exchange (both read and write) operations are much slower and should therefore be avoided whenever possible.

In a naive implementation, one may be inclined to simply have each of the GPU cores handle a single simulation by executing the inner loop of the sequential simulation algorithm, as previously illustrated in Fig. 2.60. However, this approach severely limits the computational efficiency due to the manner in which GPU cores execute a piece of code (a *kernel*). All cores that execute a given kernel simultaneously step through that kernel's lines of code. When a branching statement occurs, only one branch of the statement can be executed at a given time. Consider the following example: when an `if/else` statement is included in a GPU kernel, all cores that follow the `if` branch will execute first, while those following the `else` branch are stalled. Only after the `if` branch is completed does the other set of cores execute the `else` branch, during which time the former set of cores is stalled. As a result, such branching statements significantly limit the effective 'duty cycle' of the GPU cores. In our small-body probe simulations, the simultaneous execution of flying motion, contact motion, and collision evaluation should therefore be avoided within any given *group* of GPU cores. Instead, the simulations that are in those three phases should be grouped together and allocated to separate groups of cores. Our GPU implementation therefore includes three separate integration kernels (the 'fly' kernel, the 'collision' kernel, and the 'contact' kernel) that can be executed independently and therefore simultaneously.

Fig. 2.61 illustrates our parallelized simulation implementation. The algorithm works as follows: first, the relevant models are loaded and written onto the GPU memory. Next, the initial conditions of the simulations to be performed are loaded. Given the GPU capabilities, several hundred thousand trajectories are often simulated. Due to the limitations on available GPU memory, it holds that a large part of that memory will be taken up by the target shape and gravity models, with limited space remaining for trajectory data. Furthermore, typical GPUs have only a few thousand computation cores. As a result, not all N initial conditions to be simulated are loaded onto the GPU memory at any given time. Instead, the GPU only ever handles $n < N$ 'active' simulations at any given time.

At the first initialization of the simulation code, all n active simulations are prepared for the propagation of flying motion. The 'flying' kernel that is executed on the GPU performs integrations

for all active simulations that are in the ‘flying’ mode. Whenever one wishes to investigate or visualize the entire trajectory of a bouncing probe, the probe state should be stored at each time step provided by the integrator. This requires the CPU to read the GPU memory after every time step and thus implies a significant computational cost, relative to the overall speed of GPU code execution. Furthermore, in most cases one is only interested in the critical information of each simulated trajectory, namely, the velocity changes occurring in each target-spacecraft collision, and the final spacecraft settling state. In these cases, the kernel can be allowed to perform multiple integration steps before the GPU memory is read by the CPU. In order to allow for these two modes of propagation, a parameter k_{fly} is set and determines the number of integration steps to be executed by the kernel before the GPU memory is read by the CPU. When this data reading is finished, the CPU determines which simulations should continue flying propagation and which simulations should switch to respectively contact motion and collision evaluation. The code then continues (often simultaneously) performing k_{fly} , $k_{contact}$, and $k_{collision}$ integration steps for the respective modes of motion. After the completion of each respective set of integration steps, the CPU determines which simulations should switch propagation mode. The CPU controller also checks how many of the n simulations have been completed. Whenever $n_{done} \geq n_{update}$, the kernel integrations are paused and new simulations are added from the global list of N simulations to be performed onto the active list of n simulations currently active. This process repeats until all simulations have been completed.

In summary, the parallel GPU implementation offers a simulation speedup of roughly 100 times relative to a serial CPU implementation, for typical scenarios. On a 2018 high-performance machine with a NVIDIA GeForce GTX 1080 TI GPU, typically 20 simulations per second can be performed. This enables broad studies of lander deployment, which would be impossible with the more classical serial CPU implementation.

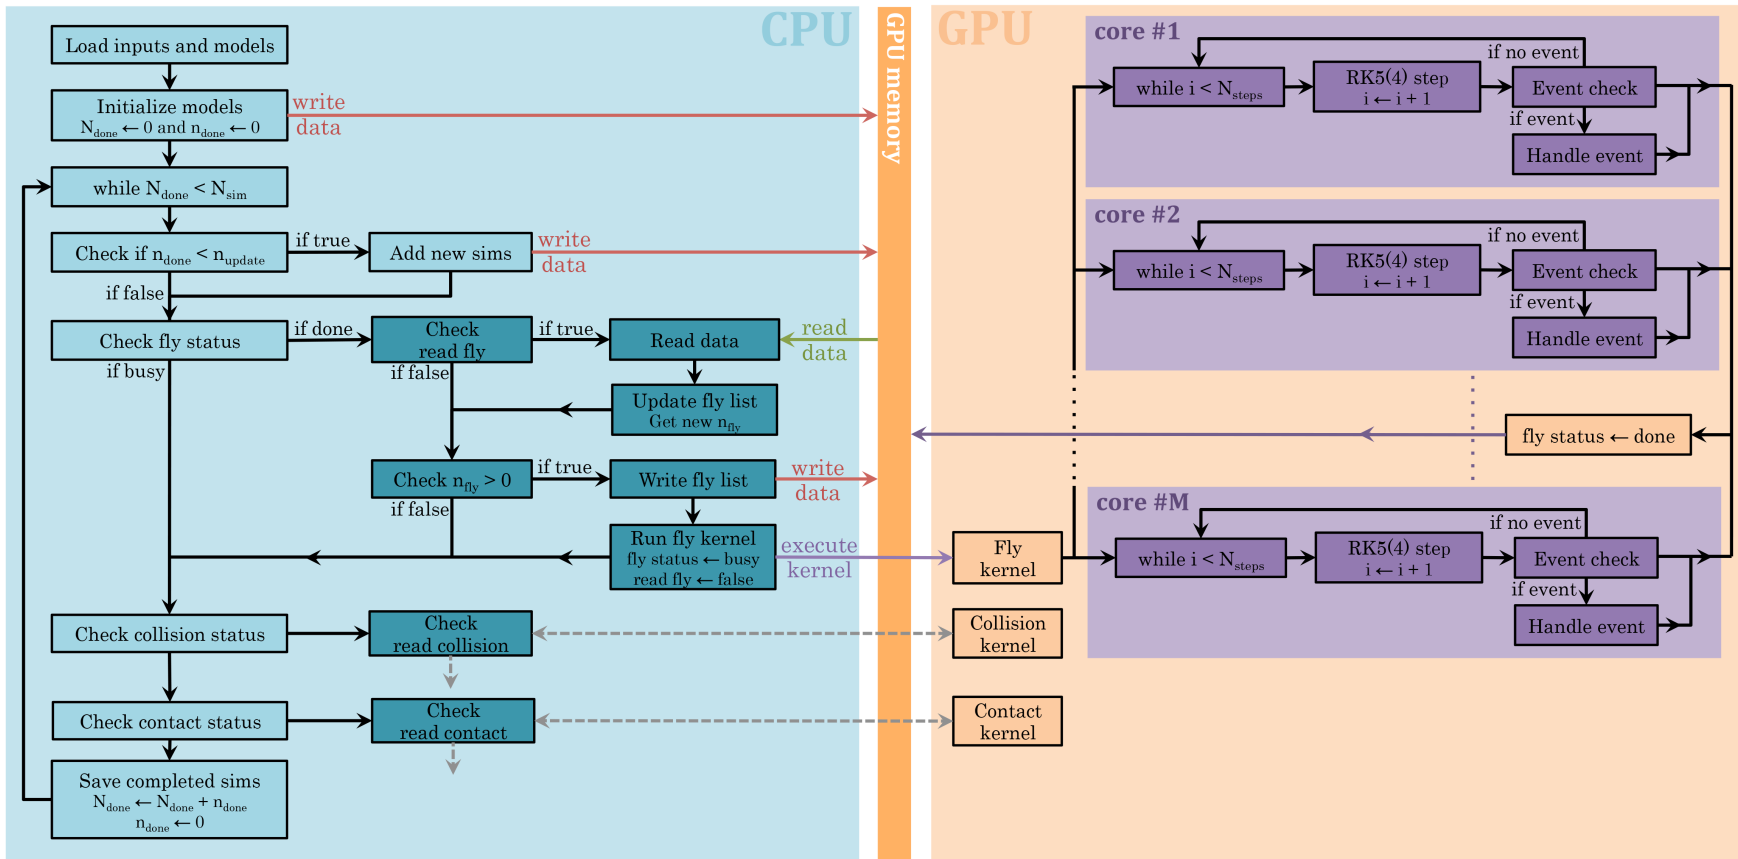


Figure 2.61: Flowchart of parallel simulation algorithm.

Chapter 3

The lift-off velocity

In June or July of 2018, the Japanese Hayabusa-2 spacecraft will arrive at asteroid Ryugu and begin exploration of the asteroid. During its stay at the asteroid, the spacecraft will deploy three hoppers (MINERVA-II) and one lander (MASCOT) [125, 43, 148]. These payloads are deployed by the spacecraft from an altitude of approximately 55 m, after which they descend to the surface and performed several bounces before coming to rest. The rovers are equipped with several scientific instruments that allow them to obtain *in situ* measurements of the surface, which are relayed to Earth via the Hayabusa-2 mothership. The probes also carry internal *momentum exchange mechanisms*, that allow them to ‘hop’ across Ryugu’s surface. By performing repeated hops, the probes are able to obtain measurements from multiple sites, which greatly increases the return of their surface exploration operations.

An important consideration in the design of these (hopping) mobility operations is the velocity imparted on a probe by its momentum device, which in turn determines the distance covered and time spent hopping [134, 118]. When a hopper is given a velocity *normal* to the local surface, it is guaranteed to lift off from that surface and perform a (possibly brief) hop, of which the complex dynamics are difficult to predict analytically. If the probe instead moves *tangential* to the local target body surface, lift-off is no longer guaranteed but requires some minimum velocity. Indeed, an arbitrary location on the small-body surface has some associated tangential velocity at which a moving particle will locally lift off from that surface and enter a period of ballistic flight. This velocity is defined as the *lift-off velocity* and varies as a function of the desired lift-off direction.

Although mobile platforms in small-body environments are likely to employ both a tangential *and* a normal velocity in their hopping operations, the assumption of a strictly tangential velocity permits an analytic definition and study of the lift-off velocity, which can be used in the determination of preliminary hopping parameters. The velocity is also useful in establishing ‘speed limits’ for surface operations of wheeled vehicles in more gravitationally tractable environments, if lift-off is not to occur. Finally, the lift-off velocity determines when rocks or dust particles, moving on a small body, will cease contact with its surface after being made mobile by geophysical processes such as landslides.

The concept of the lift-off velocity was first introduced by Staley, who studied lift-off from the long ends of a rotating asteroid in his article *Man on asteroid* [109]. Staley suggests performing “*a much more extensive analysis at other locations on an asteroid.*” The current chapter of this thesis presents precisely that analysis and is structured as follows. First, a general formulation for the lift-off velocity on the surface of an arbitrary body is derived, in Section 3.1. This formulation is reduced to three distinctly different cases: lift-off from a curved surface, a sharp ridge, and a flat plane. The magnitude of the lift-off velocity is found to be strongly dependent on the local surface curvature of a considered body. Section 3.2 therefore details methods to compute this surface curvature for various small-body shape models. The resulting lift-off velocity expressions are then applied to some simple shapes in Section 3.3, in order to establish general trends in the distribution of the lift-off velocity magnitude across the surface of a body. Finally, the expressions are applied to several small bodies of the Solar System in Section 3.4, in which global velocity distributions are provided for asteroids 1999 KW₄ Alpha, Bennu, and Eros, as well as the Mars moon Phobos.

3.1 Deriving the lift-off velocity

We define the *lift-off velocity*, V_θ , as the tangential surface velocity given to an object on a body with arbitrary shape, rotation, and gravity, at which the object will locally lift off from that body’s surface in some desired direction, $\hat{\mathbf{b}}_\theta$. The geometry of this set-up is illustrated on the left side of Fig. 3.1.

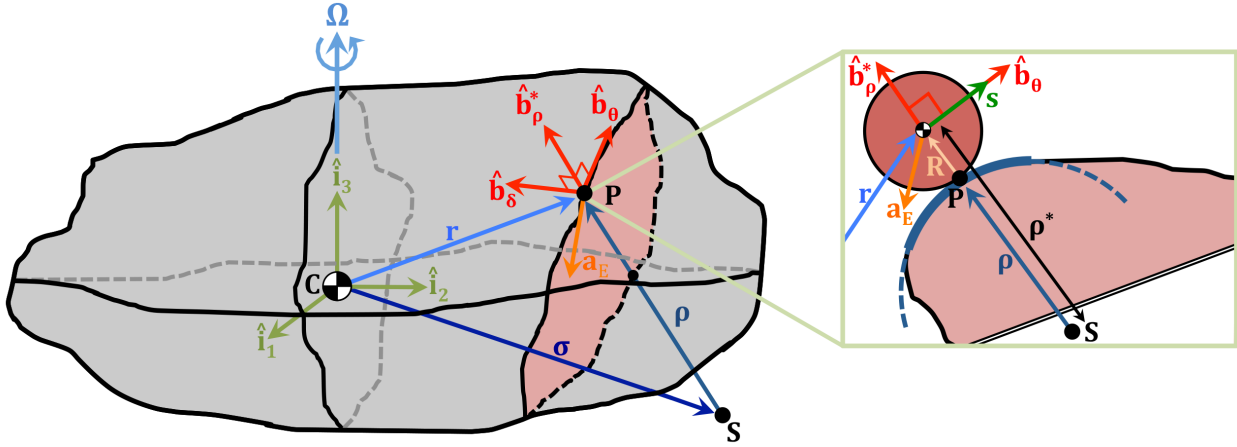


Figure 3.1: (left) Full and (right) departure plane geometry applied in the derivation of the lift-off velocity.

3.1.1 Geometry

Following the geometry of Fig. 3.1, two orthonormal reference frames are defined: the inertial frame \mathcal{I} and the rotating departure frame \mathcal{B} , as:

$$\begin{aligned} \mathcal{I} &: \{ \mathbf{C} \mid \hat{\mathbf{i}}_1, \hat{\mathbf{i}}_2, \hat{\mathbf{i}}_3 \} \\ \mathcal{B} &: \{ \mathbf{P} \mid \hat{\mathbf{b}}_\rho^*, \hat{\mathbf{b}}_\theta, \hat{\mathbf{b}}_\delta \} \end{aligned} \quad (3.1)$$

in which the $\hat{\mathbf{i}}_3$ axis is aligned with the rotation axis $\boldsymbol{\Omega}$ of the target body. The considered object, with radius vector \mathbf{R} , has its center of mass at some \mathbf{r} relative to the center of mass of the target body, and is in contact with that body at the departure point \mathbf{P} . The radius vector \mathbf{R} points from the departure point to the object center of mass, such that $\mathbf{r} = \mathbf{P} + \mathbf{R}$. With this definition, our theory holds for objects with an arbitrary shape, although Fig. 3.1 depicts a sphere for simplicity.

The object is (inertially) subjected to an external acceleration \mathbf{a}_E , which consists chiefly of the gravitational acceleration $\mathbf{g}(\mathbf{r})$ from the target body, but may also include other contributions such as solar radiation pressure or tidal effects. It is also subject to an acceleration \mathbf{a}_N from the surface normal force. The Coriolis and centrifugal accelerations are *not* included in \mathbf{a}_E but will instead appear in the final lift-off velocity expression through our choice of reference frame.

When the object is given some surface velocity \mathbf{V} , its motion will be bound to the osculating *departure plane* that is spanned by \mathbf{V} and the net acceleration \mathbf{a}_E on the object. This plane varies as a function of the direction $\hat{\mathbf{b}}_\theta$ of the surface velocity $\mathbf{V} = V\hat{\mathbf{b}}_\theta$, which can be freely chosen through rotation over some azimuthal angle η around the local surface normal $\hat{\mathbf{N}}$, but must always be perpendicular to that normal. The departure plane is shown in red in Fig. 3.1.

Although the target body may have an arbitrary shape, its surface can be approximated within the departure plane and at the departure point with some *radius of curvature*, ρ , and corresponding *center of curvature*, \mathbf{S} . The vector from the center of curvature to the object center of mass defines the $\hat{\mathbf{b}}_\rho^*$ axis; the $\hat{\mathbf{b}}_\delta$ axis right-handedly completes the rotating $\mathcal{B} : \{\mathbf{r} \mid \hat{\mathbf{b}}_\rho^*, \hat{\mathbf{b}}_\theta, \hat{\mathbf{b}}_\delta\}$ departure frame and is normal to the departure plane. The target body is assumed to be uniformly rotating about the $\hat{\mathbf{i}}_3$ axis of the inertial $\mathcal{I} : \{\mathbf{C} \mid \hat{\mathbf{i}}_1, \hat{\mathbf{i}}_2, \hat{\mathbf{i}}_3\}$ frame with angular velocity $\boldsymbol{\Omega} = \Omega\hat{\mathbf{i}}_3$, though the results can be generalized to non-principal-axis rotators.

In Fig. 3.1, the object center of mass, departure point, and center of curvature are shown as collinear for simplicity; these points are not aligned in the general case. Furthermore, it can be seen that the object center of mass is located at some distance ρ^* from the center of curvature \mathbf{S} . The value of this ρ^* , denoted *effective radius of curvature*, will be some combination of ρ and R that will be commented on shortly. Finally, in order to properly define the lift-off velocity, it is assumed that the net acceleration on the stationary object is such that it does not lift off from the small body surface by itself, *i.e.*, that $\hat{\mathbf{N}} \cdot ([\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}]\mathbf{r} - \mathbf{a}_E) \geq 0$, in which $[\tilde{\boldsymbol{\Omega}}]$ is the cross-product tensor of $\boldsymbol{\Omega}$. This expression simply states that the combined external and centrifugal accelerations must be directed against the (outwards-pointing) local surface normal, *i.e.*, that these accelerations do not pull the particle off the surface.

3.1.2 General expression

The object, located initially at $\mathbf{r}(0) = \mathbf{P} + \mathbf{R}$, is given a velocity \mathbf{V} along $\hat{\mathbf{b}}_\theta$ and will be displaced over some infinitesimal distance \mathbf{s} after a time dt . Since the particle velocity \mathbf{V} is aligned

with \mathbf{s} , one may write that:

$$\dot{\mathbf{s}} = \dot{\mathbf{s}}\hat{\mathbf{b}}_\theta = \mathbf{V} = V\hat{\mathbf{b}}_\theta \quad (3.2)$$

Applying the transport theorem to obtain the derivative of Eq. 3.2, the second derivative of \mathbf{s} is found as:

$$\ddot{\mathbf{s}} = \dot{V}\hat{\mathbf{b}}_\theta + \boldsymbol{\omega} \times (V\hat{\mathbf{b}}_\theta) \quad (3.3)$$

When the particle is given the proper lift-off velocity V_θ , it will track the surface curvature corresponding to $\{\mathbf{S}, \rho^*\}$ and perform a circular motion about the center of curvature. It therefore holds that, at lift-off, $\boldsymbol{\omega} = (V_\theta/\rho^*)\hat{\mathbf{b}}_\delta$. Substitution of this rate into Eq. 3.3 yields:

$$\ddot{\mathbf{s}} = \dot{V}_\theta\hat{\mathbf{b}}_\theta - \frac{V_\theta^2}{\rho^*}\hat{\mathbf{b}}_\rho^* \quad (3.4)$$

Combining the planar geometry on the right side of Fig. 3.1 with the full geometry on the left side, the particle position at time dt can be expressed as:

$$\mathbf{r} = \mathbf{r}(dt) = \mathbf{P} + \mathbf{R} + \mathbf{s} \quad (3.5)$$

Applying the transport theorem, the inertial derivative of \mathbf{r} is found as:

$$\dot{\mathbf{r}} = \dot{\mathbf{P}} + \dot{\mathbf{R}} + \dot{\mathbf{s}} + \boldsymbol{\Omega} \times (\mathbf{P} + \mathbf{R} + \mathbf{s}) \quad (3.6)$$

As the departure point \mathbf{P} is fixed to the target body, it must hold that $\dot{\mathbf{P}} = 0$. Assuming that the object is rigid, it also holds that $\dot{\mathbf{R}} = 0$. Substituting this into Eq. 3.6 and making use of the tensor notation $\tilde{\boldsymbol{\Omega}}$ to express the vector cross product and \cdot to express the vector dot product, one may write that:

$$\dot{\mathbf{r}} = \dot{\mathbf{s}} + [\tilde{\boldsymbol{\Omega}}] (\mathbf{P} + \mathbf{R} + \mathbf{s}) \quad (3.7)$$

Applying the transport theorem once more, the inertial acceleration of the object is found as:

$$\ddot{\mathbf{r}} = \ddot{\mathbf{s}} + [\tilde{\boldsymbol{\Omega}}]\dot{\mathbf{s}} + [\dot{\tilde{\boldsymbol{\Omega}}}] (\mathbf{P} + \mathbf{R} + \mathbf{s}) + [\tilde{\boldsymbol{\Omega}}] (\dot{\mathbf{P}} + \dot{\mathbf{R}} + \dot{\mathbf{s}}) + [\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}] (\mathbf{P} + \mathbf{R} + \mathbf{s}) \quad (3.8)$$

At the exact instant of lift-off, the infinitesimal displacement $\mathbf{s} \rightarrow 0$. Using also that $\dot{\mathbf{P}} = \dot{\mathbf{R}} = 0$, and assuming that the body is in uniform rotation such that $\dot{\tilde{\boldsymbol{\Omega}}} = 0$, Eq. 3.8 reduces to:

$$\ddot{\mathbf{r}} = \ddot{\mathbf{s}} + 2[\tilde{\boldsymbol{\Omega}}]\dot{\mathbf{s}} + [\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}] (\mathbf{P} + \mathbf{R}) \quad (3.9)$$

This expression is similar to the equations of motion of a particle in a rotating reference frame, though the set-up combining \mathbf{P} , \mathbf{R} , and \mathbf{s} is somewhat unique to our application. Recognize the appearance of the Coriolis acceleration $2[\tilde{\boldsymbol{\Omega}}]\dot{\mathbf{s}}$ and the centrifugal acceleration $[\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}](\mathbf{P} + \mathbf{R})$. These *virtual accelerations* result from our use of the rotating, non-inertial departure frame.

A particle on the surface of a gravitating body is subject to both an external acceleration \mathbf{a}_E and an acceleration \mathbf{a}_N from the normal force. Applying Newton's second law of motion to Eq. 3.9, and substituting the kinematic expressions from Eqs. 3.2 and 3.4 for $\dot{\mathbf{s}}$ and $\ddot{\mathbf{s}}$, one may write that:

$$\ddot{\mathbf{r}} = \mathbf{a}_N + \mathbf{a}_E = \dot{V}_\theta \hat{\mathbf{b}}_\theta - \frac{V_\theta^2}{\rho^{*2}} \boldsymbol{\rho}^* + 2V_\theta [\tilde{\boldsymbol{\Omega}}] \hat{\mathbf{b}}_\theta + [\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}] (\mathbf{P} + \mathbf{R}) \quad (3.10)$$

This can be rewritten to yield the normal acceleration as:

$$\mathbf{a}_N = \dot{V}_\theta \hat{\mathbf{b}}_\theta - \frac{V_\theta^2}{\rho^{*2}} \boldsymbol{\rho}^* + 2V_\theta [\tilde{\boldsymbol{\Omega}}] \hat{\mathbf{b}}_\theta + [\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}] (\mathbf{P} + \mathbf{R}) - \mathbf{a}_E \quad (3.11)$$

Eq. 3.11 expresses the acceleration required to prevent the particle from penetrating the body surface. By inspecting the direction of this acceleration with respect to the surface curvature, one can determine whether or not the particle will lift off when moving with some velocity V_θ . If the surface is locally convex, it holds that lift-off occurs when $\hat{\mathbf{b}}_\rho^* \cdot \mathbf{a}_N \leq 0$. Taking the dot product of $\hat{\mathbf{b}}_\rho^*$ onto Eq. 3.11, it holds that:

$$-V_\theta^2 - 2\rho^* V_\theta \hat{\mathbf{b}}_\theta [\tilde{\boldsymbol{\Omega}}] \hat{\mathbf{b}}_\rho^* + \rho^* \hat{\mathbf{b}}_\rho^* \cdot \left([\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}] (\mathbf{P} + \mathbf{R}) - \mathbf{a}_E \right) \leq 0 \quad (3.12)$$

in which it was used that $\boldsymbol{\rho}^* = \rho^* \hat{\mathbf{b}}_\rho^*$ and $\boldsymbol{\rho}^* \cdot \boldsymbol{\rho}^* = \rho^{*2}$. This expression shows that, for finite values of ρ^* , there always exists some V_θ such that Eq. 3.12 indeed yields a negative value. The minimum velocity required for lift-off in the $\hat{\mathbf{b}}_\theta$ direction is found from the limit case in which $\hat{\mathbf{b}}_\rho^* \cdot \mathbf{a}_N = 0$. Setting Eq. 3.12 equal to zero and solving for the lift-off velocity V_θ yields:

$$V_\theta = \pm \sqrt{\rho^{*2} \left(\hat{\mathbf{b}}_\theta [\tilde{\boldsymbol{\Omega}}] \hat{\mathbf{b}}_\rho^* \right)^2 + \rho^* \hat{\mathbf{b}}_\rho^* \cdot \left([\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}] (\mathbf{P} + \mathbf{R}) - \mathbf{a}_E \right) - \rho^* \hat{\mathbf{b}}_\theta [\tilde{\boldsymbol{\Omega}}] \hat{\mathbf{b}}_\rho^*} \quad (3.13)$$

The scalar triple product in this equation can be rewritten as:

$$\hat{\mathbf{b}}_\theta [\tilde{\boldsymbol{\Omega}}] \hat{\mathbf{b}}_\rho^* = \boldsymbol{\Omega} \cdot \left(\hat{\mathbf{b}}_\rho^* \times \hat{\mathbf{b}}_\theta \right) = \boldsymbol{\Omega} \cdot \hat{\mathbf{b}}_\delta = \hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega} \quad (3.14)$$

Note the appearance of $\boldsymbol{\Omega}$ without the cross-product tensor in this expression. This allows us to simplify Eq. 3.13 as:

$$V_\theta = \pm \sqrt{\rho^{*2} (\hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega})^2 + \rho^* \hat{\mathbf{b}}_\rho^* \cdot ([\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}] (\mathbf{P} + \mathbf{R}) - \mathbf{a}_E) - \rho^* \hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega}} \quad (3.15)$$

Eq. 3.15 provides a vectorial, frame-independent expression for the lift-off velocity at departure point \mathbf{P} and in direction $\hat{\mathbf{b}}_\theta$. Two solutions are possible, as the square root is preceded by a \pm sign. Although the lift-off direction $\hat{\mathbf{b}}_\theta$ does not explicitly appear in this expression, its choice directly determines the $\hat{\mathbf{b}}_\rho^*$ and $\hat{\mathbf{b}}_\delta$ vectors. The two solutions correspond to lift-off along the departure plane in respectively the $+\hat{\mathbf{b}}_\theta$ and $-\hat{\mathbf{b}}_\theta$ direction.

3.1.3 Lift-off from a curved surface

The computation of the effective radius of curvature ρ^* in Eq. 3.15 from ρ and R is not obvious for arbitrary geometries, and complicates an evaluation of the lift-off velocity. Fortunately, the case in which the body radius of curvature ρ is much larger than the object radius R permits a reduction of ρ^* that simplifies the lift-off velocity expression and provides useful insight. In order to carry out this simplification, the geometry of the object and the corresponding departure plane is inspected, while viewed along the $\hat{\mathbf{b}}_\theta$ axis, as illustrated in Fig. 3.2.




Figure 3.2: Geometry viewed along the $\hat{\mathbf{b}}_\theta$ axis.

We note again that although the illustration shows a spherical object, the following derivation holds for any object shape with radius vector \mathbf{R} from the departure point to the object center of mass. Although the dimensions of this figure are exaggerated, this illustration shows that the departure plane and frame differ slightly when computed at respectively \mathbf{r} or \mathbf{P} . This is true in the general case where \mathbf{S} , \mathbf{P} , and \mathbf{r} are *not* collinear. Correspondingly, the departure plane and frame will also change slightly when the object radius R is varied. These factors introduce a dependence of the lift-off velocity on the object radius R , which makes its computation somewhat involved. To simplify, the law of cosines is applied to Fig. 3.2, which yields for the radius of curvature ρ :

$$\rho^2 = \rho^{*2} + R^2 - 2\rho^* R \cos \alpha \quad (3.16)$$

in which ρ^* is again the effective radius of curvature. Assuming a locally convex surface, Eq. 3.16 can be rewritten to yield ρ^* , as:

$$\rho^* = R \cos \alpha + \sqrt{\rho^2 - R^2 \sin^2 \alpha} \quad (3.17)$$

Dividing by the radius of curvature ρ that is computed at \mathbf{P} , the *curvature ratio* can be written as:

$$\frac{\rho^*}{\rho} = \frac{R}{\rho} \cos \alpha + \sqrt{1 - \frac{R^2}{\rho^2} \sin^2 \alpha} \quad (3.18)$$

This expression can be expanded around $R/\rho = 0$ with a Taylor series as:

$$\frac{\rho^*}{\rho}(R) = 1 + \frac{R}{\rho} \cos \alpha + O\left(\frac{R}{\rho}\right)^2 + \text{H.O.T.} \quad (3.19)$$

which provides a first-order approximation to the curvature ratio for an object whose radius is small relative to the surface curvature it is resting on. The small bodies that are envisioned as targets for lander/rover missions are many orders of magnitude larger than the craft that would operate on their surfaces. For realistic scenarios, it therefore holds that the R/ρ ratio must indeed be very small, and it is found from Eq. 3.19 that $\rho^* \simeq \rho$, *i.e.*, that the effective radius of curvature is, to first order, equal to the radius of curvature. Note that this approximation is not valid at sharp edges where ρ is small, as in that case the R/ρ ratio is not small; this situation will be addressed

shortly. In a similar manner, it can be shown that the $\hat{\mathbf{b}}_\rho^*$ axis is approximately equal to the $\hat{\mathbf{b}}_\rho$ axis, by again applying the law of cosines to find an expression for the δ angle:

$$R^2 = \rho^{*2} + \rho^2 - 2\rho\rho^* \cos \delta \quad (3.20)$$

Using the earlier result that $\rho^* \simeq \rho$, the above can be rewritten as:

$$\cos \delta \simeq 1 - \frac{1}{2} \left(\frac{R}{\rho} \right)^2 \quad (3.21)$$

Since R/ρ is very small, it holds that $\cos \delta \simeq 1$, such that $\delta \simeq 0$, and $\hat{\mathbf{b}}_\rho^* \simeq \hat{\mathbf{b}}_\rho$. In words, the results that $\rho^* \simeq \rho$ and $\hat{\mathbf{b}}_\rho^* \simeq \hat{\mathbf{b}}_\rho$ state that, when the object radius is very small compared to the target body surface curvature, it holds that the departure plane, departure frame, and radius of curvature can be computed at the departure point \mathbf{P} instead of the object center of mass \mathbf{r} . This effectively removes the dependence of these quantities on the object radius R . Furthermore, it allows for Eq. 3.15 for the lift-off velocity to be reduced to:

$$V_\theta = \pm \sqrt{\rho_\theta^2 \left(\hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega} \right)^2 + \rho_\theta \hat{\mathbf{b}}_\rho \cdot \left([\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}]\mathbf{P} - \mathbf{a}_E \right) - \rho_\theta \hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega}} \quad (3.22)$$

in which was applied, by analogy, that also $\mathbf{P} + \mathbf{R} \simeq \mathbf{P}$. Once again, this shows that the radius of the object may be ignored when computing its lift-off velocity, as long as it is much smaller than the surface radius of curvature. Eq. 3.22 also shows that the lift-off velocity magnitude is chiefly determined by this radius of curvature; the task of generating lift-off velocity distributions across the surface of a target body therefore primarily consists of computing the radius of curvature distribution across that surface.

Properties of the solution The lift-off velocity expression given by Eq. 3.22 is extremely versatile: it does not restrict the shape, rotation, or gravitational model of the target body from which an object is lifting off. Instead, it merely requires a two-dimensional averaging of the body surface along the desired departure plane, which is spanned by the local gravitational acceleration vector (including rotation effects) and the desired departure direction. Due to the rotation and non-central gravity of the target body,² the object will generally not remain bound to this departure

plane. Nevertheless, the plane constructed following the above strategy corresponds to the osculating plane of motion to which the object will remain closely bound as it moves in the neighborhood of the departure point. The rotation and gravity of the target body are included in the expression simply as vectors, allowing for the application of any desired gravitational model and rotation axis. As mentioned before, the \mathbf{a}_E term may include accelerations from sources other than the target body, such as third-body attraction, tidal effects, and solar radiation pressure. Finally, it is noted that the derivation did assume the body surface to be locally convex. This assumption, and the lift-off dynamics on concave surface, will be further discussed further at the end of this section.

The lift-off expression provides two solutions, for which the square root is taken respectively with a positive or negative sign. These two solutions correspond to lift-off in the local east and west directions along the desired departure plane, respectively along and against the local surface rotational velocity. The two solutions will have different magnitudes for the general case in which $\hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega} \neq 0$, *i.e.*, when the body rotation has a component along the axis normal to the departure plane. Fig. 3.3 shows the variation of the two lift-off velocities at an arbitrary position on the surface of a rotating ellipsoid.

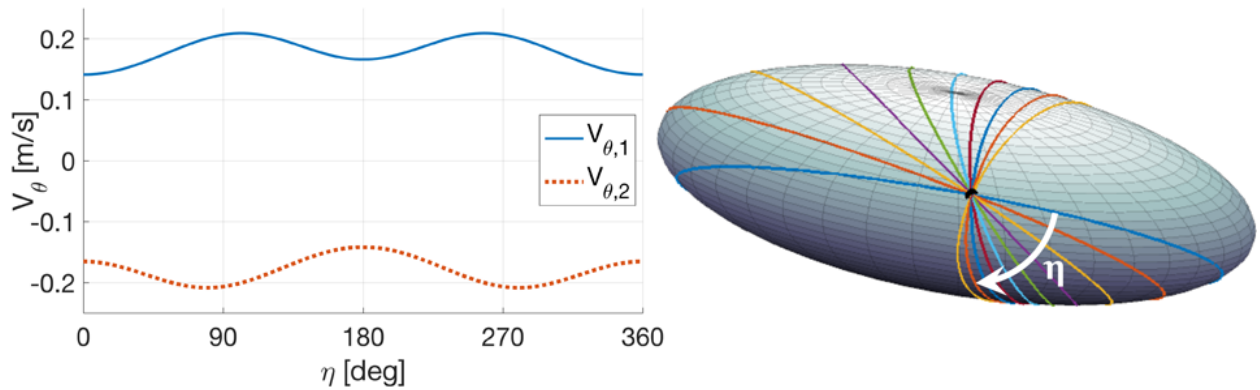


Figure 3.3: (left) Variation of the two lift-off velocities at an arbitrary position on a rotating ellipsoid, and (right) illustration of the azimuthal angle η .

In this plot, the azimuthal angle η expresses the departure direction, which is varied over 360° starting from some initial value. The methodology applied to obtain the actual curves are discussed in detail in Section 3.3; the present discussion is focused on the relative behavior of the two curves.

The two curves are mirror images of one another that are shifted by 180° along the η axis. This phenomenon is easily explained: a particular departure direction and plane are associated with some positive and negative lift-off velocities of different magnitude. When this direction is rotated by 180° , the departure plane remains unchanged but the $\hat{\mathbf{b}}_\theta$ axis, which indicates the direction of a ‘positive’ velocity, is flipped. Therefore, the magnitude of the new $V_{\theta,1}$ assumes the opposite value of $V_{\theta,2}$ at the initial orientation, and vice versa. From this analysis, it can be seen that if the minimum and maximum lift-off velocities encountered at some point are to be computed, one may either study the variation of $V_{\theta,1}$ over $0 \leq \eta \leq 360^\circ$, or study the variation of both $V_{\theta,1}$ and $V_{\theta,2}$ over $0 \leq \eta \leq 180^\circ$. The former strategy will be used in this work.

On a final note, it is repeated that the condition for lift-off used in the earlier derivation was $\hat{\mathbf{b}}_\rho \cdot \mathbf{a}_N \leq 0$. This condition sets the motion of a particle with a purely tangential velocity apart from that of particles with a normal velocity component; indeed, when a particle is given some normal velocity, this criterion is *always* satisfied.

3.1.4 Lift-off from a ridge

As discussed, Eq. 3.22 is applicable only when ρ is large relative to R . However, the case where ρ is small relative to R is also relevant, as $\rho \rightarrow 0$ when the object is moving on a sharp ridge or asperity on the target body. Although lift-off from such an asperity is likely to result in only a very brief ‘air time,’ the minimum velocity required to do so provides an important measure to the surface motion of a vehicle: if the surface is rough, for example when it is covered with small rocks, the vehicle will no longer be in continuous contact with the surface if it exceeds this ‘ridge’ lift-off velocity, V_R . This has implications to the friction experienced by the vehicle, as well as to the traction it can exert on the surface. In fact, it can be argued that the ridge velocity measure is more relevant to the motion of a surface vehicle than the curved lift-off velocity. This is true because the curved velocity is an *averaged* measure, which cannot be achieved on a real body with a wheeled vehicle. Indeed, as a vehicle exceeds the ridge lift-off velocity, the reduced traction will prevent it from achieving speeds as high as the curved lift-off velocity. The ridge lift-off velocity

therefore provides a much more relevant and applicable speed limit. The geometry of the set-up used to compute this velocity is illustrated in Fig. 3.4, where the object is seen rolling on an edge between two flat faces.

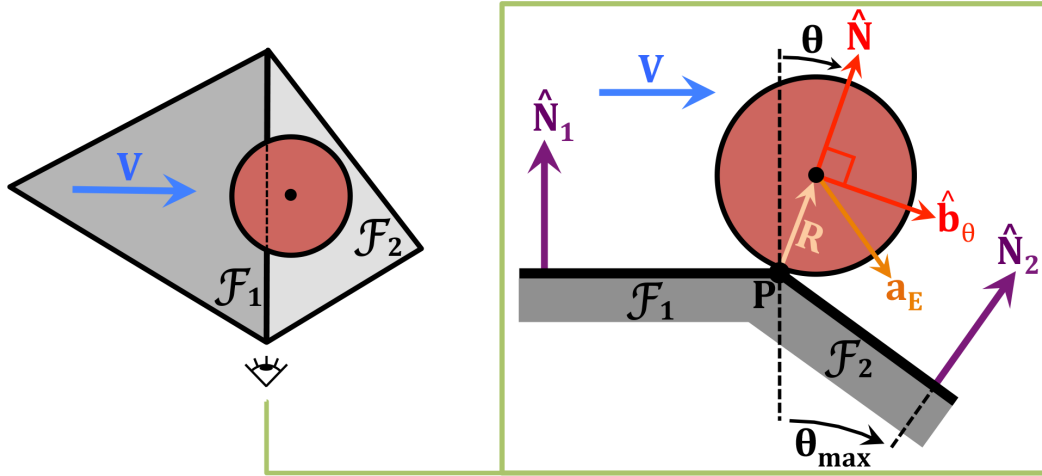


Figure 3.4: Geometry of spherical object rolling on the edge between two facets.

In order to express this ridge lift-off velocity, we return to Eq. 3.15 and remark that the surface radius of curvature $\rho = 0$, such that $\rho^* = R$. Correspondingly, the $\hat{\mathbf{b}}_\rho^*$ axis is equal to the normal $\hat{\mathbf{N}}$. Meanwhile, it still holds that $\mathbf{P} + \mathbf{R} \simeq \mathbf{P}$. This reduces Eq. 3.15 to:

$$V_R = \pm \sqrt{R^2 (\hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega})^2 + R \hat{\mathbf{N}} \cdot ([\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}]\mathbf{P} - \mathbf{a}_E) - R \hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega}} \quad (3.23)$$

Although Eq. 3.23 has two solutions, it is only the $+\sqrt{(\dots)}$ solution that is relevant, as it corresponds to lift-off in the positive $\hat{\mathbf{b}}_\theta$ direction. From Fig. 3.4, it is easy to see why the $-\sqrt{(\dots)}$ solution is dismissed: this case corresponds to lift-off when the object is rolling from facet \mathcal{F}_2 onto facet \mathcal{F}_1 , while the relevant problem is lift-off with the object rolling from facet \mathcal{F}_1 onto facet \mathcal{F}_2 . Furthermore, it can be seen from Fig. 3.4 that the normal vector $\hat{\mathbf{N}}$ changes with the angle θ , as the object proceeds to ‘roll’ on the edge. It is mostly the extremum values that are encountered at $\theta = 0$ and $\theta = \theta_{max}$ that are interesting, in which case correspondingly $\hat{\mathbf{N}} = \hat{\mathbf{N}}_1$ and $\hat{\mathbf{N}} = \hat{\mathbf{N}}_2$. The former case, in which the object exits the body surface upon encountering the edge between the two facets, is independent of the angle between the two facets.

Comparing the two expressions Although Eqs. 3.22 and 3.23 are highly similar in form, they express distinctly different types of lift-off; this difference results from the curvature metrics that are applied in the two expressions. The *curved* lift-off expression as given by Eq. 3.22 yields the velocity required for an object to track or lift off from the *averaged* body surface, through its use of the radius of curvature. In contrast, the *ridge* lift-off expression as given by Eq. 3.23 provides the velocity at which an object moving off a small convexity or protrusion will experience a brief ‘air time.’ The latter is a much more localized definition, for which the object is expected to re-impact the body surface shortly after lift-off. As an analogue on Earth, consider a car driving on a road with a speed bump: the ridge lift-off provides the velocity the car needs to just barely fly off the bump (and impact the road again shortly after), on the order of 10 m/s. In contrast, the curved lift-off velocity yields the speed required for the car to track the global shape of the Earth, approximately 7,900 m/s.

The simple example given above suggests that the curved and ridge lift-off velocities are of distinctly different magnitude. It can be shown that the respective lift-off velocity expressions are dominated by the $\sqrt{\rho \hat{\mathbf{b}}_\rho \cdot \mathbf{a}_E}$ and $\sqrt{R \hat{\mathbf{N}} \cdot \mathbf{a}_E}$ terms. Although the $\hat{\mathbf{N}}$ and $\hat{\mathbf{b}}_\rho$ vectors are generally not aligned, the angle between the two can be assumed small. To first order, the ratio between the two lift-off velocities can therefore be estimated as:

$$\frac{V_R}{V_\theta} = \sqrt{\frac{R}{\rho}} \quad (3.24)$$

Given that the R/ρ ratio is small for realistic scenarios, the lift-off velocity ratio must also be small, *i.e.*, the two lift-off velocities are of very different magnitude. Again, it holds that the curved lift-off velocity is an *averaged* velocity that is mostly applicable to smooth surfaces that are unlikely to exist on real bodies. For those, the ridge lift-off velocity is far more relevant. In fact, the phenomenon of ridge liftoff is precisely what limits the practical applicability of the curved lift-off velocity.

Finally, the earlier statement that both expressions have a dominant term is repeated; this dominance is particularly present in the ridge lift-off expression. As an example, consider Fig. 3.5,

which compares the lift-off velocity on an arbitrary location on asteroid Itokawa, computed with both the full expression as given by Eq. 3.23 and with the reduced form:

$$V_R = \sqrt{-R\hat{\mathbf{N}} \cdot \mathbf{a}_E} \quad (3.25)$$




Figure 3.5: Comparison of full and reduced ridge lift-off velocity expressions on Itokawa.

The two expressions are indeed observed to yield highly similar results, in particular for the radii representative of small-body mission hardware ($R < 1$ m). For such objects, it holds per Eq. 3.25 that the V_R/\sqrt{R} ratio is approximately constant, allowing for pseudo-normalization of the ridge lift-off velocity and for plots that are independent of the object radius R . This reduced expression will therefore be used in the subsequent results. In the curved lift-off velocity expression given by Eq. 3.22, the dependence on the azimuthal direction appears in the form of $\hat{\mathbf{b}}_\theta$. As this vector does not appear in the reduced ridge lift-off velocity expression of Eq. 3.25, the reduced velocity is independent of the azimuthal direction η . This expression also does *not* contain the \pm that biased the curved lift-off velocity relative to the local surface rotational velocity. Because of this lack of direction-dependent terms, the reduced ridge lift-off velocity expression does not display any extrema, and is independent of the azimuthal direction η . Finally, it is noted that on a vertex-facet polyhedron model, each facet with normal $\hat{\mathbf{N}}_1$ has three adjacent facets with normal $\hat{\mathbf{N}}_2$. However, the normals of adjacent facets are generally closely aligned, such that it is sufficient to simply compute the lift-off velocity with the central facet normal $\hat{\mathbf{N}}_1$.

3.1.5 Lift-off from a plane

In our derivations for the lift-off velocity from an arbitrary surface, it was implicitly assumed that the radius of curvature ρ of the local surface was some finite quantity. In the limit, the case where $\rho \rightarrow \infty$ is also relevant. This corresponds to a geometry in which the body surface is locally represented by a plane that rotates at some rate $\boldsymbol{\Omega}$. This setup is illustrated in Fig. 3.6, in which the ‘surface’ plane from which the object lifts off is shown in gray and the departure plane to which this motion is initially contained is shown in red.




Figure 3.6: Geometry for the lift-off from a rotating plane.

Reduced expression The normal acceleration experienced by an object on an arbitrary surface was previously derived as:

$$\mathbf{a}_N = \dot{V}_\theta \hat{\mathbf{b}}_\theta - \frac{V_\theta^2}{\rho^{*2}} \boldsymbol{\rho}^* + 2V_\theta [\tilde{\boldsymbol{\Omega}}] \hat{\mathbf{b}}_\theta + [\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}] (\mathbf{P} + \mathbf{R}) - \mathbf{a}_E \quad (3.26)$$

which is valid for a surface with radius of curvature ρ . As discussed above, in the limit case where the particle is moving on a plane, this curvature $\rho \rightarrow \infty$. Substituting this into Eq. 3.26, it can be seen that the V_θ^2 term disappears, reducing the normal force expression to:

$$\mathbf{a}_N = \dot{V}_\theta \hat{\mathbf{b}}_\theta + 2V_\theta [\tilde{\boldsymbol{\Omega}}] \hat{\mathbf{b}}_\theta + [\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}] \mathbf{P} - \mathbf{a}_E \quad (3.27)$$

in which it was again used that $\mathbf{P} + \mathbf{R} \simeq \mathbf{P}$. For the object to lift off from the plane, it must hold that $\hat{\mathbf{N}} \cdot \mathbf{a}_N \leq 0$, in which $\hat{\mathbf{N}}$ is the normal to the plane that replaces the $\hat{\mathbf{b}}_\rho$ vector in the \mathcal{B} -frame

and is normal to both $\hat{\mathbf{b}}_\theta$ and $\hat{\mathbf{b}}_\delta$. Taking the dot product of $\hat{\mathbf{N}}$ onto Eq. 3.27 and solving for V_θ , yields:

$$V_{\theta,\infty} = \frac{\hat{\mathbf{N}} \cdot [\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}]\mathbf{P} - \hat{\mathbf{N}} \cdot \mathbf{a}_E}{2\hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega}} \quad (3.28)$$

Properties of the limit solution Contrary to our previous results on a curved surface, Eq. 3.28 has only a single solution and states that lift-off from a rotating plane is only possible in the $+\hat{\mathbf{b}}_\theta$ direction, *i.e.*, when moving along the direction of the surface rotational velocity. If $\hat{\mathbf{b}}_\theta$ is chosen such that it points against the surface rotation, Eq. 3.28 will simply yield negative V_θ .

When the body rotation does not have a component normal to the departure plane, or when the body is entirely non-rotating, it holds that $2\hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega} = 0$, such that lift-off is impossible in both directions, as $V_\theta \rightarrow \infty$. This result can be explained through inspection of the general equation of motion of a particle in a rotating frame fixed to the target body at its center of mass:

$$\ddot{\mathbf{r}} = \mathbf{a}_E - [\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}]\mathbf{r} - 2[\tilde{\boldsymbol{\Omega}}]\dot{\mathbf{r}} \quad (3.29)$$

in which $-[\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}]\mathbf{r}$ and $-2[\tilde{\boldsymbol{\Omega}}]\dot{\mathbf{r}}$ are respectively the *centrifugal* and *Coriolis* accelerations that result from the use of a rotating, non-inertial reference frame. It is this Coriolis acceleration, the only term dependent on the particle velocity $\dot{\mathbf{r}}$, that is responsible for the planar lift-off behavior. The properties of this Coriolis acceleration are investigated in the simple geometry of a particle moving along a plane attached to the equator of a rotating sphere, as illustrated on the right side of Fig. 3.7. When the particle velocity $\dot{\mathbf{r}}$ is directed along the surface rotational velocity $\mathbf{v}_r = \boldsymbol{\Omega} \times \mathbf{r}$, the resulting Coriolis acceleration has a normal component directed *upwards* from the plane. In contrast, when $\dot{\mathbf{r}}$ is directed against \mathbf{v}_r , the Coriolis acceleration points *into* the plane. Additionally, the other two accelerations terms always yield a combined acceleration *into* the plane.

When moving in the direction of \mathbf{v}_r , there will exist some $\dot{\mathbf{r}}$ at which the Coriolis acceleration cancels out the other acceleration terms, yielding a net acceleration that pulls the object off the plane. However, when moving against \mathbf{v}_r , all acceleration components are directed *into* the plane, such that the object cannot leave the plane, no matter the magnitude of its velocity $\dot{\mathbf{r}}$. This phenomenon is illustrated on the left side of Fig. 3.7, in which the particle shown in red moves




Figure 3.7: Illustration of a particle lifting off from a rotating plane in (left) 2D and (right) 3D, where the green arrow indicates low lift-off velocity, the orange arrows indicate intermediate lift-off velocity, and the red arrows indicate infinite lift-off velocity.

along an infinite plane attached to a rotating disk.

Due to the disk rotation, the left end of the plane moves downward, while the right end moves upward. As the gravitational attraction of the disk points towards its center, the particle trajectory must curve down regardless of the direction and magnitude of its velocity, as shown by the dashed lines. When the particle moves with sufficient velocity in the $V_{\theta,1}$ direction, the plane will dip down faster than the particle trajectory does, such that the particle locally lifts off from the plane. However, when the particle moves in the $V_{\theta,2}$ direction, the plane moves up while the particle is trying to move downwards; lift-off is therefore impossible in this direction. Finally, when the disk is non-rotating, the plane is stationary and all trajectories will curve down into the plane, making lift-off impossible in *all* directions.

This simple, two-dimensional example can be extended to the full, three-dimensional case, in which lift-off is only possible if the angle between the desired lift-off direction $\hat{\mathbf{b}}_{\theta}$ and the local surface rotational velocity \mathbf{v}_r is smaller than 90° . In this case, V_{θ} has a finite value. As the angle is increased towards 90° , the body rotation component $\Omega_{\delta} \rightarrow 0$, such that $V_{\theta} \rightarrow \infty$. At precisely 90° , lift-off is again impossible in both directions, as in that case $\Omega_{\delta} = 0$. This is also illustrated

on the right side of Fig. 3.7. The familiar lift-off velocity curves were computed on the surface of an ellipsoid, for both the planar and curved cases. The ratio between the two is shown in Fig. 3.8.




Figure 3.8: Ratio of planar-to-curved lift-off velocities at an arbitrary position on a rotating ellipsoid.

This clearly shows the planar lift-off velocity growth towards infinity as the lift-off direction approaches $\eta = 90^\circ$, parallel to the rotation axis of the plane. The curves also illustrate that planar lift-off is indeed only possible when the particle departs along the local surface rotational velocity, when $-90^\circ < \eta < 90^\circ$. This is visible from the fact that the lift-off velocity ratio has a negative magnitude when $90^\circ < \eta < 270^\circ$. These numerical results confirm our analysis. The planar lift-off velocities are observed to have much greater values than their curved counterparts, even in areas where the planar velocity is well-behaved. This makes the planar lift-off velocity poorly applicable to the flat facets of a polyhedron model, instead, one must account for effective curvature that these models express, as will be discussed in Section 3.2.

Simulations To further verify the above analysis, a simple simulation of a particle moving along a plane rotating with $\Omega = 1.5 \times 10^{-4}$ rad/s, located 100 m above a pointmass with $\mu = 2.5$ m³/s², was performed. Fig. 3.9 shows the resulting motion for particles departing along with and against the rotational velocity of the plane. It can be seen that the trajectories of particles moving against the rotation indeed move into the plane and fail to achieve lift-off.




Figure 3.9: Simulated lift-off of particles on a rotating plane.

Inspecting the trajectories of the particles moving along with the rotation, it can be seen that the particle with $0.9 \times V_\theta$ moves into the plane as well; the dynamics are such that this trajectory eventually does leave the plane, though this is not physically possible due to the presence of the plane. The trajectory with $1.0 \times V_\theta$ only just achieves lift-off; the trajectory with $1.1 \times V_\theta$ lifts off from the plane at a faster rate. This behavior matches our analysis of the planar lift-off velocity expression.

3.1.6 Lift-off from a concave region

One of the major assumptions made in the preceding text was that the surface from which the particle lifts off is locally convex. If it is instead locally concave, the condition for lift-off given in Section 3.1 changes to $\hat{\mathbf{b}}_\rho \cdot \mathbf{a}_N \geq 0$. The limit velocity at which lift-off is just possible is found by again setting this inequality to zero, which yields the same result previously given for lift-off from

a convex surface:

$$V_\theta = \pm \sqrt{\rho^2 (\hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega})^2 + \rho \hat{\mathbf{b}}_\rho \cdot ([\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}]\mathbf{P} - \mathbf{a}_E)} - \rho \hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega} \quad (3.30)$$

in which, contrary to the convex case, the $\rho \hat{\mathbf{b}}_\rho \cdot ([\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}]\mathbf{P} - \mathbf{a}_E)$ term must be negative due to the requirement that a stationary particle may not already lift off due to the external acceleration and rotation of the target body. For the radical of Eq. 3.30 to be real, the $\rho^2 (\hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega})^2$ term, which is always positive, must therefore be sufficiently large. This will always be possible for a sufficiently large radius of curvature ρ . However, as ρ is decreased, there will be some limit at which the quantity in the radical is equal to zero. This is the minimum concave radius of curvature at which lift-off is just barely possible; if ρ becomes any smaller, the radical no longer has a real solution and lift-off is impossible. This minimum radius ρ_{min}^{conc} can be found by setting the radicand equal to zero:

$$\rho_{min}^{conc} = \frac{\hat{\mathbf{b}}_\rho \cdot (\mathbf{a}_E - [\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}]\mathbf{P})}{(\hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega})^2} \quad (3.31)$$

This expression is dominated by the Ω_δ^2 term in the denominator; as Ω_δ is on the order of 10^{-4} for most small bodies in the Solar system, the resulting ρ_{min}^{conc} will have a very large value. This is consistent with the discussion in the previous paragraph. The corresponding lift-off velocity can also be found from Eq. 3.30, by setting the radicand to zero and substituting for the minimum concave radius of curvature, as:

$$V_\theta^{conc} = -\rho \hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega} = \frac{\hat{\mathbf{b}}_\rho \cdot ([\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}]\mathbf{P} - \mathbf{a}_E)}{\hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega}} \quad (3.32)$$

The result that lift-off from a concave surface is only possible for very large ρ values intuitively follows from the discussion on planar lift-off situations. In particular, consider the planar lift-off simulation shown in Fig. 3.9. Comparing the scales of the x - and y -axes, it can be seen that although a particle can indeed lift off from a plane, it does so at an extremely shallow angle, despite the already large velocity given to the particle. Even if the particle velocity is increased further, there exists a limit to the upward curvature of the particle trajectory. In fact, if the velocity is increased much beyond the red trajectory, the resulting trajectory actually becomes flatter again. This is

consistent with the previous discussion, and indicates that lift-off from a concave surface is only possible for very limited geometries in which the radius of curvature ρ is sufficiently large.

The spheres, ellipsoids, and planes studied in this work are fully convex, such that there is no ambiguity on whether lift-off is possible. The small bodies of our Solar System have a more complex shape with locally convex or locally concave regions; some form of numerical averaging will be required to obtain the local curvature of the discrete models available for these bodies. In agreement with earlier results, lift-off is expected to be possible in all directions η on a locally convex region, and is expected to be (mostly) impossible on a locally concave region. Additionally, it is also possible for a region to be locally convex in some η directions, and locally concave in others, when the region is a local saddle point. In this case, lift-off is expected to be possible in the directions with convex curvature, and (mostly) impossible in the directions with concave curvature. This is illustrated in Fig. 3.10, which shows the variation of the radius of curvature and corresponding lift-off velocities on an arbitrary convex region, and an arbitrary convex-concave saddle point region. The fully concave case is not included, as in this case lift-off is generally impossible.




Figure 3.10: Radius of curvature and lift-off velocity variation for a (left) convex and (right) convex-concave region.

3.2 Surface curvature

The evaluation of the lift-off velocity on a curved surface requires knowledge of that surface's radius of curvature. This section reviews some general properties of surface curvature and applies them to develop a methodology to compute the curvature of ellipsoids and of spherical harmonics shape models. It also summarizes an existing technique for the computation of polyhedron surface curvature, which is well documented in literature.

3.2.1 General properties

Consider the warped cylinder shown on the left side of Fig. 3.11. At each point \mathbf{P} where the surface is continuous, there exists one surface normal $\hat{\mathbf{N}}$ which defines the plane tangent to the surface through \mathbf{P} . This tangent plane is spanned by two orthonormal vectors $\hat{\mathbf{e}}^+$ and $\hat{\mathbf{e}}^-$, which can be used to express some unit vector $\hat{\mathbf{e}}(\eta) = \cos(\eta) \cdot \hat{\mathbf{e}}^+ + \sin(\eta) \cdot \hat{\mathbf{e}}^-$ that lies within the tangential plane.




Figure 3.11: Illustration of (*left*) general properties of curvature and (*right*) Meusnier's theorem, on a warped cylinder.

Combined with the surface normal $\hat{\mathbf{N}}$, any $\hat{\mathbf{e}}(\eta)$ vector spans a plane that intersects the surface along some arbitrary curve. This intersection curve can be approximated in the neighborhood of \mathbf{P}

with an osculating circle that has an associated *curvature* $\kappa(\eta)$. This curvature varies as a function of the azimuthal angle η and achieves some maximum κ^+ and minimum κ^- , known as the *principal curvatures*. The *principal directions* in which these extrema occur are denoted $\hat{\mathbf{e}}^+$ and $\hat{\mathbf{e}}^-$ and must always be perpendicular [111, p. 80]. As such, they may be used as orthonormal basis for the tangent plane. Using the principal curvatures, it is also possible to define the *mean curvature* H and *Gaussian curvature* K at \mathbf{P} [111, p. 83]:

$$\begin{aligned} H &= \frac{1}{2}\kappa^+ + \frac{1}{2}\kappa^- \\ K &= \kappa^+ \cdot \kappa^- \end{aligned} \quad (3.33)$$

The reverse transformation is found as:

$$\kappa^\pm = H \pm \sqrt{H^2 - K} \quad (3.34)$$

The curvature in some direction η measured from $\hat{\mathbf{e}}^+$ can be computed from the principal curvatures using *Euler's theorem*, as [111, p. 81]:

$$\kappa(\eta) = \kappa^+ \cdot \cos^2(\eta) + \kappa^- \cdot \sin^2(\eta) \quad (3.35)$$

All curvatures mentioned so far are normal curvatures, *i.e.*, they are curvatures in a plane containing the surface normal $\hat{\mathbf{N}}$. However, the departure plane within which the curvature is computed for lift-off velocity purposes generally *does not* contain the surface normal. As such, the surface curvature in some plane offset from the normal is needed. This is illustrated on the right side of Fig. 3.11, in which the green plane and circle match the normal curvature in direction η , and the red plane and circle match the corresponding curvature of some departure plane offset at an angle δ from the local surface normal. In order to compute this curvature, *Meusnier's theorem* is applied [111, p. 75]:

$$\kappa(\eta, \delta) = \frac{\kappa(\eta)}{\cos(\delta)} \quad (3.36)$$

Finally, the curvature κ and radius of curvature ρ are simply each other's inverse, such that the conversion between the two can be achieved as:

$$\rho(\eta, \delta) = \frac{1}{\kappa(\eta, \delta)} \quad (3.37)$$

This has effectively reduced the task of computing the radius of curvature ρ to that of computing the mean and Gaussian curvatures, H and K . A parametric technique for doing so will be discussed shortly; we first briefly comment on some further properties of surface curvature.

Given that both principal curvatures κ^+ and κ^- can be independently positive or negative, three cases are possible: a surface can locally be fully convex (with $\kappa^+ > 0$ and $\kappa^- > 0$), fully concave (with $\kappa^+ < 0$ and $\kappa^- < 0$), or convex in certain directions and concave in others (with $\kappa^+ > 0$ and $\kappa^- < 0$). It was mentioned before that lift-off is generally only possible when $\kappa(\eta) > 0$. As a result, lift-off can be possible in all η directions, impossible in all η directions, or possible in some η directions and impossible in others. In the latter case, there will exist some finite minimum lift-off velocity but no maximum velocity, as the lift-off velocity grows to infinity when $\kappa(\eta)$ transitions from positive to negative values. Mathematically, this phenomenon occurs because the radicand in Eq. 3.22 is negative. On arbitrary bodies, certain regions where lift-off is partially or fully (with respect to η) impossible are therefore expected.

As mentioned, computation of the radius of curvature across some body can be reduced to computation of that body's mean and Gaussian curvatures. If the body surface is represented parametrically as some $\mathbf{r}(\phi, \lambda)$, then H and K can be obtained from the two *fundamental forms* of the surface. These two-dimensional matrices are a function of the partial derivatives of the parametric surface function $\mathbf{r}(\phi, \lambda)$ with respect to the two coordinates ϕ (latitude) and λ (longitude), as [111, p. 58]:

$$\mathbf{I} = \begin{bmatrix} E & F \\ F & G \end{bmatrix} = \begin{bmatrix} \mathbf{r}_\phi \cdot \mathbf{r}_\phi & \mathbf{r}_\phi \cdot \mathbf{r}_\lambda \\ \mathbf{r}_\lambda \cdot \mathbf{r}_\phi & \mathbf{r}_\lambda \cdot \mathbf{r}_\lambda \end{bmatrix} \quad (3.38)$$

and:

$$\mathbf{II} = \begin{bmatrix} L & M \\ M & N \end{bmatrix} = \begin{bmatrix} \mathbf{r}_{\phi\phi} \cdot \hat{\mathbf{N}} & \mathbf{r}_{\phi\lambda} \cdot \hat{\mathbf{N}} \\ \mathbf{r}_{\phi\lambda} \cdot \hat{\mathbf{N}} & \mathbf{r}_{\lambda\lambda} \cdot \hat{\mathbf{N}} \end{bmatrix} \quad (3.39)$$

in which the partial derivatives $\mathbf{r}_\phi = \partial\mathbf{r}/\partial\phi$, $\mathbf{r}_{\phi\lambda} = \partial^2\mathbf{r}/\partial\phi\partial\lambda$, etcetera. Furthermore, $\hat{\mathbf{N}}$ is the normal to the surface, which can be computed from the partial derivatives of $\mathbf{r}(\phi, \lambda)$ as [111, p. 62]:

$$\hat{\mathbf{N}} = \frac{\mathbf{r}_\phi \times \mathbf{r}_\lambda}{\|\mathbf{r}_\phi \times \mathbf{r}_\lambda\|} = \frac{\mathbf{r}_\phi \times \mathbf{r}_\lambda}{\sqrt{\mathbf{r}_\phi \cdot \mathbf{r}_\phi + \mathbf{r}_\lambda \cdot \mathbf{r}_\lambda + (\mathbf{r}_\phi \cdot \mathbf{r}_\lambda)^2}} \quad (3.40)$$

The Gaussian and mean curvatures of the surface can be computed from the coefficients of the first and second fundamental form as [111, p. 83]:

$$\begin{aligned} K &= \frac{LN - M^2}{EG - F^2} \\ H &= \frac{LG - 2NF + ME}{2(EG - F^2)} \end{aligned} \quad (3.41)$$

Although the principal curvatures κ^\pm can be established by combining Eqs. 3.41 and 3.34, the respective directions of curvatures are also required. For this, the s matrix [111, p. 80] is used:

$$s = [\mathbf{I}]^{-1}[\mathbf{II}] = \frac{1}{F^2 - EG} \begin{bmatrix} FM - GL & FN - GM \\ FL - EM & FM - EN \end{bmatrix} \quad (3.42)$$

The eigenvalues of this matrix are κ^+ and κ^- ; the corresponding two-dimensional eigenvectors $\hat{\mathbf{t}}^+ = [t_1^+, t_2^+]$ and $\hat{\mathbf{t}}^- = [t_1^-, t_2^-]$ are the principal directions expressed in the $(\mathbf{r}_\phi, \mathbf{r}_\lambda)$ basis; they can be converted into the Cartesian (x, y, z) basis to yield the $\hat{\mathbf{e}}^\pm$ directions as:

$$\hat{\mathbf{e}}^\pm = t_1^\pm \cdot \hat{\mathbf{r}}_\phi + t_2^\pm \cdot \hat{\mathbf{r}}_\lambda \quad (3.43)$$

This finally allows us to compute the radius of curvature within some arbitrary plane (spanned by the external acceleration \mathbf{a}_E and some chosen lift-off direction $\hat{\mathbf{b}}_\theta$ at some departure point \mathbf{P}) and to obtain the corresponding lift-off velocity, assuming that some parametric $(\mathbf{r}_\phi, \mathbf{r}_\lambda)$ representation of the body surface is available.

3.2.2 Ellipsoid

The curvature computation methodology outlined above is easily applied to tri-axial ellipsoids. The surface of such a shape can be parameterized as:

$$\mathbf{r}(\phi, \lambda) = \begin{bmatrix} a \cos \phi \cos \lambda \\ b \cos \phi \sin \lambda \\ c \sin \phi \end{bmatrix} \quad (3.44)$$

in which a , b , and c are the ellipsoid semi-axes, and ϕ and λ are the latitude and longitude of some surface point. Computation of the mean and Gaussian curvatures requires the first-order partial derivatives of this radius equation, which are found as:

$$\mathbf{r}_\phi = \frac{\partial \mathbf{r}}{\partial \phi} = \begin{bmatrix} -a \sin \phi \cos \lambda \\ -b \sin \phi \sin \lambda \\ c \cos \phi \end{bmatrix} \quad \text{and} \quad \mathbf{r}_\lambda = \frac{\partial \mathbf{r}}{\partial \lambda} = \begin{bmatrix} -a \cos \phi \sin \lambda \\ b \cos \phi \cos \lambda \\ 0 \end{bmatrix} \quad (3.45)$$

Similarly, the second-order partials are found as:

$$\mathbf{r}_{\phi\phi} = \frac{\partial^2 \mathbf{r}}{\partial \phi \partial \phi} = \begin{bmatrix} -a \cos \phi \cos \lambda \\ -b \cos \phi \sin \lambda \\ -c \sin \phi \end{bmatrix} \quad \text{and} \quad \mathbf{r}_{\lambda\lambda} = \frac{\partial^2 \mathbf{r}}{\partial \lambda \partial \lambda} = \begin{bmatrix} -a \cos \phi \cos \lambda \\ -b \cos \phi \sin \lambda \\ 0 \end{bmatrix} \quad (3.46)$$

$$\mathbf{r}_{\phi\lambda} = \frac{\partial^2 \mathbf{r}}{\partial \phi \partial \lambda} = \begin{bmatrix} a \sin \phi \sin \lambda \\ -b \sin \phi \cos \lambda \\ 0 \end{bmatrix} \quad (3.47)$$

These expressions can be substituted into Eq. 3.40 to yield the surface normal and into Eqs. 3.38 and 3.42 to obtain the surface curvatures on a tri-axial ellipsoid. This method will be used in the analysis of the lift-off velocity distribution across the surfaces of various ellipsoids, in Section 3.3.

3.2.3 Spherical harmonics

The spherical harmonics shape model was introduced in Section 2.2. Recall that this model parameterizes the shape surface with a harmonic series as:

$$\mathbf{r}(\phi, \lambda) = r(\phi, \lambda) \cdot \begin{bmatrix} \cos \phi \cos \lambda \\ \cos \phi \sin \lambda \\ \sin \phi \end{bmatrix} \quad (3.48)$$

in which the radius $r(\phi, \lambda)$ is computed as:

$$r(\phi, \lambda) = \sum_{j=0}^J \sum_{k=0}^j (A_{jk} \cdot \cos(k\lambda) + B_{jk} \cdot \sin(k\lambda)) \cdot P_{jk}(\sin \phi) \quad (3.49)$$

We repeat that A_{jk} and B_{jk} are the degree j and order k spherical harmonic coefficients, while J is the maximum degree of the considered shape model. Furthermore, P_{jk} is the degree j , order k associated Legendre function. This function is most efficiently computed recursively, as [128, p. 593]:

$$\begin{cases} P_{n,0} = \frac{1}{n} ((2n-1)xP_{n-1,0} - (n-1)P_{n-2,0}) & n \geq 2 \\ P_{n,m} = P_{n-2,m} + (2n-1)\sqrt{1-x^2}P_{n-1,m-1} & m \neq 0, m < n \\ P_{n,n} = (2n-1)\sqrt{1-x^2}P_{n-1,m-1} & n \neq 0 \end{cases} \quad (3.50)$$

in which the first three terms are equal to [128, p. 593]:

$$P_{0,0}(x) = 1 \quad \text{and} \quad P_{1,0}(x) = x \quad \text{and} \quad P_{1,1} = \sqrt{1-x^2} \quad (3.51)$$

This spherical harmonic expansion is a *continuous* surface representation: it can be evaluated at any (ϕ, λ) and thus produce a shape of arbitrary resolution (even though the degree J of the model is limited by numerics). In analogy to the ellipsoid, the surface curvature on a spherical harmonics shape can be computed by substituting the partial derivatives of the parametric $\mathbf{r}(\phi, \lambda)$ into Eqs. 3.38 and 3.42. Taking the partial derivatives of Eq. 3.48, it is found that:

$$\mathbf{r}_\phi = \begin{bmatrix} \frac{\partial r}{\partial \phi} c\phi c\lambda - rs\phi c\lambda \\ \frac{\partial r}{\partial \phi} c\phi s\lambda - rs\phi s\lambda \\ \frac{\partial r}{\partial \phi} s\phi + rc\phi \end{bmatrix} \quad \text{and} \quad \mathbf{r}_\lambda = \begin{bmatrix} \frac{\partial r}{\partial \lambda} c\phi c\lambda - rc\phi s\lambda \\ \frac{\partial r}{\partial \lambda} c\phi s\lambda + rc\phi c\lambda \\ \frac{\partial r}{\partial \lambda} s\phi \end{bmatrix} \quad (3.52)$$

$$\mathbf{r}_{\phi\phi} = \begin{bmatrix} c\phi c\lambda \left(\frac{\partial^2 r}{\partial \lambda^2} - r \right) - 2 \frac{\partial r}{\partial \lambda} c\phi s\lambda \\ c\phi s\lambda \left(\frac{\partial^2 r}{\partial \lambda^2} - r \right) + 2 \frac{\partial r}{\partial \lambda} c\phi c\lambda \\ \frac{\partial^2 r}{\partial \lambda^2} s\phi \end{bmatrix} \quad \text{and} \quad \mathbf{r}_{\lambda\lambda} = \begin{bmatrix} c\phi c\lambda \left(\frac{\partial^2 r}{\partial \lambda^2} - r \right) - 2 \frac{\partial r}{\partial \lambda} c\phi s\lambda \\ c\phi s\lambda \left(\frac{\partial^2 r}{\partial \lambda^2} - r \right) + 2 \frac{\partial r}{\partial \lambda} c\phi c\lambda \\ \frac{\partial^2 r}{\partial \lambda^2} s\phi \end{bmatrix} \quad (3.53)$$

$$\mathbf{r}_{\phi\lambda} = \begin{bmatrix} \frac{\partial^2 r}{\partial \phi \partial \lambda} c\phi c\lambda - \frac{\partial r}{\partial \phi} c\phi s\lambda - \frac{\partial r}{\partial \lambda} s\phi c\lambda + rs\phi s\lambda \\ \frac{\partial^2 r}{\partial \phi \partial \lambda} c\phi s\lambda + \frac{\partial r}{\partial \phi} c\phi c\lambda - \frac{\partial r}{\partial \lambda} s\phi s\lambda - rs\phi c\lambda \\ \frac{\partial^2 r}{\partial \phi \partial \lambda} s\phi + \frac{\partial r}{\partial \lambda} c\phi \end{bmatrix} \quad (3.54)$$

in which $c\phi = \cos \phi$ and $s\phi = \sin \phi$. The radius $r = r(\phi, \lambda)$ can be computed by Eq. 3.49; its partial derivatives are found as:

$$\left\{ \begin{array}{l} \frac{\partial r}{\partial \phi} = \sum_{j=0}^J \sum_{k=0}^j [A_{jk} \cdot c(k\lambda) + B_{jk} \cdot s(k\lambda)] \cdot \frac{\partial P_{jk}(s\phi)}{\partial s\phi} \cdot c\phi \\ \frac{\partial r}{\partial \lambda} = \sum_{j=0}^J \sum_{k=0}^j [-kA_{jk} \cdot s(k\lambda) + kB_{jk} \cdot c(k\lambda)] \cdot P_{jk}(s\phi) \\ \frac{\partial^2 r}{\partial \phi^2} = \sum_{j=0}^J \sum_{k=0}^j [A_{jk} \cdot c(k\lambda) + B_{jk} \cdot s(k\lambda)] \cdot \left[\frac{\partial^2 P_{jk}(s\phi)}{\partial s^2 \phi} \cdot c^2 \phi - \frac{\partial P_{jk}(s\phi)}{\partial s\phi} \cdot s\phi \right] \\ \frac{\partial^2 r}{\partial \lambda^2} = \sum_{j=0}^J \sum_{k=0}^j [-k^2 A_{jk} \cdot c(k\lambda) - k^2 B_{jk} \cdot s(k\lambda)] \cdot P_{jk}(s\phi) \\ \frac{\partial^2 r}{\partial \phi \partial \lambda} = \sum_{j=0}^J \sum_{k=0}^j [-kA_{jk} \cdot s(k\lambda) + kB_{jk} \cdot c(k\lambda)] \cdot \frac{\partial P_{jk}(s\phi)}{\partial s\phi} \cdot c\phi \end{array} \right. \quad (3.55)$$

Finally, the evaluation of these r partial requires the partial derivatives of the Legendre polynomials.

From Eqs. 3.50 and 3.51, the first derivatives may be found as:

$$\begin{array}{l} \frac{\partial P_{0,0}}{\partial x} = 0 \quad \text{and} \quad \frac{\partial P_{1,0}}{\partial x} = 1 \quad \text{and} \quad \frac{\partial P_{1,1}}{\partial x} = \frac{-x}{\sqrt{1-x^2}} \\ \left\{ \begin{array}{l} \frac{\partial P_{n,0}}{\partial x} = \frac{1}{n} \left((2n-1)P_{n-1,0} + (2n-1)x \frac{\partial P_{n-1,0}}{\partial x} - (n-1) \frac{\partial P_{n-2,0}}{\partial x} \right) \quad \text{if } n \geq 2 \\ \frac{\partial P_{n,m}}{\partial x} = \frac{\partial P_{n-2,m}}{\partial x} - (2n-1) \frac{x}{\sqrt{1-x^2}} P_{n-1,n-1} \\ \quad + (2n-1) \sqrt{1-x^2} \frac{\partial P_{n-1,m-1}}{\partial x} \quad \text{if } m \neq 0, m < n \\ \frac{\partial P_{n,n}}{\partial x} = -(2n-1) \frac{x}{\sqrt{1-x^2}} P_{n-1,n-1} + (2n-1) \sqrt{1-x^2} \frac{\partial P_{n-1,n-1}}{\partial x} \quad \text{if } n \neq 0 \end{array} \right. \end{array} \quad (3.56)$$

Similarly, the second derivatives are found as:

$$\begin{array}{l} \frac{\partial^2 P_{0,0}}{\partial x^2} = 0 \quad \text{and} \quad \frac{\partial^2 P_{1,0}}{\partial x^2} = 0 \quad \text{and} \quad \frac{\partial^2 P_{1,1}}{\partial x^2} = \frac{-1}{(1-x^2)^{3/2}} \\ \left\{ \begin{array}{l} \frac{\partial^2 P_{n,0}}{\partial x^2} = \frac{1}{n} \left(2(2n-1) \frac{\partial P_{n-1,0}}{\partial x} + (2n-1)x \frac{\partial^2 P_{n-1,0}}{\partial x^2} - (n-1) \frac{\partial^2 P_{n-2,0}}{\partial x^2} \right) \quad \text{if } n \geq 2 \\ \frac{\partial^2 P_{n,m}}{\partial x^2} = \frac{\partial^2 P_{n-2,m}}{\partial x^2} - (2n-1)(1-x^2)^{-3/2} P_{n-1,m-1} \\ \quad - (2n-1) \frac{x}{\sqrt{1-x^2}} \frac{\partial P_{n-1,m-1}}{\partial x} - (2n-1) \frac{x}{\sqrt{1-x^2}} \frac{\partial P_{n-1,m-1}}{\partial x} \\ \quad + (2n-1) \sqrt{1-x^2} \frac{\partial^2 P_{n-1,m-1}}{\partial x^2} \quad \text{if } m \neq 0, m < n \\ \frac{\partial^2 P_{n,n}}{\partial x^2} = -(2n-1)(1-x^2)^{-3/2} P_{n-1,n-1} - (2n-1) \frac{x}{\sqrt{1-x^2}} \frac{\partial P_{n-1,n-1}}{\partial x} \\ \quad - (2n-1) \frac{x}{\sqrt{1-x^2}} \frac{\partial P_{n-1,n-1}}{\partial x} + (2n-1) \sqrt{1-x^2} \frac{\partial^2 P_{n-1,n-1}}{\partial x^2} \quad \text{if } n \neq 0 \end{array} \right. \end{array} \quad (3.57)$$

This figure shows the *1-ring neighborhood* of vertex \mathbf{P}_i , which consists of all edges \mathcal{E}_j and facets \mathcal{F}_j that contain vertex \mathbf{P}_i . Each of the facets has an associated area A_j (highlighted in green), which is computed differently depending on the shape of the facet [65]:

$$A_j = \begin{cases} \frac{1}{8} \cdot \left(\|\mathbf{b}_j\|^2 \cot \beta_j + \|\mathbf{c}_j\|^2 \cot \gamma_j \right) & \text{if } \mathcal{F}_j \text{ is non-obtuse} \\ \frac{1}{4} \cdot \|\mathbf{b}_j \times \mathbf{c}_j\| & \text{if } \alpha_j > \pi/2 \\ \frac{1}{8} \cdot \|\mathbf{b}_j \times \mathbf{c}_j\| & \text{if } \beta_j > \pi/2 \text{ or } \gamma_j > \pi/2 \end{cases} \quad (3.59)$$

The area A_i associated with vertex \mathbf{P}_i is then simply computed as the sum of the areas A_j of all facets \mathcal{F}_j from the 1-ring neighborhood of \mathbf{P}_i [65]:

$$A_i = \sum_{\forall \mathcal{F}_j} A_j \quad (3.60)$$

Using this area and the angle α_j defined in Fig. 3.12, the Gaussian curvature K_i of vertex \mathbf{P}_i can be computed through a summation over the facets from its 1-ring neighborhood, as [65]:

$$K_i = \frac{1}{A_i} \left(2\pi - \sum_{\forall \mathcal{F}_j} \alpha_j \right) \quad (3.61)$$

Similarly, the mean curvature H_i can be found through a summation over all edges \mathcal{E}_j in the 1-ring neighborhood of \mathbf{P}_i . To obtain H_i , the mean curvature normal \mathbf{H}_i is first computed as [65]:

$$\mathbf{H}_i = \frac{1}{2A_i} \sum_{\forall \mathcal{E}_j} \|\mathbf{P}_j - \mathbf{P}_i\| (\cot \theta_j + \cot \psi_j) \quad (3.62)$$

This mean curvature normal contains information about both the vertex normal $\hat{\mathbf{N}}_i$ and the mean curvature H_i , which are found as [65]:

$$\hat{\mathbf{N}}_i = \frac{\mathbf{H}_i}{\|\mathbf{H}_i\|} \quad \text{and} \quad H_i = \frac{\|\mathbf{H}_i\|}{2} \quad (3.63)$$

We note that this vertex normal $\hat{\mathbf{N}}_i$ is the discrete counterpart of the normal vector computed with spherical harmonics in Eq. 3.40. The principal curvatures are then found from K_i and H_i as [65]:

$$\kappa_i^\pm = H_i \pm \sqrt{\Delta_i} \quad \text{with} \quad \Delta_i = \max(H_i^2 - K_i, 0) \quad (3.64)$$

We note that this expression matches Eq. 3.34 for the principal curvatures. In order to compute the associated principal directions, the symmetric curvature tensor B_i of vertex P_i is first computed, which is found by minimizing the error E_i [65]:

$$E_i = \sum_{\forall \mathcal{E}_j} w_j \left(\hat{\mathbf{d}}_j^T B_i \hat{\mathbf{d}}_j - \kappa_j \right)^2 \quad \text{with} \quad B_i = \begin{bmatrix} a_i & b_i \\ b_i & c_i \end{bmatrix} \quad (3.65)$$

In this expression, $\hat{\mathbf{d}}_j$ is the two-dimensional unit projection of the edge $\mathcal{E}_j = \mathbf{P}_j - \mathbf{P}_i$ onto the tangent plane of \mathbf{P}_i ; the latter is normal to $\hat{\mathbf{N}}_i$. Furthermore, κ_j is the normal curvature of the surface at \mathbf{P}_i in the direction of the edge \mathcal{E}_j , which can be computed as [65]:

$$\kappa_j = \frac{2(\mathbf{P}_i - \mathbf{P}_j) \cdot \hat{\mathbf{N}}_i}{\|\mathbf{P}_i - \mathbf{P}_j\|} \quad (3.66)$$

Finally, the weights w_j are defined as [65]:

$$w_j = \frac{1}{A_i} \left[\frac{1}{8} (\cot \theta_j + \cot \psi_j) \cdot \|\mathbf{P}_i - \mathbf{P}_j\|^2 \right] \quad (3.67)$$

Once the curvature tensor B_i has been determined by this minimization, the principal directions are found as the eigenvectors of B_i . Meyer et al. note that although the eigenvalues of B_i are the corresponding principal curvatures, the numerical accuracy of these eigenvalues is typically lower than that of Eq. 3.64 [65]. Eq. 3.64 is therefore used to compute the principal curvatures and obtain the corresponding directions from B_i .

Comparison with spherical harmonics In theory, the spherical harmonic and polyhedron shape models of a given body should yield matching curvatures. However, this is not necessarily true in practice, where it is found that the possible difference between the two models depends on the resolution of the considered models. This resolution is quantified for the spherical harmonic model by the maximum degree and order J of the shape model; for the polyhedron model it is represented by the total number of vertices in the model. In order to illustrate these differences, the degree and order $J = 17$ model of the Mars moon Phobos is used, as developed by [142]. The corresponding polyhedron model can be constructed by sampling the spherical harmonic model; a comparison of the two models was previously shown in Fig. 2.2.

We have computed the relevant curvatures of Phobos's spherical harmonic model using the previously derived expressions. These computations were then repeated on the corresponding polyhedron model, sampled at various resolutions. To illustrate the convergence of the curvatures of the two models, the difference in Gaussian curvature between the spherical harmonic model and the various polyhedron models is shown in Fig. 3.13. These differences are scaled by the Gaussian curvature of Phobos's reference sphere. It can be seen that the differences shrink as the resolution of the polyhedron model is increased, though a fairly high resolution is required to obtain differences on the order of 1% (from top left to bottom right, the models have a resolution of 642, 2,562, 10,242, and 40,962 vertices). This is an intuitive result: in order to capture information about very local changes on a discretized surface, it must be sampled at a correspondingly fine resolution.




Figure 3.13: Scaled differences in Gaussian curvature between the spherical harmonic and various polyhedron models of Phobos.

In this application to Phobos, the spherical harmonics model was available beforehand; the corresponding polyhedron shape models could be sampled to verify its convergence to the ‘true’ spherical harmonics shape. However, in most cases the opposite situation is true: only the polyhedron shape model exists, and the corresponding spherical harmonics model must be constructed in order to apply it to compute curvature distributions at an arbitrary location on the surface. To carry out this model conversion, a multipass iterative residual fit (MIRF) was developed, which fits the spherical harmonic coefficients A_{jk} and B_{jk} as used in Eq. 3.49 to a set of (x, y, z) vertices. This MIRF was implemented following the strategy by Shen & Chung [106], who applied it to a different form of the spherical harmonic expression. For more information, the reader is referred to their work. The implementation of this filter was verified by sampling the spherical harmonic shape model of Phobos, and recovering the coefficients from the sampled shape with our filter.

In this test, the filter is fed a polyhedron shape constructed from spherical harmonics without the addition of any noise. As a result, the finite number of coefficients of the original model can be recovered to numerical precision. This is not true for arbitrary shapes, which may require an infinite number of coefficients. However, as a limited number of vertices are available as input data to the filter, aliasing issues may occur when trying to fit high-degree coefficients. Due to the non-uniform spacing of most small-body polyhedron shape models, this limit is difficult to predict analytically [98]. Instead, numerical tests are used when applying our filter to fit polyhedron shape models.

Finally, it is repeated that one main difference between the two models relates to their continuity. As the spherical harmonic shape model is continuous, it can provide the surface curvature at any given point on the surface. The polyhedron model, which is a discrete shape, can only do so at the vertices of the considered shape model. This makes the spherical harmonic model more fit for producing smooth distributions.

3.3 Applications to simple shapes

Following the derivation of the lift-off velocity on a curved body, some general properties of the expression that apply to any curved body were discussed. This discussion is now extended by investigating the global lift-off velocity distribution across some simple shapes, namely, the sphere and several tri-axial ellipsoids, in order to gain further insight into the properties of the lift-off velocity.

3.3.1 Applications to the sphere

The curved lift-off velocity expression, as given by Eq. 3.15, can be simplified significantly when studying departure from the surface of a sphere rather than some arbitrary body. On a sphere, the radius of curvature ρ is always equal to the radius of the sphere and the departure point $\mathbf{P} = \rho \hat{\mathbf{b}}_\rho$. Ignoring external sources other than the gravity of the (uniform-density) sphere, the acceleration \mathbf{a}_E must always be aligned with the $\hat{\mathbf{b}}_\rho$ axis, such that $\hat{\mathbf{b}}_\rho \cdot \mathbf{a}_E = g$.

Reduced expression Substituting these simplifications into Eq. 3.15 allows us to write for the lift-off velocity on the surface of a sphere:

$$V_{\theta,s} = \pm \sqrt{\rho^2 (\hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega})^2 + \rho^2 \hat{\mathbf{b}}_\rho \cdot ([\tilde{\boldsymbol{\Omega}}][\tilde{\boldsymbol{\Omega}}]\mathbf{b}_\rho) - \rho g - \rho \hat{\mathbf{b}}_\delta \cdot \boldsymbol{\Omega}} \quad (3.68)$$

which can be reduced to:

$$V_{\theta,s} = \pm \sqrt{-\rho^2 \Omega_\theta^2 - \rho g - \rho \Omega_\delta} \quad (3.69)$$

In this, the angular velocity components are equal to:

$$\Omega_\rho = \Omega \sin \phi \quad \text{and} \quad \Omega_\theta = \Omega \sin \eta \cos \phi \quad \text{and} \quad \Omega_\delta = \Omega \cos \eta \cos \phi \quad (3.70)$$

in which $\phi = \sin^{-1}(z/\rho)$ is the latitude of the departure point and the azimuthal angle η indicates the departure direction. In Fig. 3.14, the effect of changes in ϕ and η on the resulting departure plane orientation is illustrated. Substituting the angular velocity components of Eq. 3.70 into Eq. 3.68 finally yields for the lift-off velocity on the surface of a sphere:

$$V_{\theta,s} = \pm \sqrt{-\rho^2 \Omega^2 \sin^2 \eta \cos^2 \phi - \rho g - \rho \Omega \cos \eta \cos \phi} \quad (3.71)$$

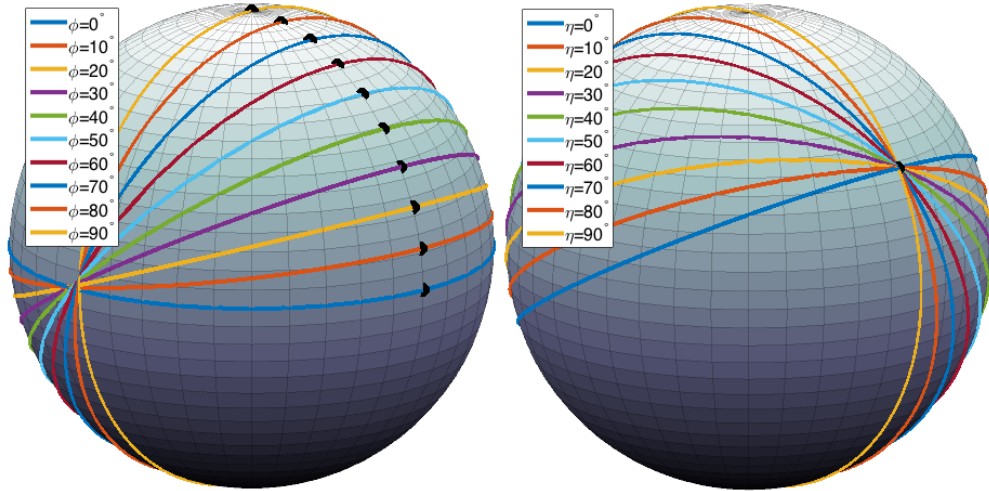


Figure 3.14: Lift-off plane geometry on the surface of a sphere for variations in (left) departure point latitude ϕ with fixed $\eta = 0^\circ$ and (right) departure direction η with fixed $\phi = 30^\circ$.

The dwarf planet Ceres is approximately spherical with a mean radius $\rho = 473$ km, $g = -0.29$ m/s², and $\Omega = 1.923 \times 10^{-4}$ rad/s. Using Eq. 3.71 to compute the lift-off velocity on Ceres's equator, where $\phi = 0^\circ$ and $\eta = 0^\circ$, as:

$$V_{\theta,s,eq} = \pm\sqrt{-\rho g} - \rho\Omega \Rightarrow \begin{cases} V_{\theta,1} = +279 \text{ m/s} \\ V_{\theta,2} = -461 \text{ m/s} \end{cases} \quad (3.72)$$

These velocities are respectively the prograde and retrograde velocities, relative to Ceres's surface, that correspond to an inertial velocity of 370 m/s, which is the circular velocity $V_c = \sqrt{\mu/\rho}$ at Ceres's equator. In a similar fashion, at the poles it holds that $\phi = \pm 90^\circ$. Substituting this into Eq. 3.71 yields for the polar lift-off velocity on Ceres:

$$V_{\theta,s,pol} = \pm\sqrt{-\rho g} \Rightarrow \begin{cases} V_{\theta,1} = +370 \text{ m/s} \\ V_{\theta,2} = -370 \text{ m/s} \end{cases} \quad (3.73)$$

These two solutions are equal in magnitude and opposite in direction. This result is expected, as a particle located at the poles of a sphere is not imparted any velocity by the sphere's rotation, and therefore does not experience the directionality bias that is present when lifting off from the equatorial plane.

Simulations To further validate our results for the lift-off velocity on the surface of a sphere, several simulations of the motion of a particle on a rotating sphere are performed, with $\rho = 200$ m, $\mu = 2.5$ m³/s², and $\Omega = 1.5 \times 10^{-4}$ rad/s. Fig. 3.15 shows the motion of particles departing from an arbitrary position on this sphere. It is noted that these results (as all others shown in this chapter) are plotted in a body-fixed, rotating reference frame.




Figure 3.15: Simulated lift-off of a particle from the surface of a rotating sphere.

This figure shows trajectories of particles with $0.9V_\theta$, $1.0V_\theta$, and $1.1V_\theta$, lifting off in both the $+\hat{\mathbf{b}}_\theta$ and $-\hat{\mathbf{b}}_\theta$ directions. It can be seen that the particles with the lower and higher velocities respectively move into and away from the sphere surface. The particles that are given exactly $1.0V_\theta$ indefinitely track the surface, although the trajectories that do not move in the equatorial plane slowly curve away from the targeted lift-off plane. This behavior is expected, and is the result of the rotation of the sphere. As all of the simulations with $1.0V_\theta$ track the sphere surface, it is concluded that the expressions correctly capture the lift-off velocity on a sphere.

Lift-off velocity distribution With Eq. 3.71 validated, a more broad investigation of the lift-off velocity distribution across the surface of a sphere may be performed. To that end, Fig. 3.16 shows the variation of the lift-off velocity V_θ as a function of the azimuthal angle η , for a number of latitudes ϕ .




Figure 3.16: Lift-off velocity curves for variations in ϕ and η on the surface of a rotating sphere.

These results show that, when the particle is located on either poles of the sphere (where $\phi = 90^\circ$), the lift-off velocity is constant and independent of the angle η . As the particle position is shifted away from the poles, a minimum and maximum lift-off velocity appear; they occur when the particle velocity is directed parallel to the equatorial plane, respectively along with or against the local rotational velocity. The magnitude difference between the extrema increases as the departure point is shifted closer to the equator. This, too, is expected, as a particle closer to the equator on a sphere is located further from the body rotation axis and experiences a greater rotational velocity from the body surface, thus decreasing the minimum and increasing the maximum lift-off velocities.

Finally, the lift-off velocity across the surface of the sphere is computed, while varying the departure direction η between 0° and 360° and determine the minimum and maximum velocity encountered at each surface point. The distributions of these minima and maxima, as well as the difference between their magnitudes, are shown in Fig. 3.17, plotted on a latitude-longitude grid. As expected, the lift-off velocity distribution is symmetric with respect to the equatorial plane and independent of the departure longitude λ , due to the sphere's symmetry properties.

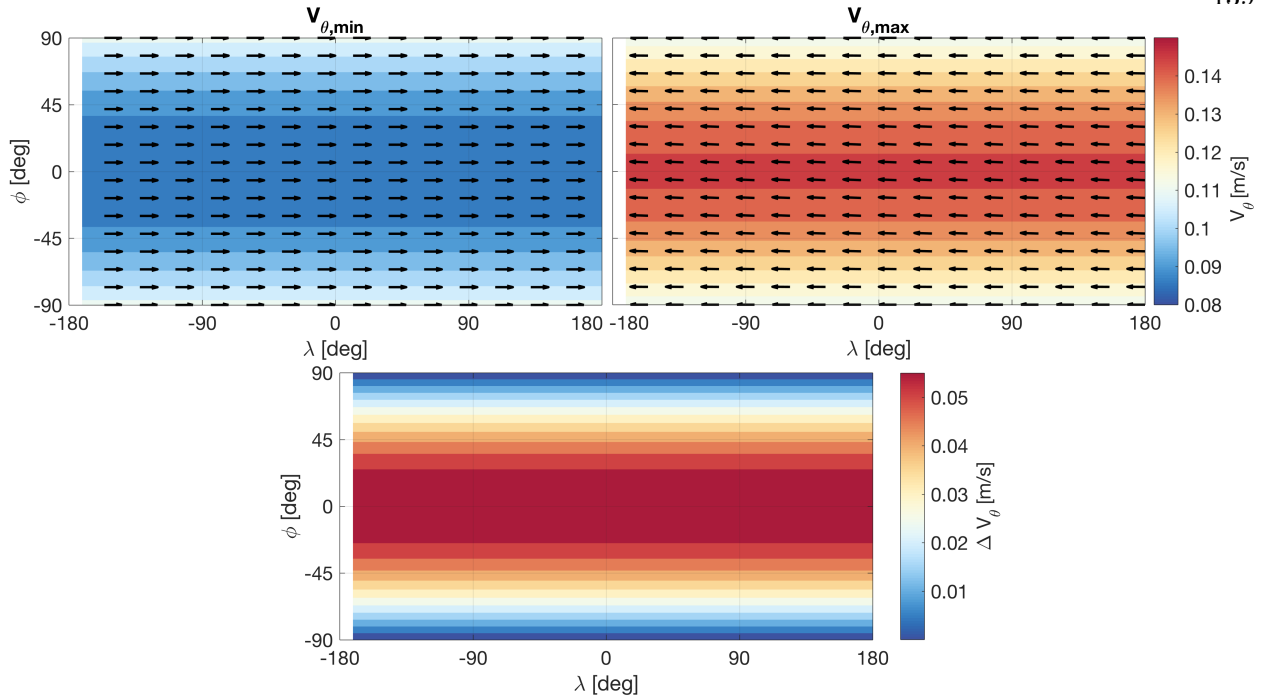


Figure 3.17: (left) Minimum lift-off velocity, (right) maximum lift-off velocity, and (bottom) extremum difference distribution across the surface of a rotating sphere. The arrows indicate the directions in which the respective extrema occur.

Furthermore, the minimum lift-off velocity increases with increasing latitude, whereas the maximum velocity decreases, due to the aforementioned distribution of the rotational velocity across the body surface. The largest differences between the minimum and maximum velocity occur at the equator; the difference is zero at the poles. All minimum velocities occur when $\eta = 0^\circ$; the maximum velocities occur for $\eta = 180^\circ$, in agreement with the earlier statement that the minima and maxima occur when the lift-off velocity is directed respectively along or against the local rotational velocity.

3.3.2 Applications to ellipsoids

The results presented above for the lift-off velocity distribution on a rotating sphere provide a great validation point for our formulations, but are not particularly applicable to the motion of particles on the surface of actual small bodies, due to their characteristically non-spherical shapes.

A more pragmatic application is the (tri-axial) ellipsoid, which is frequently used as a first estimate of the shape of a small body and has been used by previous authors to study the surface motion on small bodies [39]. In Section 3.2, the partials of the parametric $\mathbf{r}(\phi, \lambda)$ necessary to compute the curvature across an ellipsoid were provided. The gravitational acceleration of a constant-density ellipsoid with its axes aligned with an $\{x, y, z\}$ axis system may be computed as [102]:

$$\begin{aligned} g_x &= -\frac{3\mu x}{2} \int_{\Lambda(\mathbf{r})}^{\infty} \frac{du}{(a^2 + u)\Delta(u)} \\ g_y &= -\frac{3\mu y}{2} \int_{\Lambda(\mathbf{r})}^{\infty} \frac{du}{(b^2 + u)\Delta(u)} \\ g_z &= -\frac{3\mu z}{2} \int_{\Lambda(\mathbf{r})}^{\infty} \frac{du}{(c^2 + u)\Delta(u)} \end{aligned} \quad (3.74)$$

in which μ is the gravitational parameter of the ellipsoid. The parameter $\Lambda(\mathbf{r})$ is defined through $\Phi(\mathbf{r}, \lambda) = 0$ in which:

$$\Phi(\mathbf{r}, u) = \frac{x^2}{a^2 + u} + \frac{y^2}{b^2 + u} + \frac{z^2}{c^2 + u} \quad (3.75)$$

On the surface of the ellipsoid, it holds that $\Lambda(\mathbf{r}) = 0$. Finally, the $\Delta(u)$ function is defined as:

$$\Delta(u) = \sqrt{(a^2 + u)(b^2 + u)(c^2 + u)} \quad (3.76)$$

This model may be evaluated using Carlson's elliptic integrals; for a more detailed discussion, the reader is referred to [102]. Combined with the method for curvature computation, this allows for an evaluation of the lift-off velocity across the surface of an ellipsoid.

Simulations To verify whether the strategy outlined above indeed yields the correct velocity, several simulations of particles departing from a rotating ellipsoid were performed, with $a = 200$ m, $b = 100$ m, $c = 50$ m, $\mu = 2.5 \text{ m}^3/\text{s}^2$, and $\Omega = 1.5 \times 10^{-4}$ rad/s. Once again, each simulation is performed with $0.9V_\theta$, $1.0V_\theta$, and $1.1V_\theta$. The results can be found in Fig. 3.18.

In all cases, the trajectories with $1.0V_\theta$ provide the best surface tracking; both the trajectories with $0.9V_\theta$ and $1.1V_\theta$ diverge from the local surface faster. This behavior matches that of particles lifting off from the surface of a sphere; however, the trajectories of particles with $1.0V_\theta$ do show some differences. As the radius of curvature of a sphere is constant across its surface, the $1.0V_\theta$ particles indefinitely track that surface. A particle lifting off from a generic ellipsoid instead experiences a

change in the local radius of curvature as it moves across the ellipsoid surface. As a result, the $1.0V_\theta$ lift-off will track the surface only in some neighborhood of the initial departure point. As the particle moves away from this departure point, its trajectory will diverge from the surface, as is visible near the edges of the simulations of Fig. 3.18. The rate of divergence depends on the rate of change of ρ and on the local gravity, both of which vary across the ellipsoid surface. Nevertheless, it is concluded from these simulations that the provided strategy correctly yields the lift-off velocity on a general ellipsoid.




Figure 3.18: Simulated lift-off for particles on the surface of a general ellipsoid.

Lift-off velocity distribution The lift-off velocity variation across the surface of three ellipsoids is studied: an *oblate* ellipsoid with $a = b = 200$ m and $c = 100$ m, a *prolate* ellipsoid with $a = b = 150$ m and $c = 200$ m, and a *general* ellipsoid $a = 200$ m, $b = 100$ m, and $c = 50$ m. The three ellipsoids are illustrated (at the same scale) in Fig. 3.19 and are given $\mu = 2.5 \text{ m}^3/\text{s}^2$, and $\Omega = 1.5 \times 10^{-4} \text{ rad/s}$. The distributions of minimum and maximum radius of curvature are shown on a longitude-latitude grid in Figs. 3.20 through 3.22, in which the respective directions of extremum curvature are marked with arrows. In all ellipsoids, there are points at which the minimum and maximum ρ have similar values, and points where there is a very large difference between the two. Due to the multiplier-like effect of ρ in Eq. 3.15, the surface distribution of the lift-off velocity on these ellipsoids is expected to closely follow that of the radius of curvature.




Figure 3.19: The tested oblate, prolate, and general ellipsoids.

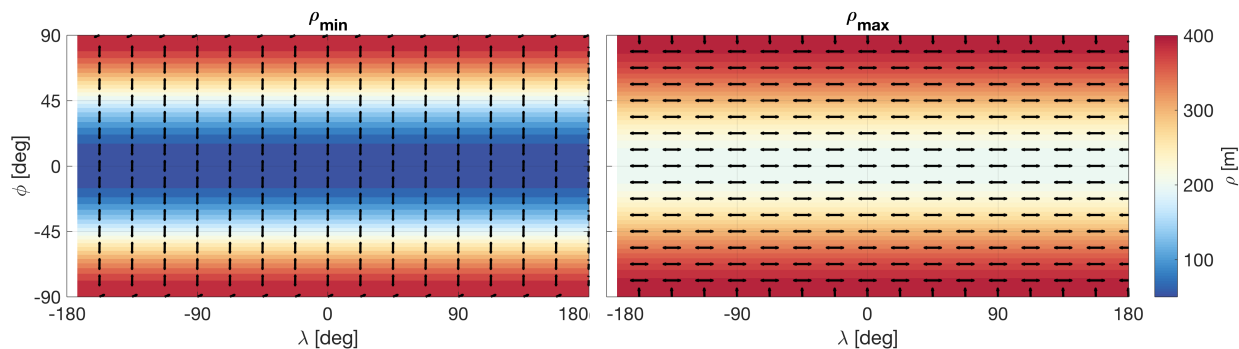


Figure 3.20: (left) Minimum and (right) maximum ρ across the surface of the oblate ellipsoid.




Figure 3.21: (left) Minimum and (right) maximum ρ across the surface of the prolate ellipsoid.




Figure 3.22: (left) Minimum and (right) maximum ρ across the surface of the general ellipsoid.

Using the ρ distributions, the lift-off velocity V_θ is computed on the three ellipsoids while varying η to obtain the extremum velocities encountered at each point, as well as the respective directions in which they appear. The resulting distributions are shown in Figs. 3.23 through 3.25. Several patterns can be distinguished, in particular when inspecting the directions in which the respective extrema occur. Similar to the lift-off velocities on the surface of a sphere, the minimum and maximum velocities are observed to always have a component respectively along or against the local surface rotational velocity. However, contrary to the sphere, the directions of the velocity extrema are not necessarily parallel to the equatorial plane. Furthermore, it is found that the extremum direction varies as a function of latitude on the oblate and prolate ellipsoids, and as a function of both latitude and longitude on the general ellipsoid. On the general ellipsoid, a clear pattern in the minimum lift-off velocity direction as a function of the leading/trailing edges of the body is also seen, with the lowest velocities occurring on the far ends of the body. Due to the ellipsoid's symmetry with respect to the equator, the minimum/maximum lift-off velocity holds for lift-off in both a north and south direction.




Figure 3.23: (left) Minimum and (right) maximum V_θ across the surface of the oblate ellipsoid.




Figure 3.24: (left) Minimum and (right) maximum V_θ across the surface of the prolate ellipsoid.




Figure 3.25: (left) Minimum and (right) maximum V_θ across the surface of the general ellipsoid.

To further investigate these phenomena, the variation in lift-off velocity and radius of curvature is plotted as a function of the azimuthal angle η , for a number of latitudes on the considered ellipsoids. This is first done for the oblate spheroid in Fig. 3.26, in which the lift-off velocity is independent of the longitude λ of the departure point. In it, the minimum and maximum lift-off velocities for each considered latitude are marked with respectively dots and diamonds.

From the V_θ distribution, all maximum lift-off velocities are observed to occur at $\eta = 180^\circ$; the direction opposite the local surface rotational velocity, which happens to also be the point where the greatest ρ values are encountered. The behaviour of the minimum V_θ points is more complex. At high latitudes, the minima occur at $\eta = 0^\circ$. Inspecting the ρ curves, it is found that at such high latitudes, there is little variation in the radius of curvature when η is changed. As a result, the lift-off velocity is mostly dependent on the local rotational velocity, with $\eta = 0^\circ$ being aligned with the that velocity. When moving to lower latitudes, a minimum is observed in the the radius of curvature curves near $\eta = 90^\circ$. However, this does not correspond to the V_θ minimum. It appears that the minimum lift-off velocity results from a more intricate combination of the radius of curvature and local surface rotational velocity, and matches the trends previously seen in Fig. 3.23.




Figure 3.26: (top) V_θ and (bottom) ρ vs. η on the surface of the oblate ellipsoid.

A different behavior is seen in the V_θ distribution of the prolate ellipsoid, as shown in Fig. 3.27, where the minimum lift-off velocities occur at minimum ρ with $\eta = 0^\circ$, and the maximum velocities display a more complex behavior. Finally, the curves corresponding to the general ellipsoid are

shown in Fig. 3.28. Here, a complex behavior in the maximum lift-off velocity is seen, though it still closely adheres to the ρ distribution, and matches the trends seen in Figs. 3.24 and 3.25.




Figure 3.27: (top) V_θ and (bottom) ρ vs. η on the surface of the prolate ellipsoid.

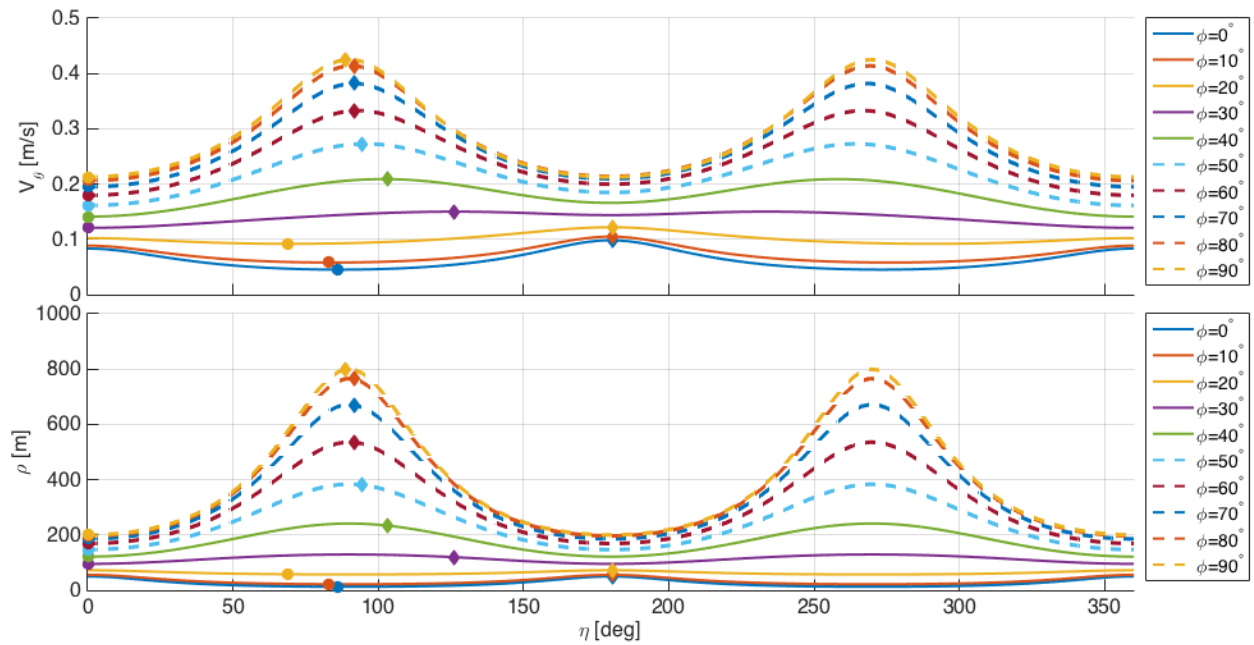


Figure 3.28: (top) V_θ and (bottom) ρ vs. η on the surface of the general ellipsoid.

3.4 Applications to Solar System bodies

The simple applications to the sphere and ellipsoids provided above have yielded insight into the general properties and trends of lift-off velocity distributions. With this understanding, the lift-off velocity on actual solar system bodies can be investigated. Here, such investigations were performed for the asteroids 1999 KW₄ Alpha, Bennu, and Eros, as well as the Mars moon Phobos.

3.4.1 1999 KW₄ Alpha

The 1999 KW₄ binary asteroid system consists of a fast-spinning, diamond-shaped primary and a smaller, tidally locked secondary. As the system is thought to be representative of the binary asteroid population, results of the lift-off dynamics on its bodies are of interest to a lot of these systems. Furthermore, 1999 KW₄ is the only binary asteroid system to have been mapped at an appreciable resolution, using radar imagery from an Earth flyby [86]. The resulting shape model lends itself to an application of the lift-off theory, with a particular application to the primary body, known as (1999 KW₄) Alpha. This interest stems from Alpha's high spin rate, which causes its equatorial region to rotate at near-orbital velocity. 1999 KW₄ Alpha is therefore expected to exhibit relatively low lift-off velocities for an object of its mass. Table 3.1 provides Alpha's properties [85].

Table 3.1: Characteristic properties of asteroid 1999 KW₄ Alpha [85].

Parameter	Symbol	Value
Gravitational parameter	μ	157.04 m ³ /s ²
Rotation period	T	2.7645 hr
Semi-major axis	a	708.5 m
Semi-intermediate axis	b	680.5 m
Semi-minor axis	c	591.5 m
Vertices in shape model	n_P	4,586
Facets in shape model	$n_{\mathcal{F}}$	9,168

Polyhedron shape Fig. 3.29 provides the minimum and maximum lift-off velocities computed on each of the 4,586 vertices of 1999 KW₄ Alpha's polyhedron shape model. In this computation, effects from the small and distant Beta body were ignored. The color coding of the facets in this figure is indicative of the corresponding lift-off velocity magnitudes. When the region sur-

rounding a vertex is black, lift-off is impossible, *i.e.*, the region is concave. If a region is black in plots of both the minimum and maximum velocities, the region is fully concave (with both $\kappa^+ < 0$ and $\kappa^- < 0$). If a region is black only in the maximum velocity plot, the region is mixed convex/concave (with $\kappa^+ \cdot \kappa^- < 0$). The ridge lift-off velocity distribution is also shown in Fig. 3.30. These results will be discussed shortly.




Figure 3.29: (*left*) Minimum and (*right*) maximum V_θ across the polyhedron model of 1999 KW₄ Alpha.




Figure 3.30: Ridge lift-off velocity across the polyhedron model of 1999 KW₄ Alpha.

Spherical harmonics shape In order to apply the expressions for spherical harmonic curvature to 1999 KW₄ Alpha, the spherical harmonic shape model must first be constructed from the polyhedron model with the MIRF. As previously discussed, there will be some optimum degree J at which the spherical harmonic shape provides the ‘best’ fit to the polyhedron shape. In order to quantify the quality of our fit, 100 polyhedron vertices are excluded from the ‘training’ set that is fed to the filter. The MIRF is then run to obtain the A_{jk} and B_{jk} coefficients and combine those with Eq. 3.49 to construct the spherical harmonic equivalent of the excluded polyhedron vertices. By computing the RMS of the radius difference of the excluded points while varying J , the accuracy of the spherical harmonic fit can be quantified. Fig. 3.31 plots the RMS of both the vertices included in the fit and the vertices excluded from the fit.




Figure 3.31: RMS error of spherical harmonics fit of 1999 KW₄ Alpha.

Inspecting this figure, one can see that the RMS of the vertices included in the MIRF fit continues to decrease as higher spherical harmonic degrees J are fit. This could indicate that the quality of the fit is improving when these higher degrees are added. However, the RMS of the excluded points shows a minimum at degree $J = 32$ at an RMS error of 3.1 m. This indicates that aliasing problems occur when higher degrees are used; the analysis is therefore done with the degree-and-order 32 spherical harmonic model. Fig. 3.32 plots both the original polyhedron shape model and the $J = 32$ spherical harmonic model of 1999 KW₄ Alpha. The two indeed appear very similar at this scale.




Figure 3.32: (*left*) Polyhedron and (*right*) $J = 32$ spherical harmonic model of 1999 KW₄ Alpha.

For a more detailed comparison, consider to the XZ-plane intersection shown in Fig. 3.33. This reveals that the spherical harmonic shape indeed mostly matches the polyhedron, though there are some local variations that are averaged out by the spherical harmonics. Due to this, correspondingly small local differences are expected in the surface curvatures and lift-off velocities computed on the spherical harmonic model, though the general trends should certainly be preserved. The minimum and maximum lift-off velocity distribution across 1999 KW₄ Alpha's surface as computed with the $J = 32$ spherical harmonic model can be found in Fig. 3.34.

Discussion Reviewing the minimum lift-off velocity distributions, lift-off is found to be possible, in at least some local tangential directions, from most regions on 1999 KW₄ Alpha's surface, at velocity magnitudes between roughly 0.1 and 0.5 m/s. The majority of the regions have lift-off velocities on the lower end of this range; this is especially true for the equatorial ridge, which has a very small lift-off velocity. This is the result of 1999 KW₄ Alpha's high rotation rate: its equatorial ridge rotates at near-orbital velocity, such that very little additional 'push' is required for an object to lift off. Also due to this high rotation rate, the maximum lift-off velocity is much higher and in the 1.0 to 1.5 m/s range. As the surface radius of curvature in the north-south direction is small on the equatorial ridge, this region's maximum lift-off velocity is correspondingly low. Finally, it is found from the maximum velocity distribution that roughly half of 1999 KW₄ Alpha's surface




Figure 3.33: XZ-plane intersection of the polyhedron and spherical harmonic model of 1999 KW₄ Alpha.




Figure 3.34: (*left*) Minimum and (*right*) maximum lift-off velocity across the spherical harmonic model of 1999 KW₄ Alpha.

has local directions in which lift-off is impossible. This indicates that the respective regions are locally concave in some azimuthal directions η . If a region appear black in both the minimum and the maximum lift-off velocity distributions, this region is locally concave in *all* azimuthal directions. The relative ease with which an object on the ridge may displace itself and enter an orbital arc can

be both advantageous and disadvantageous. From a geophysical point of view, these results indicate that dust particles or rocks displaced from mid-to-high latitude regions and moving towards the equatorial region are very likely to lift off, as their velocity will be naturally increased due to the geopotential difference between the respective regions.

Lift-off from a ridge or asperity is possible from all locations on the surface, regardless of their convexity. This makes sense: even in a locally concave region, it is possible for an object to roll off a small rock and experience a very brief air time. The reader is reminded that Fig. 3.30 displays the *pseudo-normalized* form of the ridge lift-off velocity, such that the lift-off for an object with radius $R = 1$ m will occur at roughly 1.80 cm/s, but at only 0.57 cm/s for an object with radius $R = 0.1$ m. Although this might be high enough for the operation of a wheeled surface vehicle, the general challenge of delivering a payload to the surface of such a fast-rotating asteroid remains. It is repeated that it will be difficult for the vehicle to achieve speeds much greater than V_R ; the reduction in traction at speeds greater than V_R will limit the maximum speed that can be achieved using the wheels. Finally, it is noted that the ratios between the ridge and curved lift-off velocities agree with those predicted by Eq. 3.24.

Comparing the lift-off velocities computed on 1999 KW₄ Alpha's polyhedron with their spherical harmonic counterparts, matching magnitudes and general trends are observed. However, small local differences between the two models exist, and lift-off from a particular position may be possible on one model but impossible on the other. This dichotomy is a result of the inherent surface smoothing applied in the spherical harmonic fit, as previously illustrated in Fig. 3.33. Nonetheless, the spherical harmonic model is useful as it provides a continuous distribution of the lift-off velocity, while the polyhedron model can only provide velocities at the model vertices.

Fast-spinning asteroids such as 1999 KW₄ Alpha tend to be oblate in shape, resulting in a lower geopotential at the equator than at the poles. Indeed, as mentioned above, a particle displaced from a high latitude will naturally tend to move towards the equatorial region. It is therefore logical to ask what the minimum latitude is from which a particle must have departed, given that it lifted off before reaching the equatorial region. Here, this question is addressed by

constructing a *minimum latitude of release* map, for which we first cycle over a grid of longitudes and compute the minimum north-south lift-off velocity for each longitude point, by varying the latitude. On 1999 KW₄ Alpha, the line connecting these points follows the equatorial ridge, as marked by the dashed line in Fig. 3.35. For each point along this minimum line, we then compute the latitude from which an object must have departed in a north-south direction in order to have lifted off by the time it reaches the point of minimum lift-off velocity along that longitude, by making use of the Jacobi constant J :

$$J = \frac{1}{2} \dot{\mathbf{r}} \cdot \dot{\mathbf{r}} - \frac{1}{2} (\boldsymbol{\Omega} \times \mathbf{r}) \cdot (\boldsymbol{\Omega} \times \mathbf{r}) - U(\mathbf{r}) = \frac{1}{2} V^2 - W(\mathbf{r}) \quad (3.77)$$

in which $U(\mathbf{r}) \geq 0$ is the gravitational potential and $W(\mathbf{r}) = U(\mathbf{r}) + \frac{1}{2} \Omega^2 (x^2 + y^2)$ is the geopotential (assuming uniform rotation about the z -axis). When a particle departs from rest from some \mathbf{r}_0 and moves to a position \mathbf{r}_1 with lower geopotential, it will have gained a velocity v_1 that can be computed by equating the Jacobi constants at both positions, *i.e.*:

$$W(\mathbf{r}_0) = \frac{1}{2} \{ V_1^2 + W(\mathbf{r}_1) \} \quad (3.78)$$

in which $0 < \{ \leq 1$ expresses an energy loss due to friction when converting geopotential energy into kinetic energy. At $\{ = 1$, there is no energy loss. Rewriting Eq. 3.78 for the velocity V_1 yields:

$$V_1 = \sqrt{2\{ (W(\mathbf{r}_0) - W(\mathbf{r}_1)) \}} \quad (3.79)$$

By setting V_1 equal to the minimum north-south lift-off velocity, the latitude where $W(\mathbf{r}_0)$ satisfies Eq. 3.79 can be found. This allows for a computation of the line of minimum latitude of release for a given value of $\{$. Fig. 3.35 shows some sample lines of minimum latitude of release for 1999 KW₄ Alpha. For each considered value of $\{$, the corresponding line shows the minimum latitude from which an object, such as a grain or rock, must have departed if it lifted off before or when reaching the equatorial ridge. When $\{$ is decreased, the object loses more energy due to friction, such that it can depart from a higher latitude and experience a greater change in geopotential before achieving lift-off velocity. Quite large changes in $\{$ are required for the lines to shift to higher latitudes; this is

due to the fact that 1999 KW₄ Alpha's equatorial region is at much lower geopotential than its polar regions, due to the fast rotation rate of the former. Although the matching between the dissipative parameter f and more intuitive parameters such as the coefficient of friction and restitution is not obvious, these results suggest that objects departing from 1999 KW₄ Alpha's polar regions can be expected to lift off before reaching the equator.




Figure 3.35: Lines of minimum latitude of release on 1999 KW₄ Alpha.

3.4.2 Bennu

Asteroid Bennu was discovered in 1999 and is a potential Earth impactor, as it is expected to perform several close passes to Earth in the coming century. The asteroid can be reached from Earth with a relatively low ΔV and is the target of OSIRIS-REx, the third New Frontiers mission. This mission is set to perform remote sensing operations at Bennu in 2019, culminating in the acquisition of a sample that will return to Earth in 2023 [57]. Bennu has been observed extensively, resulting in a medium-resolution shape model that can be used for lift-off velocity computations. Although OSIRIS-REx will not perform surface mobility operations, the lift-off velocity results are still relevant to geophysical processes. Table 3.2 provides Bennu's characteristic properties [105].

As the strategy applied to compute the lift-off velocity distributions of Bennu is identical

Table 3.2: Characteristic properties of asteroid (101955) Bennu [105].

Parameter	Symbol	Value
Gravitational parameter	μ	5.2 m ³ /s ²
Rotation period	T	4.29746 hr
Semi-major axis	a	258.5 m
Semi-intermediate axis	b	250.0 m
Semi-minor axis	c	230.0 m
Vertices in shape model	n_P	1,348
Facets in shape model	$n_{\mathcal{F}}$	2,692

to that applied for 1999 KW₄ Alpha, the relevant figures are simply presented, followed by a brief discussion afterwards. The minimum and maximum lift-off velocities computed on Bennu's polyhedron model can be found in Fig. 3.36. Similarly, the pseudo-normalized ridge velocity can be found in Fig. 3.37. The spherical harmonic fitting of the polyhedron shape model displays a minimum RMS of 2.49 m at $J = 16$; the RMS variation as a function of the degree J can be found in Fig. 3.38. Finally, the spherical harmonic lift-off velocity distributions are shown in Fig. 3.39.

The minimum lift-off velocity on Bennu is similar across most of its surface at roughly 0.1 m/s, save for a few fully concave regions where lift-off is impossible in all directions. Interestingly, this matches the minimum velocity on 1999 KW₄ Alpha, despite the latter being roughly 30 times more massive than Bennu. This again illustrates the fast rotation rate of 1999 KW₄ Alpha: the relatively high mass and fast rotation of 1999 KW₄ Alpha result in the same minimum lift-off velocity as on the relatively low mass and slow rotation of Bennu.

Although Bennu also has an equatorial ridge, it is much less pronounced than 1999 KW₄ Alpha's with no distinctly lower lift-off velocities. The maximum lift-off velocity is on the order of 0.5 m/s, much lower than on 1999 KW₄ Alpha; its distribution shows that again about half of the surface is locally concave in some directions. The equatorial ridge is more visible in the ridge lift-off velocity, which is found to be lower on Bennu's ridge than on the rest of the body. For an object with radius $R = 1$ m, this velocity is approximately 9 mm/s, for an object with radius $R = 0.1$ m it reduces to 2.8 mm/s. The velocities computed on the spherical harmonic models again match those computed on the polyhedron, though distinctly fewer concave regions exist on the former.




Figure 3.36: (left) Minimum and (right) maximum V_θ across the polyhedron model of Bennu.




Figure 3.37: Ridge lift-off velocity across the polyhedron model of Bennu.




Figure 3.38: RMS error of spherical harmonics fit of Bennu.




Figure 3.39: (*left*) Minimum and (*right*) lift-off velocity across Benu's spherical harmonic model.

The lines of minimum latitude of release were also constructed, using the spherical harmonic shape model, as shown in Fig. 3.40. Inspecting this figure, it is found that objects lifting off at the equator can originate from relatively high latitudes, in particular when compared to 1999 KW₄ Alpha.




Figure 3.40: Lines of minimum latitude of release on Benu.

3.4.3 Eros

Asteroid Eros is the second largest near-Earth asteroid and was explored by the NEAR-Shoemaker spacecraft in 2000. This mission extensively mapped the asteroid’s surface, resulting in a global, high-resolution shape model which reveals its highly non-spherical shape. Table 3.3 provides the characteristic properties of Eros [69, 70]. Following the same approach taken for the previous bodies, the minimum and maximum lift-off velocity distribution were computed across Eros’s polyhedron shape model, as shown in Fig. 3.41, and the ridge lift-off velocity in Fig. 3.42.

Table 3.3: Characteristic properties of asteroid (433) Eros [69, 70].

Parameter	Symbol	Value
Gravitational parameter	μ	446,250 m ³ /s ²
Rotation period	T	5.2703 hr
Semi-major axis	a	17.35 km
Semi-intermediate axis	b	8.70 km
Semi-minor axis	c	7.00 km
Vertices in shape model	n_P	25,350
Facets in shape model	$n_{\mathcal{F}}$	49,152

When applying the MIRF to fit a spherical harmonic model to Eros’s polyhedron shape model, it is found that the RMS error remains large, for both low and high considered degree J . This phenomenon is illustrated in Fig. 3.43, where the original shape model is shown, as well as the $J = 18$ and $J = 40$ spherical harmonic fits. Although both fits coarsely capture Eros’ shape, severe aliasing issues are present, which deteriorate as J increases. This seems to suggest that spherical harmonics cannot be used to represent a body as non-spherical as Eros.

In an attempt to mitigate these aliasing issues, Eros’s polyhedron model is divided into five sectors (see Fig. 3.44); one spherical harmonic fit is performed for each separate sector. The RMS error of the resulting spherical harmonic fit is shown in Fig. 3.45, where a number of vertices were again excluded from the fit in order to quantify its performance. A minimum error of approximately 48 m is observed at degree $J = 24$; the corresponding shape model can be found in Fig. 3.46. This sectorized spherical harmonic model indeed matches the original polyhedron model, save for a few small artifacts. Fig. 3.47 shows the lift-off velocity distributions computed with this model.



Figure 3.41: (left) Minimum and (right) maximum lift-off velocity across Eros's polyhedron model.




Figure 3.42: Ridge lift-off velocity across Eros's polyhedron model.

Reviewing the polyhedron lift-off velocity results, it is found that lift-off is possible in at least some directions from most regions on Eros's surface, at minimum velocities between 1 and 2 m/s, with the lowest velocities occurring at the far ends of the body. The maximum lift-off velocities are on the order of 5 m/s, though large portions of the surface are partially or fully concave, such that lift-off is not possible in at least some azimuthal directions. The distributions appear quite 'spotty,' due to the high resolution of the shape models and corresponding local shape variations.




Figure 3.43: (*left*) Polyhedron, (*middle*) $J = 18$, and (*right*) $J = 40$ spherical harmonic model of Eros.




Figure 3.44: Sectors of the Eros polyhedron shape that are separately fitted.




Figure 3.45: RMS error of five-sector spherical harmonics fit of Eros.




Figure 3.46: (*left*) Polyhedron and (*right*) $J = 24$ sectored spherical harmonic shape model of Eros.



Figure 3.47: (*left*) Minimum and (*right*) maximum lift-off velocity across the $J = 24$, sectored spherical harmonic model of Eros.

From the ridge lift-off velocity distribution, it is found that a brief air time is possible at roughly 7 cm/s for an object with $R = 1$ m and at roughly 2 cm/s for an object with $R = 0.1$ m. This ridge velocity varies more distinctly than on the previously discussed bodies, with distinctly lower velocities occurring inside craters, where the local surface slope is high.

Finally, reviewing the results on the spherical harmonic shape model, much larger differences than those observed in comparisons on the previous bodies are found. The minimum velocity

distribution appears less spotty, having much more clustered regions where lift-off is impossible. This is a direct result from the inherent smoothing in the construction of the spherical harmonic fit. However, it is also noticed that the velocity magnitude is on average about a meter per second higher than on the polyhedron model. The transition zones between two adjacent sectors are distinctly visible as having a lower lift-off velocity than the surrounding regions. In an attempt to make the model more continuous at the boundaries, the separate sector fits were fed some vertices from outside of their boundaries, but this did not significantly improve our results. It appears that, due to Eros's strongly non-spherical shape, the surface smoothing inherent in the spherical harmonic filtering is stronger than in the previously studied, more spherical bodies.

3.4.4 Phobos

As a final application, the lift-off velocity distribution is computed across Phobos's degree-and-order $J = 45$ shape model. In this, the Hill-Clohessy-Wiltshire equations are used to capture the gravitational effects from Phobos as well as the tidal force from Mars:

$$\ddot{\mathbf{r}} = \omega^2 \begin{bmatrix} 3x \\ 0 \\ -z \end{bmatrix} + \omega \begin{bmatrix} 2\dot{y} \\ -2\dot{x} \\ 0 \end{bmatrix} + \nabla U(\mathbf{r}) \quad (3.80)$$

Figs.3.48 and 3.49, provide the minimum and maximum lift-off velocity distribution on Phobos's surface. A global minimum lift-off velocity $v_{\theta, min} = 1.60$ m/s is found. In other words, any surface vehicle operating on the surface of Phobos can safely achieve a surface velocity of 1.60 m/s without ceasing contact with the *averaged* surface. Most of Phobos's surface has a minimum lift-off velocity of at least 2 m/s. These minimum velocities generally occur for motion towards the east, along the rotational velocity of Phobos's surface. When an object moves towards the west, against the surface rotation, the lift-off velocity is much higher. In fact, most of Phobos's surface has a maximum velocity upwards of 10 m/s, well in excess of the L_1 return velocity. This suggests that vehicle can operate at relatively high velocities in directions opposite Phobos's surface. Additionally, the normalized ridge lift-off velocity distribution is also provided in Fig. 3.50. From this, it is found

that a wheeled vehicle with $R = 0.1$ m has $V_R \approx 2$ cm/s, while an astronaut with hip height $R = 0.9$ m has $V_R \approx 6$ cm/s. As mentioned before, it is this ridge lift-off velocity that is limiting to the operation of a surface vehicle: it will begin to briefly cease contact with the surface at this velocity, and will have difficulties reaching higher speeds due to a reduction in surface traction.




Figure 3.48: Minimum lift-off velocity on Phobos's surface.




Figure 3.49: Maximum lift-off velocity on Phobos's surface.




Figure 3.50: Normalized ridge lift-off velocity on Phobos's surface.

Chapter 4

Parametric studies

The integration of the modeling elements summarized in Chapter 2 yields a framework capable of performing high-fidelity simulations of probes with an arbitrary shape that are deployed to some target small body, complete with statistical surface features. These simulations have various applications, the most obvious of which is the validation of a deployment strategy. Given the release conditions of a probe, *i.e.*, its position, velocity, attitude, and corresponding uncertainties, a set of Monte Carlo-type simulations can predict the success rate at which that deployment strategy delivers probes to the small-body surface. Analyzing the resulting trajectories in further detail, one may similarly predict relevant deployment statistics such as the settling time and surface dispersion of the considered probe. Such studies are useful when investigating specific mission scenarios, *e.g.*, the deployment of the rovers onboard *Hayabusa2*, for which the combined effects of uncertainty in all environmental parameters can be predicted. This application is considered the next chapter.

Utility also lies in the *individual* variation of the relevant environmental (and probe) parameters, as those results may guide the probe design and surface modeling methods. For example, by studying the effects of the internal mass distribution of a probe on its motion, guidelines for the mechanical design of surface exploration probes may be obtained. Similarly, by individually studying the effects of, *e.g.*, the surface coefficient of restitution or the density of surface rocks, the relative strength with which they affect the probe deployment statistics can be determined. This, in turn, may help prioritize studies and observations of the corresponding surface parameters. These *parametric studies* are the focus of the present chapter, which is structured as follows.

Section 4.1 investigates the motion of spherical landers. Although surface exploration probes can and do have more complex shapes (that can be handled by the provided simulation methodology) this shape restriction significantly reduces the probe design space, thereby reducing the number of parameters to be considered in the parametric studies. More specifically, the effect of the surface interaction coefficients and the presence of (polyhedral) surface rocks are investigated. The resulting trends in the motion of a spherical probe provide a basis from which more detailed, shape-varying investigations of complex probes may depart. This is the focus of Section 4.2, which provides a detailed analysis of how probes with different shapes behave when deployed to the small-body surface. For these shapes, the surface interaction coefficients are also varied and their effects compared to those observed for spherical probes. The effect of the internal mass distribution of the complex probes on their settling behavior is then investigated. As a third and final part of this chapter, Section 4.3 investigates the effect of local surface topography on the motion of a deployed probe. Although a simple study investigating the effect of polyhedral rocks is included in Section 4.1 on spherical probes, Section 4.3 uses fully procedural features and investigates how their presence affects the motion of the complex probes. Finally, the section also provides an analysis of how the presence of regolith on the small-body surface influences deployment statistics, using a simple restitution-modulation scheme.

4.1 Spherical probes

Small-body exploration probes can and do have arbitrary shapes; expressions for the contact interactions between such a probe and the small-body surface are complex, as discussed in Chapter 2. The collisions of spherical probes, on the other hand, are far simpler since such collisions must always be *central*. The distinction between central and eccentric collisions was introduced in Section 2.5; a central collision occurs when the collision vector \mathbf{r} points through the spacecraft center of mass. As a result, the normal and friction forces are decoupled. This significantly simplifies the impact behavior compared to eccentric collisions, in which these forces *are* coupled.

The spherical probe, assumed rigid and undeformable, can thus be described entirely using only a radius R , mass m , and inertia matrix $[\mathbf{I}]$. As the latter is assumed to be spherically symmetric, it must be a scalar matrix of the form:

$$[\mathbf{I}] = jmR^2 \cdot \mathbb{I}_3 \quad (4.1)$$

in which \mathbb{I} is the identity matrix. The *reduced inertia* j describes the internal mass distribution of the sphere. It holds that $0 \leq j \leq 2/3$ for a spherically symmetric body, with $j = 0$ for a point mass and $j = 2/3$ for a thin shell. Both j and $[\mathbf{I}]$ increase when mass is moved outwards from the center of the sphere towards its shell. Fig. 4.1 illustrates the probe structure for varying j values, assuming a constant lander mass m . Given the radius R of some spherical probe and a set of internal components, j is dependent on the internal arrangement of these components. Lower j values may be achieved by placing high-density components, such as batteries, at the center of the probe, and vice versa.




Figure 4.1: Illustration of mass configurations of a spherical probe.

Due to variations in small-body morphology, strategies that result in successful deployment to some target may fail when naïvely applied to some other target. To establish the ‘difficulty’ of deploying a probe to the surface of a given body, the amended potential field can be used. More specifically, when the zero-velocity curve of the minimum amended potential lies relatively far away from the surface, the corresponding environment is *favorable* for lander deployment [117]; asteroid 25143 Itokawa is an excellent example of such a target. Itokawa will be used for sensitivity analyses in dynamically favorable environments; two nominal deployment scenarios to this target are first set up.

4.1.1 Nominal scenarios

Asteroid Itokawa was visited by the Hayabusa spacecraft in 2007 when it performed remote sensing operations at various altitudes and carried out a touch-and-go sample acquisition maneuver. It produced a global, high-resolution shape model of the asteroid, which is one reason why Itokawa is an excellent target environment for high-fidelity simulation of surface exploration probes. Table 4.1 provides some characteristic properties of the body [35]; the shape model was obtained from Gaskell et al. [36]. Table 4.2 also lists the nominal spacecraft, surface interaction, and rock parameters that are used in simulation. Unless noted otherwise, the interaction coefficients of the spacecraft with respect to the rocks are equal to those of the spacecraft with respect to the polyhedral surface. The nominal rock distribution parameters correspond to $K \approx 14$ rocks larger than $s_{min} = 5$ cm per square meter of surface area. Two different scenarios will be investigated; 1,000 simulations are performed for each scenario. These simulations use the polyhedron model for shape representation, with procedurally-generated polyhedral surface rocks.

Table 4.1: Characteristic properties of asteroid 25143 Itokawa [35, 36].

Parameter	Value	Parameter	Value
Gravitational parameter μ	2.341 m ³ /s ²	Semi-major axis a	268 m
Density σ	1.90 g/cm ³	Semi-intermediate axis b	147 m
Rotation period T	12.132 hr	Semi-minor axis c	105 m
Vertices in shape model n_P	98,306	Facets in shape model $n_{\mathcal{F}}$	196,608

Table 4.2: Nominal spacecraft and environmental parameters.

Parameter	Value	Parameter	Value
Lander mass m	5 kg	Rock power index α	2.8
Lander radius R	0.125 m	Minimum rock size s_{min}	0.05 m
Lander reduced inertia j	0.40	Maximum rock size s_{max}	0.50 m
Coefficient of restitution e	0.65	Rock initialization number k_0	2.50×10^{-3}
Coefficient of friction f	1.00	Rock initialization size s_0	2.50 m
Coefficient of rolling resistance C_{RR}	0.035	Minimum rock sink ratio r_{min}	-0.4
		Maximum rock sink ratio r_{max}	0.25

Table 4.3: Deployment parameters of Itokawa scenarios.

Parameter	High-altitude value	Low-altitude value
Release position \mathbf{x}	[500, 0, 0] m	[0, 0, 200] m
Position 3σ uncertainty $\Delta\mathbf{x}$	[30, 30, 30] m	[15, 15, 15] m
Release velocity \mathbf{v}	[-0.03, 0, 0] m/s	[0, 0, -0.05] m/s
Velocity 3σ uncertainty $\Delta\mathbf{v}$	[0.01, 0.01, 0.01] m/s	[0.01, 0.01, 0.01] m/s




Figure 4.2: Overview of the nominal Itokawa deployments.

High-altitude deployment In the high-altitude deployment scenario, probes are released approximately 200 m above Itokawa’s surface (≈ 1 body radius), resulting in first impacts centered around its neck region. This scenario corresponds to probes being released from a mothership on a (low-velocity) hyperbolic flyby; the corresponding deployment parameters are listed in Table 4.3. The results of this nominal deployment are included in Figs. 4.2 through 4.4 and are briefly reviewed here. The release, first impact, and settling locations are shown in three-dimensional form on the left side of Fig. 4.2. Corresponding two-dimensional latitude-longitude maps are shown on the left side of Fig. 4.3, in which a heat map shows the clustering of final settling positions and a contour map shows the relevant locations as scattered points. In all plots except the heat map, the surface coloring is indicative of the local geopotential, computed as $V = \Omega^2(x^2 + y^2) - U(\mathbf{x})$, that represents the energy of the local surface. Dark regions have a high geopotential, while bright regions have a low geopotential. Due to the natural dynamics of the small-body environment, particles will tend towards areas of low geopotential, similar to how a ball will settle at the bottom of a hill rather than on its crest. This is clearly visible in Fig. 4.3, in which the concentration of probe settling

locations is densest near the $(\lambda = -20^\circ, \phi = -20^\circ)$ region. This low-slope region is known as *Muses Sea* and is covered with fine regolith particles [73]. In the two-dimensional maps, Muses Sea is the large area covering the bottom half of the plot.

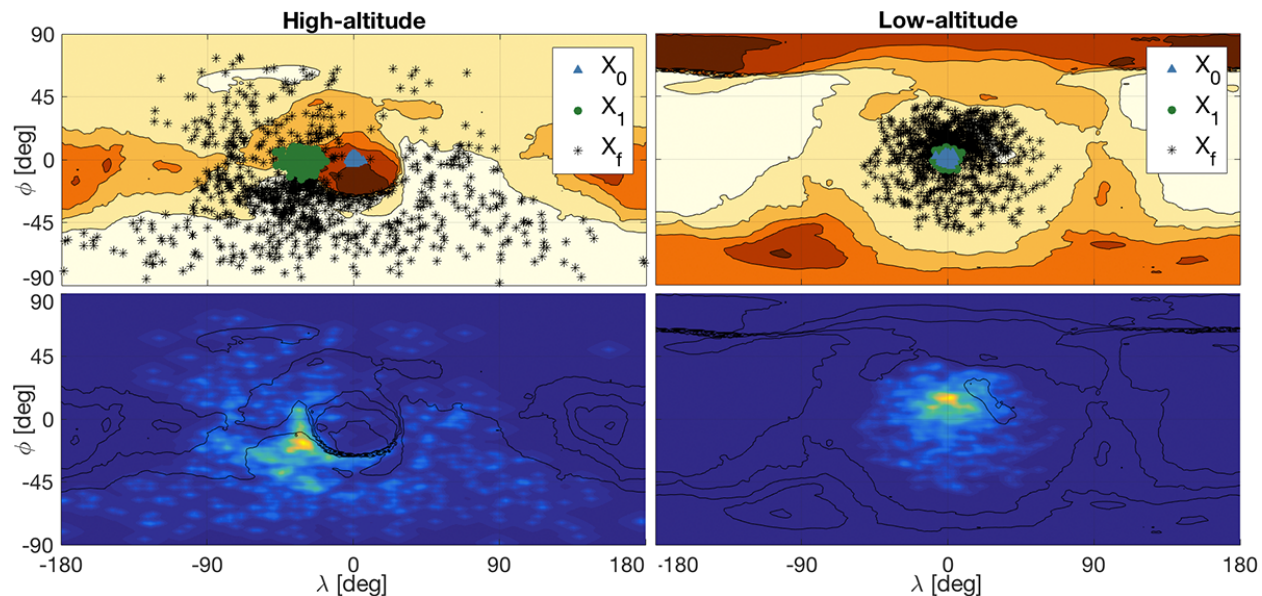


Figure 4.3: Heat map and position scatter plot of nominal Itokawa deployments.

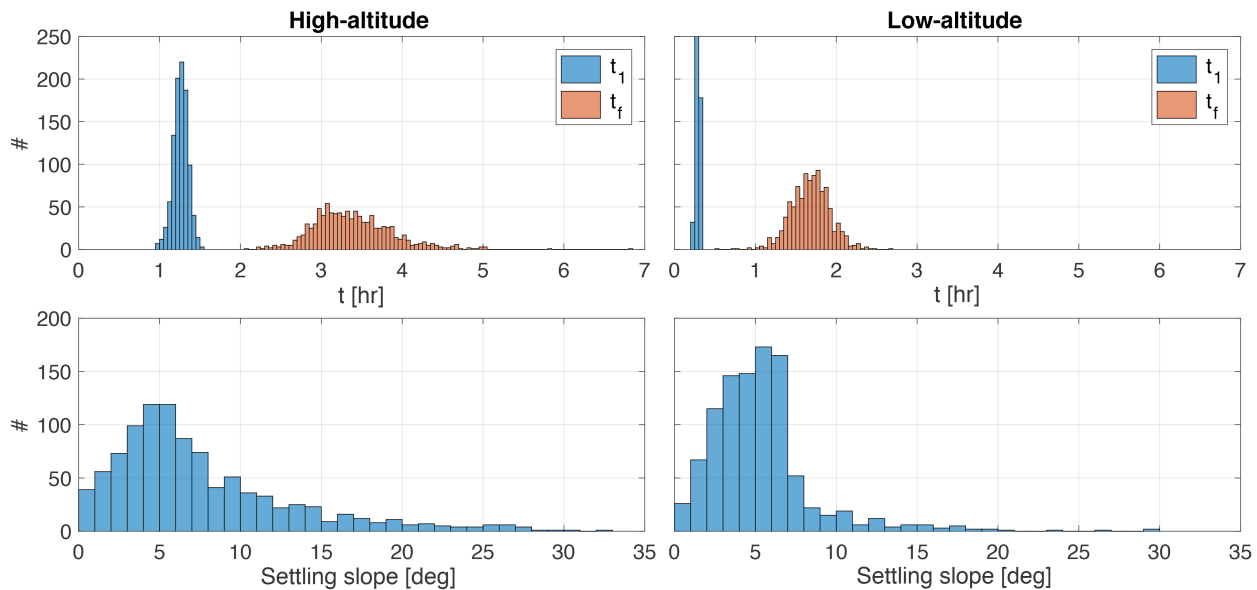


Figure 4.4: Statistics of nominal Itokawa deployments.

Nevertheless, significant scattering of the settling positions across Itokawa’s surface is observed; this illustrates the chaotic nature of motion near/on the small-body surface. The first impact time t_1 and settling time t_f distributions shown on the left side of Fig. 4.4 tend towards a Gaussian distribution. The first impact times are more clustered, *i.e.*, that they have a smaller standard deviation than the settling time. Finally, the settling slope of the simulated trajectories is measured. This slope is defined as the angle between the net acceleration on the probe, as given by Eq. 2.2, and the surface normal $\hat{\mathbf{N}}$ of the closest facet to which the probe comes to rest. If the probe comes to rest while in contact with two facets, the average of the two corresponding surface normals is used. Note that this may be different from the ‘effective’ settling slope at which a probe rests, which is created through the local geometry of the surface features, including rocks, that the probe is touching. The settling slope distribution shows that most probes come to rest on slopes less than 15° , though some settle in higher-slope regions.

Low-altitude deployment In the low-altitude deployment scenario, probes are released approximately 70 m above Itokawa’s rotation axis (≈ 0.35 body radii); this mimics release from a hovering mothership. The relevant initial parameters are included in Table 4.3; the results can be found in Figs. 4.2 through 4.4. Note that since the settling positions are clustered around Itokawa’s north pole, a standard latitude-longitude (ϕ, λ) plot is visually poor. Instead, an alternate definition (ϕ^*, λ^*) centered around the Z -axis (north pole) is used. As expected, the low-altitude simulations settle on the surface much quicker than the high-altitude scenario, which is visible in the statistics of Fig. 4.4. As a result, the probes’ range of motion is much smaller, such that their settling locations are more densely clustered, which is clear from Fig. 4.3. Since the nominal release point lies above the *Sagamihara* low-slope, low-geopotential region, the probes’ settling positions remain restricted to this region. The settling slopes are correspondingly small, particularly in comparison to the high-altitude deployment. The settling time and slope distributions shown in Fig. 4.4 are unimodal, with the settling time distributions resembling a Gaussian and the settling slope distributions more similar to a gamma distribution. For conciseness, standard boxplots are therefore used to represent these distributions in the following parameter variations.

4.1.2 Effect of surface interaction coefficients

The applied target-spacecraft contact model accounts for forces and torques generated by restitution, friction, and rolling resistance. The magnitude of these interactions is chiefly determined by their respective coefficient, *i.e.*, by e , f , and C_{RR} . This simple model is an abstraction of the much more complex interactions that occur between the probe and the granular, regolith-covered small-body surface. In order to investigate precisely how the deployment statistics are affected by these coefficients, the two nominal scenarios are repeated while independently varying the three coefficients. The provided results include the settling time and settling slope statistics, collated in respectively Figs. 4.7 and 4.8, in which the high-altitude results are shown solid and the low-altitude results are transparent. Heat maps of the settling positions are also provided in Figs. 4.9 through 4.11, in which the high-altitude scenario is shown on the left and the low-altitude results are shown on the right. These results are now reviewed.

Coefficient of restitution The nominal high- and low-altitude deployment scenarios are repeated with coefficient of restitution values $e = [0.35; 0.50; 0.65; 0.80]$. Inspecting the settling time statistics of Fig. 4.7, it is clear that the coefficient of restitution has a strong effect on the probe dynamics. More specifically, high e values are found to correspond to long settling times and vice versa. This is true for both the high- and low-altitude scenarios; both show variations in the mean settling time of up to 100% within the range of tested e values. This strong effect is expected: the coefficient of restitution dictates the amount of normal velocity damping during collisions, with lower e values indicating that more energy is dissipated in each collision, such that probes settle more quickly. Before reviewing the settling slope statistics in Fig. 4.8, consider the heat maps of Fig. 4.9. These maps show a stronger clustering of settling positions when e is decreased. This matches the settling time statistics: if probes settle faster, their range of motion will be smaller, such that the settling locations are more clustered. In both Itokawa scenarios, the probes are drawn towards regions of low geopotential.




Figure 4.5: Settling time statistics for varying surface interaction coefficients.




Figure 4.6: Settling slope statistics for varying surface interaction coefficients.

In the high-altitude scenario, this region is Muses Sea. When e is low, probes have a limited range of motion, such that few simulations actually reach Muses Sea; the majority settle just north of the region. When e is increased, the probe range of motion follows and simulations are drawn into Muses Sea. In Fig. 4.9, this behavior is reflected in the bright, densely-clustered spots of settling locations, which move south into Muses Sea when e increases. For high e values, this clustering is accompanied by significant scattering across the rest of Itokawa's surface. Similar trends are present in the low-altitude scenario, although the clustering is consistently denser than in the high-altitude scenario, as probes deployed from lower altitude naturally have a smaller range of motion. The settling positions remain constrained to the Sagamihara low-potential region; their scattering

increases with increasing e . These heat maps provide a useful background when inspecting the settling slope statistics of Fig. 4.8. In it, the mean settling slope increases when e is decreased. This result matches the settling position trends: when e is low, the probe range of motion is small and few simulations reach the relatively flat Muses Sea. They instead settle (often against rocks) in higher-sloped regions. When e is increased, the range of motion increases, and landers are drawn towards the flat Muses Sea. This explains why the mean settling slope decreases when e increases. These results show the strong effect of the coefficient of restitution on probe deployment statistics.

Coefficients of rolling resistance and friction The coefficient of rolling resistance C_{RR} governs the rate of angular velocity and rotational energy dissipated during collisions and contact. Reviewing Figs. 4.7 through 4.10 for the effects of $C_{RR} = [0.018; 0.035; 0.070]$ variations, it can be seen that they are similar to those observed when varying e , although the relationship is inverse. Indeed, higher C_{RR} values are found to lead to faster landings and correspondingly more densely clustered settling positions and higher settling slopes. These results match our expectations: when C_{RR} is higher, the rate of rotational energy dissipation is higher, such that probes settle faster.

The Coulomb friction force and torque drive the linear and angular lander velocity into a no-slip state where $\mathbf{v} = \boldsymbol{\omega} \times \mathbf{r}$. The magnitude of f dictates how quickly this synchronization occurs: at low f , synchronization will take longer to occur than at high f . Nonetheless, when reviewing the results of $f = [0.2; 1.0; 2.5]$ variation in Figs. 4.7, 4.8, and 4.11, no distinguishable difference between the low, intermediate, and high values of f are found. This indicates that even low $f \approx 0.2$ is sufficient to achieve stick in all collisions.

Surface heterogeneities In the nominal deployment scenarios, the interaction coefficients and rock distribution are assumed constant across the target surface. However, this assumption does not necessarily hold. For example, the Muses Sea and Sagami-hara low-geopotential regions on asteroid Itokawa were observed to be significantly smoother than the rest of the surface and covered with fine regolith. Given this dichotomy in surface composition, it is reasonable to expect Muses Sea and Sagami-hara to interact with a lander differently than other, more rocky regions.

Here, it is investigated whether surface heterogeneities such as those encountered on Muses




Figure 4.7: Settling time statistics for varying surface interaction coefficients.




Figure 4.8: Settling slope statistics for varying surface interaction coefficients.




Figure 4.9: Heat map of settling locations for varying e .

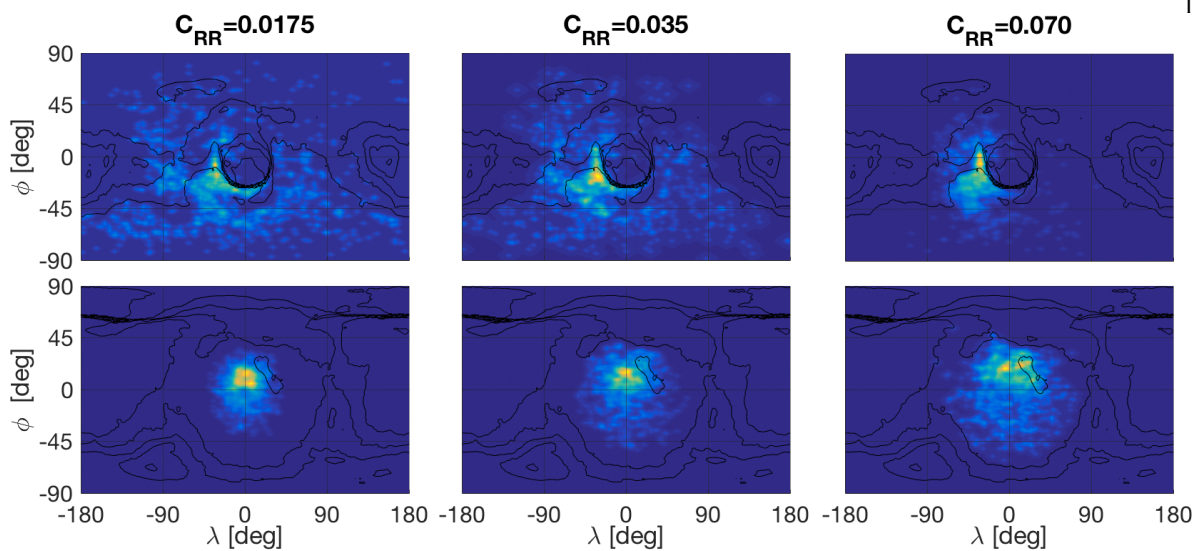


Figure 4.10: Heat map of settling locations for varying C_{RR} .

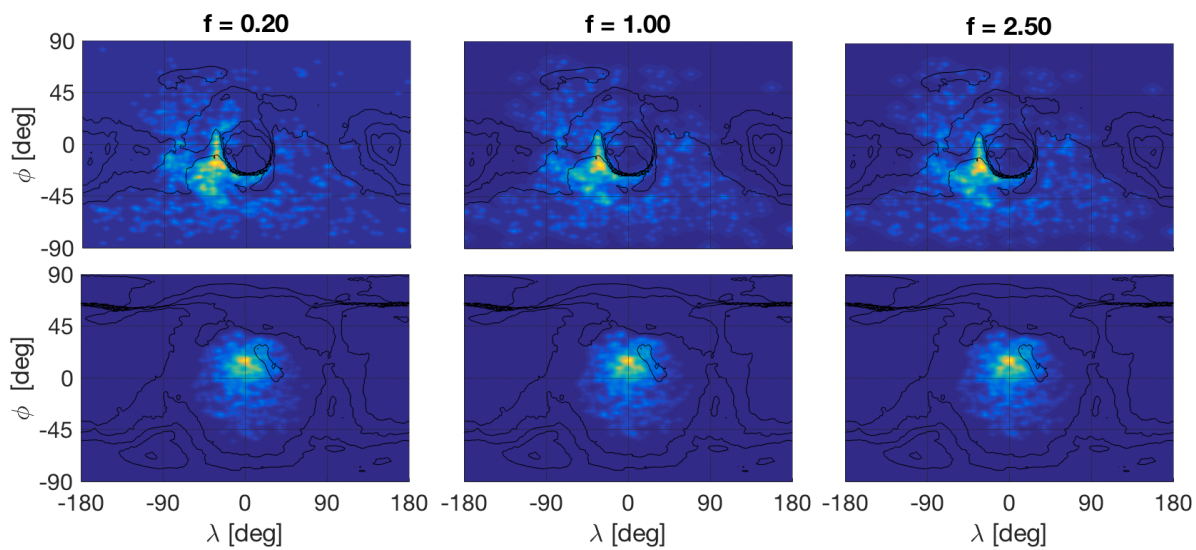


Figure 4.11: Heat map of settling locations for varying f .




Figure 4.12: Shaded Muses Sea and Sagamihara regions on Itokawa.

Sea affect the statistics of probe deployment. Using the surface geopotential, it is straightforward to select the Muses Sea and Sagamihara regions, as marked in Fig. 4.12. The two regions are given a coefficient of restitution of $e = 0.35$ and a rock initialization number of $k_0 = 0.625 \times 10^{-3}$; the latter indicates a rock density a quarter that of the nominal scenarios. The high-altitude deployment scenario is then repeated; the results are shown in Figs. 4.13 and 4.14. It is stressed that the intention in doing this is not to model the *actual* heterogeneities on Itokawa, but rather to investigate the effects, if any, of the presence of heterogeneous areas on probe deployment.




Figure 4.13: (top) Settling time and (bottom) settling angle on nominal and heterogeneous Itokawa surfaces.




Figure 4.14: Settling position heat maps of nominal and heterogeneous Itokawa surfaces.

Reviewing the settling time distribution of Fig. 4.13, it is found that probes settle slightly faster in the heterogeneous scenario than in the nominal scenario. This is expected: it was shown earlier that the high-altitude scenario draws probes into the Muses Sea region. As this region is given a lower coefficient of restitution in the heterogeneous scenario, probes therefore settle faster. Due to this lower e value, probes are also more likely to settle in Muses Sea, as is obvious from Fig. 4.14. This is true despite our earlier observations that shorter settling times generally result in probes settling before reaching Muses Sea. It is easy to see why Fig. 4.13 makes sense regardless: the lower e values that result in shorter settling times are present only in Muses Sea. In other words, probes have a relatively large range of motion across most of Itokawa's surface, but a much smaller range of motion in Muses Sea. Once they reach the region, they are therefore unlikely to escape. This is further supported by the settling slope distributions: the mean settling slope of the heterogeneous scenario is lower, matching the geography of Muses Sea. These results show that local variations in the surface interaction coefficients and rocks can indeed affect lander dynamics, and may change the degree to which certain regions act as a basin of attraction or rejection.

Discussion: Surface interaction coefficients Our simulations show that the coefficient of restitution is by far the most significant determinant of trends in the deployment of spherical probes, with higher energy dissipation rates resulting in faster landing with smaller surface dispersion. These results suggest that, at a minimum, it is important to obtain an accurate estimate of the e value on the surface of a considered target. Deviation from the true value may result in simulation trends that are biased in terms of settling time, settling slope, and surface dispersion. Furthermore, it was seen that local variations in the coefficient of restitution may transform regions into basins of rejection or attraction. One possible approach to account for these regional variations would be to use visual observations of the small-body surface to estimate the e distribution across the surface, based on the surface brightness.

Although the coefficient of rolling resistance has a similar effect on the settling statistics, its influence is smaller. This suggests that, although it is relevant to obtain an accurate C_{RR} estimate, priority should be given to studies of the coefficient of restitution. De Blasio and Saeter

experimentally measured rolling resistance coefficient values for small beads on Earth [27]. Van wal et al. have performed similar experiments for spherical objects in the size range of small-body landers, rolling on a granular bed [130].

On a final note, it is repeated that no distinguishable effects on deployment statistics (of these *spherical* probes) were found when varying the coefficient of friction f . This suggests that even relatively small f values are sufficient to drive the tangential contact point velocity into stick where $\mathbf{v} = \boldsymbol{\omega} \times \mathbf{r}$. It thus appears that, with regard to deployment of a spherical probe, little to no priority should be given to studies of friction on the small-body surface.

4.1.3 Effect of surface rocks

The discussion of the persistent polyhedral rocks model in Section 2.3 provided various details on the advantages that such a model offers over simpler stochastic models. In their work, Tardivel et al. performed a simple comparison in which a single deployment scenario was performed both with and without their stochastic rock model. They found that between the two sets *“the basins of the final resting positions are completely different”* and concluded that *“the presence of small rocks for surface motion on an asteroid [...] cannot be neglected”* [118]. It will be verified whether this claim holds when using a persistent rock model. Furthermore, it will be investigated in more detail how the density of surface rocks affects probe deployment. This is of interest to mission planning, as an understanding of these effects may provide requirements on the accuracy of orbital observations by a mothership from which the rock distribution parameters are obtained. For example, if the particular number of rocks has little influence on deployment statistics, a coarse estimate of the rock field density may be sufficient.

One advantage of the persistent rock model over the simpler stochastic model is the capability to easily tweak the rock distribution parameters. In a stochastic model, this would require re-running numerous simple simulations, analyzing the statistics, and then using those to perform the full simulations. With the persistent model, one can simply adjust the value of the relevant parameters, proceed with simulation, and investigate the effect on probe dynamics, as done here.

Number of rocks To investigate how the number of surface rocks affects lander deployment, the k_0 parameter in Eq. 2.22 is simply varied. This determines the number of rocks K per unit surface area; the two deployment scenarios are repeated for different values of $K = [4; 7; 14; 28; 56; 86]$ rocks/m². Fig. 4.15 illustrates what the small-body surface looks like for the tested K values; the spherical probe with $R = 0.125$ m is included for scale. The statistics of deployment for these cases can be found in Figs. 4.16 through 4.18.




Figure 4.15: Small-body surface for varying number of rocks.




Figure 4.16: Settling time and slope statistics for varying K .




Figure 4.17: Heat map of settling locations for varying K in high-altitude Itokawa scenario.

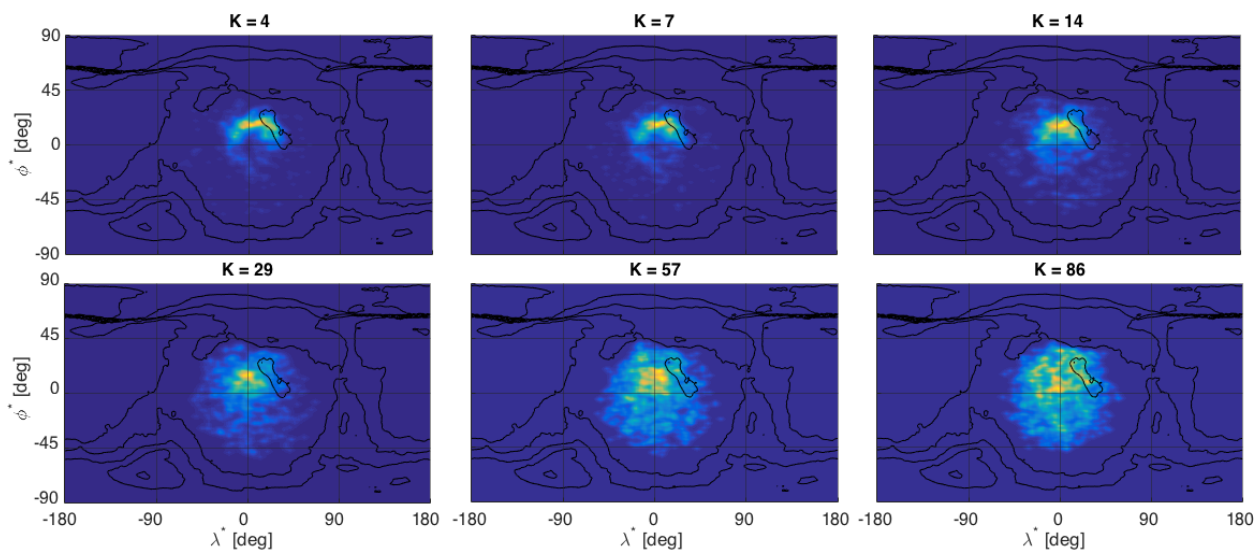


Figure 4.18: Heat map of settling locations for varying K in low-altitude Itokawa scenario.

Reviewing the settling time statistics, the presence of rocks on the small-body surface is indeed seen to have a strong effect on probe deployment, in agreement with Tardivel et al. [118]. More specifically, notice the settling time decreasing as the number of rocks, K , increases. This result may be explained using geometrical arguments: when a spherical probe performs repeated impacts on a smooth surface, each collision is performed relative to (approximately) the same surface normal. The restitution impulses that damp the probe's normal velocity therefore always act in the same direction. Any tangential velocity must be damped by rolling resistance, which dissipates

energy at much slower rates. When rocks are added onto the flat surface, successive collisions generally occur with different surface normals, such that restitution is able to dissipate velocity in different directions. This increase in dissipation is far greater than what can be dissipated by rolling resistance. The restitution impulse thus dissipates energy much faster on a ‘rough’ surface than on a smoother surface. As a result, an increasingly rocky surface yields shorter probe settling times, which is clearly visible in both the high- and low-altitude scenarios. This trend is present through a range of increasing K values, though a limit appears beyond which any further increase in K no longer affects the settling time statistics. Indeed, the results of Fig. 4.16 suggest that this limit lies around $K \sim 57$ rocks/m². Returning to the sample surface topographies of Fig. 4.15, it can be seen that this K value roughly corresponds to surface-rock saturation relative to the spherical probe size. In other words, at values beyond $K \sim 57$ rocks/m² for the applied rock model, there are no gaps in which the spherical probe can fit without touching any rocks. When K is small, the probe mostly impacts only the polyhedral small-body surface. As K increases, the probe impacts both the polyhedral surface and rocks on that surface. When K reaches its saturation limit, the probe can no longer impact the polyhedral surface, and instead always impacts a rock. Any further increase in K thus no longer effects any change, such that the settling time trends remain unchanged.

Inspecting the settling slope statistics of Fig. 4.16, it is found that it increases as the number of rocks is increased. This result makes sense, as shorter settling times were found to correspond to steeper settling slopes in our previous sensitivity analyses. This effect is likely reinforced by the fact that a persistent rock model allows probes to settle on steep slopes while resting against (a) rock(s), see also our previous illustration in Fig. 2.13. In agreement with the settling time results, a further increase in K beyond its saturation limit does not notably change the settling slope statistics.

Interestingly, reviewing the surface dispersions in Figs. 4.17 and 4.18, opposite trends are found when comparing the K variations in the high- and low-altitude scenarios. More specifically, it can be seen that for the high-altitude scenario, the lander dispersion *decreases* when K is increased. In the low-altitude scenario, the opposite occurs and the dispersion *increases* when K is increased. This difference is somewhat surprising, as the trends in settling time and settling slope statistics

do agree between the two scenarios. This behavior is explained as follows: in the high-altitude scenario, the first impacts occur at relatively high velocities, such that probes have a fairly large range of motion. The inclusion of surface rocks increases the energy dissipation rate, thus reducing the settling time and the probe range of motion. The surface dispersion thus decreases with higher K values. In contrast, probes in the low-altitude scenario impact the small-body surface with a much lower velocity that is mostly normal to the surface due to the relative flatness of the Sagami-hara region on Itokawa. The trajectories thus naturally show little tangential motion, as visible in the sample trajectory of Fig. 4.2. When rocks are added to the surface, they ‘perturb’ the smoothness of Sagami-hara, such that probe trajectories show much more tangential motion. This effect appears to outweigh the increased energy dissipation rate from the denser rock field, as the dissipation increases with growing K . In both scenarios, the saturation K value is again apparent, as the dissipation trends for $K = 57$ rocks/m² and $K = 86$ rocks/m² are very similar.

Rock interaction coefficients In the sensitivity analysis of the rock number K , both the small-body surface and the rocks on that surface were given the same interaction coefficients e , f , and C_{RR} . Although this is a useful assumption in the type of ‘mock’ scenarios used in our sensitivity analyses, it is unlikely to match the true topography of a target body. Furthermore, as mentioned in Section 2.3, there are various reasons why the regolithic surface and monolithic rocks may be expected to have different coefficients of restitution. Although the development of an accurate localized model is well beyond the scope of this paper, we do wish to obtain an estimate of what the effects of these local variations may be. As a test, the high- and low-altitude scenarios are therefore repeated with varying coefficient of restitution on the rocks, while maintaining the nominal e value on the polyhedral surface. For conciseness, only show the settling time variation is shown as result, in Fig. 4.19. Note that this plot merges the results for varying rock restitution with those presented earlier where e is varied globally. More specifically, the solid boxplots correspond to the *global* variation of e , whereas the transparent boxplots correspond to simulations where e is varied only on the rocks. The high- and low-altitude scenarios are shown on respectively the top and bottom of the figure.




Figure 4.19: Settling time statistics for global and rock-only variation of e .

The probe deployment trends observed when performing the rock-only variation of e match those seen when varying the coefficient globally. More specifically, higher e values correspond to longer settling times, in both the high- and low-altitude scenarios. The shift in settling time distribution is not as dramatic for the rock-only variation as it is for the global variation. These results make sense when returning to Fig. 4.15 and inspecting the surface samples with an intermediate number of rocks. In the corresponding simulations, probes alternate between impacting rocks and impacting the polyhedral surface. It is therefore expected that a rock-only e variation has less effect on the deployment statistics than a global variation. Nevertheless, these results show that differences in interaction coefficients between the polyhedral surface and its rocks may indeed have a significant effect on the motion of a probe. This illustrates that accurate maps of the interaction coefficient variation may be important in order to obtain accurate statistics on probe deployment.

Discussion: Rocks The results of our sensitivity analysis of the effects of surface rocks on lander deployment agrees with the (coarse) results of Tardivel et al. and confirm that the presence of surface rocks indeed strongly affects the motion of a probe. In our detailed investigations, it was found that if rocks are ignored in the surface modeling of a small body, the effective rate at which energy is dissipated along a spherical probe's deployment trajectory will be biased towards lower rates, due to reduced variability in the surface normal. As a result, the expected settling

time will be longer, with observed variations of up to 33% in two sample scenarios on asteroid Itokawa. Furthermore, a lack of surface rocks prevents probes from settling in high-slope regions, as previously illustrated in Fig. 2.13. As the small-body surface becomes more and more densely populated with rocks, probes are able to settle in increasingly steeper regions on that surface. This behavior continues up to a limit K value, where the surface becomes saturated with rocks relative to the size of the probe, *i.e.*, the probe can no longer ‘touch’ the surface but always impacts a rock. Although these trends appear in both the high- and low-altitude Itokawa deployment scenarios, the observed probe dispersion trends of these scenarios shows some differences. In the high-altitude scenario where probes have a large range of motion, the dispersion decreases as more rocks are added. In the low-altitude scenario where probes have a much smaller range of motion across a flat and slow-sloped region, the inclusion of rocks increases the probe surface dispersion. This suggests that the effects of surface rocks on probe dispersion are dependent on the considered deployment scenario and the topography of the targeted landing area.

Given that interaction heterogeneities may exist between the regolithic small-body surface and the more monolithic rocks on that surface, a simple sensitivity analysis was also performed, in which the coefficient of restitution, e , is given different values on the polyhedral surface and on its rocks. Our results show that such a dichotomy may indeed affect deployment statistics. Combining this with the ‘surface heterogeneity’ results presented earlier, this indicates that accurate maps of the variation in interaction coefficients and rock population across the small-body surface may be necessary in order to obtain accurate estimates of the lander deployment dynamics. For deployments in which probes have a small range of motion, for example after release from a mothership hovering at low altitude, such a variation may be small and unlikely to significantly affect the dynamics. However, for high-altitude release where probes have a large range of motion, localized variations in interaction coefficients and rock density will create basins of attraction or rejection, such that localized modeling may be required. Mothership observations of the surface brightness could be combined with computations of the surface slope and geopotential to this end. Nonetheless, the resolution to which such local variations should be determined remains an open question. Many

further details regarding the surface rock model may be considered, but go well beyond the modeling techniques applied in this paper. For example, Hayabusa observed clear trends in the orientation of rocks on certain regions of the Itokawa surface [73].

4.1.4 Summary: Deployment sensitivities in favorable environments

Slowly rotating bodies such as asteroid Itokawa present a favorable environment for the deployment of a scientific lander package. Its amended potential field is such that low-energy releases of a lander can be carried out from appreciably high altitudes that yield a minor risk to an orbiting mothership. Deployment is similarly possible from a mothership that is hovering inertially. For both scenarios, sensitivity analyses were performed to investigate how relevant environmental parameters that describe the probe-target system affect the dynamics of a spherical probe deployed to Itokawa. Investigating the effects of surface interaction coefficients, it was found that the coefficient of restitution, e , is the major determinant of lander deployment statistics. This suggests that, from a probe-deployment perspective, studies estimating the interaction properties of small-body surfaces should focus on developing accurate models of the coefficient of restitution. Although a similar argument can be made for the coefficient of rolling resistance, its effect is of lesser magnitude. No notable effects were found when varying the coefficient of friction, suggesting that little to no priority should be given to studies to determine its particular value.

Similarly, it was shown that the presence of rocks on the small-body surface significantly affects the motion of a lander, confirming preliminary results by Tardivel et al. Using our sensitivity analyses, it was found that the density of surface rocks, *i.e.*, the number of rocks per unit surface area, strongly correlates with the expected probe settling time. Probes settle faster when the small-body surface is more densely populated with rocks, as they effectively create a randomization of the surface normal. This increases the probe energy dissipation rate, resulting in faster settling. Similarly, the presence of rocks enables probes to settle on steeper slopes, as they effectively ‘get stuck’ between rocks. A saturated rock density appears, at which the small-body surface appears completely covered in rocks relative to the probe size. A further increase in rock density beyond

this saturation limit does not further affect the deployment statistics. Investigating heterogeneities in the coefficient of restitution and number of rocks, it was found that local variations can have a significant effect on the deployment statistics and appearance of any basins of attraction and rejection. These results indicate the value of constructing maps of the surface interaction coefficients and rock model across the small-body surface.

4.1.5 Deployment sensitivities in challenging environments

The above discussion revealed how the small-body surface properties affect the motion of a deployed probe. These sensitivity analyses were performed on asteroid Itokawa, whose dynamical environment is favorable for probe deployment. Both the high- and low-altitude scenarios are energetically guaranteed to settle on the small-body surface, without significant orbital motion around the body. When deploying to this kind of target, variations in the relevant environmental parameters simply increase or decrease deployment statistics such as the settling time and surface dispersion. Many other targets have dynamical environments that are far more challenging; consider for example the binary system 1999 KW₄. This system consists of a fast-spinning, diamond-shaped primary body and a smaller, tidally locked secondary. Although deployment to the secondary is feasible, the delivery of a probe to the surface of the primary body, known as *Alpha*, is challenging due to its fast rotation rate. Indeed, 1999 KW₄ Alpha's equatorial region rotates at near-orbital velocities, such that a small velocity increase is sufficient to (temporarily) launch a particle away from the surface. In such environments, naïve deployment strategies may fall short, and carefully designed strategies are necessary in order to successfully deploy a probe to the surface. Furthermore, as small velocity increments may significantly alter the success rate of deployment to these challenging targets, high deployment sensitivities to variations in the environmental parameters may be expected. These sensitivities will not only affect statistics such as the settling time, but will likely determine the overall success rate of a deployment strategy. In this section, simple analyses are used to investigate how to optimally deploy in a dynamically challenging environment, verify the criterion using simulations, and probe some of the deployment sensitivities. This is done

with asteroid 1999 KW₄ Alpha as target; some of Alpha’s characteristic properties are provided in Table 4.4 [85, 87]. Note that 1999 KW₄ is the only binary asteroid system to have been mapped at appreciable resolution, using radar imagery from an Earth flyby, which makes it an appropriate target for high-fidelity simulations using our methodology.

Table 4.4: Characteristic properties of asteroid 1999 KW₄ Alpha [85, 87].

Parameter	Value	Parameter	Value
Gravitational parameter μ	157.04 m ³ /s ²	Semi-major axis a	708.5 m
Density σ	1.974 g/cm ³	Semi-intermediate axis b	680.5 m
Rotation period T	2.7645 hr	Semi-minor axis c	591.5 m
Vertices in shape model n_P	4,586	Facets in shape model $n_{\mathcal{F}}$	9,168

Tools required to examine this problem The dynamical environment of 1999 KW₄ Alpha is indeed very different from Itokawa’s, due to its fast rotation and axisymmetric shape. As a result, the inner realm of the asteroid, *i.e.*, the region of space energetically bound to the asteroid, just barely extends beyond the surface of the body. 1999 KW₄ Alpha’s poles are in fact energetically “free” to escape the asteroid. Deploying from the inner realm, as done in the Itokawa scenarios, would thus require a mothership to descend to only a few dozen meters above the equator of the asteroid. Such proximity operations are often seen as hazardous to the mothership spacecraft and are therefore unlikely to be considered in a mission.

To further illustrate this challenging environment, consider a position on the equatorial ridge, where the surface rotational and circular orbital velocities are nearly equal. Using an equatorial radius of $R \approx 730$ m, a surface rotational velocity $V_s = \Omega R = 0.4607$ m/s and a surface circular velocity $V_c = \sqrt{\mu/R} = 0.4638$ m/s are found. In practice, this near-balancing causes the equatorial surface slopes to be considerably high and somewhat randomly oriented, pointing along the direction of the residual acceleration, generally north or south. This phenomenon makes it difficult, if not impossible, for an object to settle on 1999 KW₄ Alpha’s equatorial ridge. The surface slopes across 1999 KW₄ Alpha’s facets are shown in Fig. 4.20; these slopes were computed using a polyhedral gravity model. Relatively flat regions can be found just south of the equatorial ridge, which are promising for the deployment of a probe. This sub-equatorial region, which sits at a radius

of $R \approx 690$ m, indeed shows a wider gap between the surface rotational and orbital velocities $V_s = 0.4355$ m/s and $V_c = 0.4755$ m/s. These numbers, although computed with a point-mass gravity assumption, suggest that a particle is indeed much more likely to settle in the sub-equatorial region than it is on the equatorial ridge.




Figure 4.20: Surface slopes on 1999 KW₄ Alpha.

These topographical aspects are lost when using a simplified model such as a (smooth) sphere or ellipsoid. A sphere with the density and rotational period provided in Table 4.4 does not present a particular challenging environment: its equatorial region is well-suited for landing. This trend persists when using an ellipsoid. A simple spin-up of these bodies in order to mimic the difficulty of landing does not modify this situation; the centrifugal acceleration will ‘instantly’ overcome the the gravitational acceleration in the equatorial region, rather than showing distinct differences between the equatorial ridge and the rest of the surface. This is contrary to what is evidenced when inspecting the variation in slopes on 1999 KW₄ Alpha. Hence, a simplified model wrongfully concludes that landing is either straightforward or impossible.

1999 KW₄ Alpha presents a challenging target for landing precisely because of the combination of its bicone shape and fast rotation rate. Research has shown that this shape in fact results from the rotation rate [104]. Although certain models may yield insight into deployment to bodies such as 1999 KW₄ Alpha, simulation with a high-fidelity model is necessary in order to properly exploit 1999 KW₄ Alpha's details: the potential field irregularities, slope variations, and random surface asperities. These details determine whether a probe is grounded or returned into orbit.

Impact geometry Though not mentioned explicitly, our previous deployment results on asteroid Itokawa showed that a probe's first impact on the small-body surface is responsible for the majority of the energy dissipation along a trajectory. In order for a probe to successfully settle on the surface of 1999 KW₄ Alpha within some acceptable time, it is thus argued that the first impact must dissipate sufficient energy. Given 1999 KW₄ Alpha's challenging dynamical environment as discussed above, we wish to investigate which impact geometries maximize the first-impact energy dissipation, and test the resulting criterion numerically using our simulation framework. Although it was noted that simple, analytical models cannot uniquely determine if landing is possible, they may provide insight into the dynamics of a single impact and into what constitutes a better or worse deployment strategy. To illustrate what is considered as *impact geometry*, consider Fig. 4.21.




Figure 4.21: Velocity hodograph for impacts on a (*left*) flat and (*right*) rocky surface.

Both sides of this figure illustrate the geometry of a particle impacting a flat surface with (the same) *incoming* velocity \mathbf{V}_1 . The small-body surface rotates with some velocity \mathbf{V}_s , which defines the local horizontal. For fast rotators such as 1999 KW₄ Alpha, the surface rotational velocity V_s is only slightly smaller than the surface orbital velocity V_c . Inertial velocities are shown as solid lines, while surface-relative velocities are shown as dashed lines. During impact, the normal component of the *surface-relative* probe velocity is damped by restitution; effects from friction and rolling resistance are ignored. As the surface normal $\hat{\mathbf{N}}$ differs in the flat and rocky scenarios of Fig. 4.21, the energy dissipation in the two scenarios is expected to differ correspondingly. In other words, it is argued that the *outgoing* probe velocity \mathbf{V}_2 is dependent on both the initial release conditions as well as the local surface geometry at which the first impact occurs. This impact geometry is generalized for equatorial impacts to a spherical target with point-mass gravity in Fig. 4.22.




Figure 4.22: Geometry of first impact of a particle onto a rotating, spherical target.

In order to compute the energy dissipated during first impact of a particle, it is assumed that the particle is released at apoapse $r_A > R$ of some Keplerian orbit with periapse $r_P < R$ below the target surface. This orbit is illustrated on the left side of Fig. 4.22. For such an orbit, the true anomaly θ at impact may be computed as:

$$\theta = \arccos\left(\frac{a(1 - \epsilon^2) - R}{R}\right) \quad \text{with} \quad a = \frac{r_A + r_P}{2} \quad \text{and} \quad \epsilon = \frac{r_A - r_P}{r_A + r_P} \quad (4.2)$$

in which a and ϵ are respectively the semi-major axis and eccentricity of the release orbit. The flight path angle α is the angle between the *inertial* incoming particle velocity ${}^{\mathcal{N}}\mathbf{V}_1$ and the local surface horizontal. It can be computed using the ϵ and θ as:

$$\alpha = \arccos\left(\frac{1 + \epsilon \cos \theta}{\sqrt{1 + 2\epsilon \cos \theta + \epsilon^2}}\right) \quad (4.3)$$

in which, at impact, it must hold that $\alpha < 0$. The inertial incoming velocity magnitude ${}^{\mathcal{N}}V_1$ can be found from the orbital energy equation. Combining this with the flight path angle α , the inertial incoming velocity vector can be written as:

$${}^{\mathcal{N}}\mathbf{V}_1 = \sqrt{\mu\left(\frac{2}{R} - \frac{1}{a}\right)} \cdot \begin{bmatrix} \cos \alpha & \sin \alpha \end{bmatrix}^T \quad (4.4)$$

At the point of impact, the target surface is inclined at some angle β relative to the local horizontal. The corresponding surface normal and tangent directions may be written as:

$$\hat{\mathbf{N}} = \begin{bmatrix} -\sin \beta & \cos \beta \end{bmatrix}^T \quad \text{and} \quad \hat{\mathbf{T}} = \begin{bmatrix} \cos \beta & \sin \beta \end{bmatrix}^T \quad (4.5)$$

The *surface-relative* incoming particle velocity ${}^{\mathcal{S}}\mathbf{V}_1$ is computed using the inertial surface rotational velocity ${}^{\mathcal{N}}\mathbf{V}_s$ as:

$${}^{\mathcal{S}}\mathbf{V}_1 = {}^{\mathcal{N}}\mathbf{V}_1 - {}^{\mathcal{N}}\mathbf{V}_s \quad \text{with} \quad {}^{\mathcal{N}}\mathbf{V}_s = \begin{bmatrix} \Omega R & 0 \end{bmatrix}^T \quad (4.6)$$

During impact, restitution damps the normal component of this velocity, such that the surface-relative *outgoing* particle velocity is equal to:

$${}^{\mathcal{S}}\mathbf{V}_2 = \left({}^{\mathcal{S}}\mathbf{V}_1^T \hat{\mathbf{T}}\right) \hat{\mathbf{T}} - e \left({}^{\mathcal{S}}\mathbf{V}_1^T \hat{\mathbf{N}}\right) \hat{\mathbf{N}} \quad (4.7)$$

in which e is the coefficient of restitution. It is assumed that the settling time and surface dispersion of the particle are minimized when ${}^{\mathcal{S}}V_2$ is minimized. To gain insight into what release conditions minimize this velocity, the surface inclination β is fixed. The velocity ${}^{\mathcal{S}}V_2$ is then computed while varying the periape and apoapse radii of the release orbit. This analysis is performed using the parameters of 1999 KW₄ Alpha, for which R is set at $R = 690$ m; this is the approximate radius of 1999 KW₄ Alpha's sub-equatorial, low-sloped region. Contour plots show the corresponding ${}^{\mathcal{S}}V_2$

values for surface inclination values of $\beta = [0; 30; 45]$ deg in Fig. 4.23. These results show that when the surface is not inclined, *i.e.*, when $\beta = 0$ deg, ${}^S V_2$ is always minimized when $r_P = R$, regardless of the orbit apoapse. When β increases, an optimum $0.95R < r_{P,min} < 1.0R$ appears. Although the particular $r_{P,min}$ value shows a slight dependence on r_A and β , the variation is minimal. This suggests that, regardless of the local surface properties and release orbit apoapse, the orbit periapse should be set just beneath the target surface. This agrees with the results of [117]. To gain further insight into these results, the apoapse is fixed at $r_A = 2.0R$; ${}^S V_2$ is computed while varying the periapse radius r_P and local surface inclination β . The results of this variation are shown in Fig. 4.24. The region where ${}^S \mathbf{V}_1^T \hat{\mathbf{N}} > 0$ in the bottom right corner of the figure has been cut off, as those impact geometries are not physically possible.

Two optimum lines can be found in this plot. The first optimum line $\beta(r_P)$, shown with a dashed line, indicates the surface inclination β that maximizes the energy dissipation for each value of r_P . The bottom of Fig. 4.24 plots the corresponding flight path angle α and the impact angle η (defined in Fig. 4.22 as the angle between $\hat{\mathbf{N}}$ and ${}^S \mathbf{V}_1$). All optima $\beta(r_P)$ have an impact angle $\eta = 0$ deg, indicating that these optima correspond to normal impacts. These results make intuitive sense: by choosing some apoapse and periapse radii, the incoming particle velocity ${}^S \mathbf{V}_1$ is fixed. Given that normal restitution is the only mechanism for energy dissipation, the maximum dissipation then occurs when the surface inclination β is such that ${}^S \mathbf{V}_1$ is normal to $\hat{\mathbf{N}}$. Although friction was not included in this simple analytical model, it was previously noted that the energy dissipation contribution of friction is much smaller than that of restitution (for a spherical probe, if the coefficient of restitution is not particularly large); it is thus argued that these results do not change when including friction. This optimum line provides useful guidance in the design of deployments where the periapse radius r_P is constrained by, *e.g.*, operational and hardware considerations. In that situation, the $\beta(r_P)$ line can be used to target surface regions with a local inclination that yields maximum energy dissipation. A second optimum line $r_P(\beta)$, shown with a solid line, indicates the periapse radius r_P that maximizes the energy dissipation for each β value. In agreement with the earlier results of Fig. 4.23, it is found that these optima all occur for periapse




Figure 4.23: Outgoing surface-relative particle velocity for varying surface inclinations.




Figure 4.24: Outgoing surface-relative particle velocity for varying periapse and surface inclination.

values of $0.95R < r_P < 1.0R$. The corresponding α and η angles are plotted on the right side of the figure. This optimum line provides slightly different guidance: if operational considerations constrain the first-impact region, the $r_P(\beta)$ line provides the corresponding periapse radius that should be used in order to maximize energy dissipation on that first impact.

Combining these results, it is seen that the highest energy dissipation occurs when both r_P and β are high. This suggests that probes should be released on orbits with a periapse just beneath the target surface, with the first impact occurring in regions that rise up vertically relative to the local surface horizontal and against the surface velocity. Physically, this corresponds to probes impacting a steep cliff or large rock. Landing on fast rotators would thus benefit from the presence

of surface rocks and boulders. Nevertheless, the dissipated energy remains nearly constant when β decreases from 90 deg to 0 deg. This suggests that a high periapse value is the most substantial element in guaranteeing high energy dissipation; the presence of rocks plays only a minor role at such periapse values. Although it is preferable to release probes on a descent orbit with $r_P \rightarrow 1.0R$, such a release is unlikely to be carried out on an actual space mission. Uncertainties in the probe release mechanism result in errors in the apoapse velocity of a probe, such that a release with nominal $r_P \sim 1.0R$ may end up with a periapse *outside* of the target surface, resulting in failed deployment. Probe release is therefore likely to be performed with lower periapse radii; in such cases, the local surface inclination has a stronger effect on the dissipated energy. The optimum $\beta(r_P)$ line should then be used in the release planning and targeting procedures. Note that when the periapse radius is decreased, regions with low surface inclination become more favorable for landing. In other words, when deploying probes on an orbit with a lower periapse, the presence of rocks (which may cause high β) may be *detrimental* to the success rate of probe deployment. This suggests that the use of accurate release mechanisms is beneficial, as it allows for release orbits with higher periapse radii that dissipate more energy and are much less sensitive to the local surface topography including rocks. The observed trends remain near-constant when the apoapse radius is increased; the ${}^S V_2$ values will simply be higher as probes are released with greater initial energy. When varying the coefficient of interaction e , as illustrated in Fig. 4.25, the trends do show notable change.

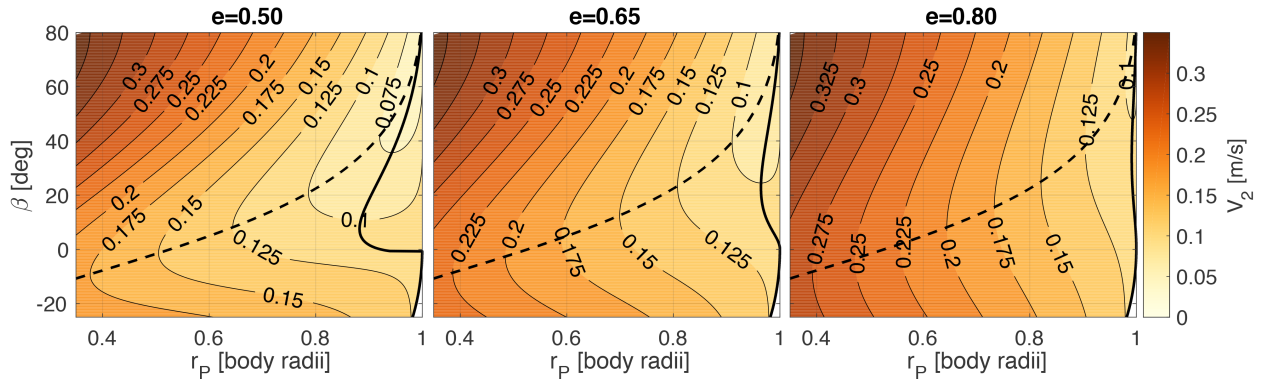


Figure 4.25: Outgoing surface-relative particle velocity for varying coefficient of restitution.

In addition to the obvious result that lower e values result in greater energy dissipation, it is found that the $r_P(\beta)$ line shifts to lower r_P values when e decreases. In other words, when the coefficient of restitution decreases, the optimum periapse also decreases. This makes it more likely for a release mechanism to successfully target that optimum r_P without risking $r_P > 1$ with release uncertainties. It is thus useful to target regions expected to have a low coefficient of restitution and to use probe hardware designs that minimize the coefficient, as done in the target marker design for Hayabusa and Hayabusa2 [100]. On a final note, notice that the optimum $\beta(r_P)$ line remains unchanged with varying e ; this must be true as this optimum line is based purely on geometrical considerations that are not affected by e .

Simulations The above analysis yields insight into the energy dissipation of a particle impacting a fast-spinning asteroid. Although these results may be used as guidelines in the design of a probe deployment strategy, the analysis was developed using a strong assumption whose validity has yet to be tested: it was claimed that trends in the first-impact energy dissipation of a particle moving in a central gravity field provide predictions about the full trajectory of a probe settling on a small-body. In order to validate this assumption, probe deployment to asteroid 1999 KW₄ Alpha is simulated using the methodology discussed in Chapter 2. Most importantly, the claim is tested that high release-orbit periapses maximize energy dissipation, due to near-normal first impacts on the small-body surface, resulting in faster and more successful deployment. These tests are initially performed using a ‘smooth’ 1999 KW₄ Alpha, and then repeated on a rocky version to verify if the presence of rocks is indeed detrimental to lander deployment. Overall, our analytic claims made above match those made by Tardivel et al. [117]. The main goal of the subsequent simulations is therefore to validate whether these predictions continue to hold when applied to the simulation framework discussed in this thesis, since this framework uses more advanced modeling methods than the work by Tardivel et al.

Periapse radius Our deployment analysis suggests that probe release orbits with a low periapse have a high initial energy and fail to dissipate it quickly enough. More specifically, low-periapse orbits impact the surface at relatively shallow flight path angles, such that the post-impact

tangential probe velocity remains quite high. The probe range of motion will be high, allowing them to continue moving a significant distance across the small-body surface. To test this claim, probe deployment to 1999 KW₄ Alpha is simulated, following release along the positive x -axis with an apoapse radius $r_A = 2.0R$, again with $R = 690$ m and using 1999 KW₄ Alpha's gravitational parameter $\mu = 157.04 \text{ m}^3/\text{s}^2$. In all simulations, effects from the small secondary body are ignored, to avoid the time-variability this body introduces. Four sets of simulations are performed, with periapse radii of $r_P = [0.05, 0.70, 0.90, 1.00]R$. The low-periapse orbit with $r_P = 0.05R$ is released with an inclination of $i = 25$ deg to target the equatorial ridge; all other releases are performed with $i = 0$ deg and target the low-sloped, sub-equatorial region mentioned before. The exact position, velocity, and corresponding uncertainties at release can be found in Table 4.5. Three hundred simulations are performed for each of the four considered deployment scenarios. Fig. 4.26 shows three-dimensional plots of 10 arbitrary trajectories from each scenario, in which 1999 KW₄ Alpha's surface is colored according to its local surface slope, as previously shown in Fig. 4.20. The release, first impact, and final positions of the four scenarios are also marked on an equal-area Mollweide projection in Fig. 4.27.

Table 4.5: Deployment parameters of 1999 KW₄ Alpha scenarios.

Release periapse r_P	$0.05R$	$0.70R$	$0.90R$	$1.00R$
Release position \mathbf{x} [m]	[1380, 0, 0]	[1380, 0, 0]	[1380, 0, 0]	[1380, 0, 0]
Position 3σ uncertainty $\Delta\mathbf{x}$ [m]	[15, 15, 15]	[15, 15, 15]	[15, 15, 15]	[15, 15, 15]
Release velocity \mathbf{v} [m/s]	[0, -0.8034, 0.0315]	[0, -0.6280, 0]	[0, -0.6052, 0]	[0, -0.59550]
Velocity 3σ uncertainty $\Delta\mathbf{v}$ [m/s]	[0.01, 0.01, 0.01]	[0.01, 0.01, 0.01]	[0.01, 0.01, 0.01]	[0.01, 0.01, 0.01]

Inspecting the results of the low-periapse ($r_P = 0.05R$) release on the top left of both Figs. 4.26 and 4.27, observe that the corresponding first impacts indeed occur on 1999 KW₄ Alpha's equatorial ridge, with a relatively shallow flight path angle of $\alpha \sim 45$ deg. As restitution is unable to damp the considerable tangential component of this velocity, the post-impact tangential velocity remains quite high. In agreement with our predictions, landers thus show considerable motion across the asteroid surface following first impact; some trajectories even encircle the entire body more than once. Note that in this scenario, many simulations failed to settle on the small-body

surface altogether, within the allowed 10 hr simulated time. Some of the final position markers in the $r_P = 0.05R$ plot are therefore projections of *orbital* positions onto the small-body surface.

Although some of the low-periapse trajectories do settle on 1999 KW₄ Alpha's surface within 10 hrs, this release is considered to be unacceptable for numerous reasons. Due to the large range of motion, settling region predictions cannot be performed with any relevant accuracy. It is therefore impossible to predict the expected illumination conditions of the landing site or to target specific regions on 1999 KW₄ Alpha's surface. Moreover, the observed 'circumnavigating' trajectories prevent the scheduling of mothership communication times. The tracking of a lander following this type of trajectory is unlikely to be possible, requiring an active search for the eventual landing site by the mothership. Finally, as settling times are (relatively) long, battery power on-board the probe is effectively wasted (on, *e.g.*, housekeeping and thermal operations) as most scientific instruments on-board the probe can only produce relevant results once settled on the surface. These arguments confirm our earlier claims that low-periapse lander release to fast-spinning targets results in 'failed' deployments.

The other scenarios with higher release periapses perform far better. Seeking to maximize the periapse radius while guaranteeing 'successful' first impacts, consider the $r_P = 0.90R$ scenario shown on the bottom left of Figs. 4.26 and 4.27. From the three-dimensional plot, it can be seen that trajectories first impact the asteroid with a slightly (in the target frame) retrograde velocity, though the impact is mostly normal to the local surface. As a result, restitution dissipates a relatively high amount of energy, such that little tangential velocity remains after impact. The probe range of motion thus remains small, with probes quickly settling following this first impact. The settling positions shown in Fig. 4.27 for this release thus show considerable clustering.

Given that (normal) restitution is the main energy dissipation mechanism, the $r_P = 0.70R$ scenario has been included, as probes released on such an orbit impact the surface more normally than in the $r_P = 0.90R$ scenario. Inspecting the results, it can be seen that the $r_P = 0.70R$ dispersion in probe settling locations is slightly increased relative to the $r_P = 0.90R$ scenario. Although impacts indeed occur more normally with $r_P = 0.70R$, the corresponding impact velocities




Figure 4.26: Sample deployment orbits to 1999 KW₄ Alpha with varying periaapse radius.

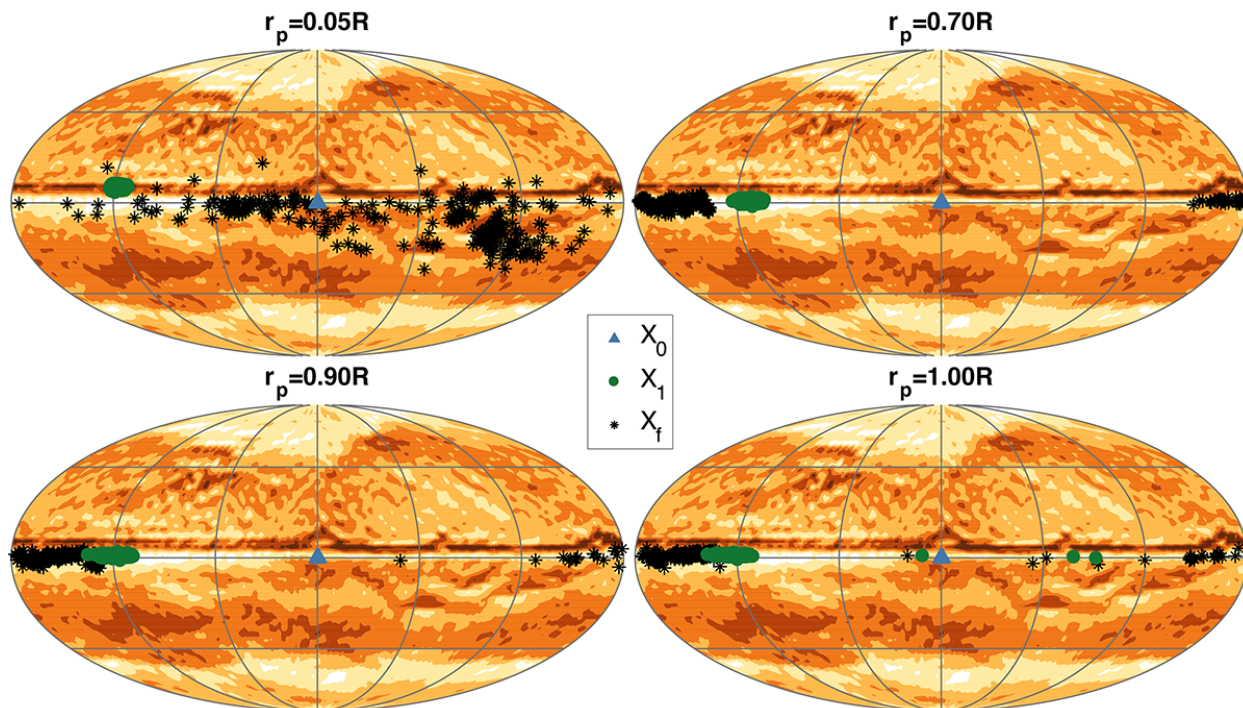


Figure 4.27: Surface dispersion of deployments to 1999 KW₄ Alpha with varying periaapse radius.

are higher, such that the post-impact probe range of motion is also higher. Finally, we consider the ‘optimum’ predicted by our analysis and test release on an orbit with $r_P = 1.00R$. Although the first-impact velocity of these trajectories is reduced even further, Fig. 4.26 shows that due to the uncertainties at release, some of the trajectories skim the surface and fail to impact the surface at the targeted sub-equatorial region. Although other trajectories do perform very well, this risk of ‘missed’ impact renders the $r_P = 1.00R$ release unsatisfactory.

These results are also reflected in the energy trends of the considered trajectories, as shown on the left side of Fig. 4.28. In this figure, boxplots are used to mark the energy at release (rightmost boxes) and at first impact (leftmost boxes) of the four deployment scenarios. Inspecting the figure, it is found that the low-periapse scenario releases probes with a significantly higher initial energy than all other scenarios. It is therefore intuitive to expect these trajectories to have a considerably higher range of motion and corresponding settling time. The three high-periapse trajectories have similar energies; a clear trend of lower energies for higher periapses is seen. This agrees with the earlier statement that although the $r_P = 0.70R$ trajectories impact the surface more normally than the $r_P = 0.90R$ trajectories, their velocities at first impact are higher, resulting in a larger post-impact range of motion. Finally, the angle η between the incoming (target-frame) probe velocity and the normal vector with respect to which the first impacts occur is also plotted, as shown on the right side of Fig. 38. The optimal scenario, with $r_P = 0.90R$, has probes impacting at relatively shallow angles of $\eta \sim 30$ deg, which roughly matches the optimum conditions of the solid line previously shown in Fig. 4.24.




Figure 4.28: Energy and impact angle for varying r_P deployments to 1999 KW₄ Alpha.

Although the release trajectories would likely benefit from further fine-tuning, in particular when taking into account effects from the secondary body KW₄ Beta (see also [117]), these results indicate that our simple, geometrical analysis of deployment to a fast-spinning body indeed provides relevant guidelines. Despite the analysis being performed for a single impact of a particle in a central gravity field, its conclusions extend to the global trajectory of a probe deployed to a small body. They show that proper release trajectory design is paramount in dynamically challenging environments such as 1999 KW₄ Alpha. A low-periapse orbit, which mimics slow release from an inertially hovering mothership, performs poorly with probes showing considerable motion around and across the asteroid surface. Release on a high-periapse orbit results in successful deployment with appreciable clustering of final settling positions. When choosing a high-periapse release orbit, it is important to take into account the uncertainties at release, to guarantee probes do not skim the surface and ‘miss’ the first impact altogether.

Presence of rocks During the earlier sensitivity analyses in Itokawa’s favorable dynamical environment, the effect of surface rocks on probe deployment was investigated. In such an environment, the (number) density of surface rocks was found to modify relevant statistics such as the probe settling time and surface dispersion. Given the landing challenges demonstrated on 1999 KW₄ Alpha, we are eager to investigate how the optimal, low-periapse release scenario is affected by surface rocks. The analytical model predicted that an increase in the local surface inclination β is somewhat detrimental to the energy dissipation in deployment trajectories with the optimum $r_P = 0.90R$. In order to verify this, the simulations are repeated while including rocks on the surface, testing rock density numbers of $K = [0, 14, 28]$ rocks/m². The resulting surface dispersions are plotted in Fig. 4.29; the measured energies and impact angles are shown in Fig. 4.30. Finally, Fig. 4.31 also shows scatter plots relating the impact angle η and the corresponding settling times.




Figure 4.29: Surface dispersion of deployments to 1999 KW₄ Alpha with varying rock densities.




Figure 4.30: Energy and impact angle for varying rock density deployments to 1999 KW₄ Alpha.




Figure 4.31: Settling time versus impact angle for varying rock density deployments to 1999 KW₄ Alpha.

Fig. 4.29 shows that rocks on 1999 KW₄ Alpha's surface indeed affect the probe surface dispersion. Although many trajectories with rocks perform similar to the rock-less scenario, in which probes settle due west of the first impact location, several others do show circumnavigatory behavior similar to that of the low-periapse scenario. Given that the energy at release is the same between the three scenarios, the observed differences must result from the impact geometry. Inspecting the

impact angle η as plotted on the right side of Fig. 4.30, the angle increases significantly as more rocks are added to the surface, although the mean angle remains unchanged at its corresponding smooth-surface value. This result agrees with the prediction of our analytical model that increases in the local surface inclination β deteriorates the first-impact energy dissipation. This occurs because trajectories with low initial energy have a high periapse and impact the (smooth) target surface normally. When rocks are added to the surface, the effective impact normal is disturbed, such that probes obtain an effective tangential velocity, reducing the energy dissipated on first impact. Furthermore, the local orientation of rocks may be such that probes are scattered in a north/south direction. This scattering may send the probes on circumnavigatory trajectories as shown earlier following low-periapse release. The settling time trends of Fig. 4.31 are also testament to this fact: when the number of rocks is increased, the effective impact angles spread over a wide range of values, and more trajectories are shown to have a maximum ‘settling time’ of 10 hrs, indicating that they failed to settle on the surface within this time.

These results indicate that the presence of surface rocks indeed affects the success rate of a deployment scenario at a challenging target such as 1999 KW₄ Alpha, confirming the results of [117]. They may increase the local surface inclination at first impact and may also scatter probes in a north/south direction that limits the dissipated energy and increases the probe range of motion. This underscores the importance of properly modeling the target small-body surface, and suggests the targeting of smooth surface regions when possible. On asteroid Itokawa, low-sloped regions were observed to be covered with fine regolith (possibly covering up any rocks present in the area). If the targeted low-slope, sub-equatorial region on 1999 KW₄ Alpha shows the same smooth terrain, it would make an excellent target for probe deployment from relatively high altitudes. If the region is instead found to be rocky, it may be necessary to lower the deployment apoapse in order to further reduce the probe energy at release and improve the deployment success rate.

Summary: Deployment sensitivities in challenging environments It was shown that fast-rotating bodies, such as 1999 KW₄ Alpha, are challenging targets for the deployment of a probe. Due to its fast rotation rate, low-energy deployments from equilibrium points of 1999

KW₄ Alpha's amended potential field require release from a few tens of meters above the surface; a risky operation for a mothership. Using a simple analytical model of the gravitational field and probe first-impact geometry, deployment from higher altitudes was investigated. Although release orbits with a low periaipse impact the surface at high velocities and fail to have probes settle within appreciable time, orbits with a high periaipse of $r_P \sim 0.90R$ are shown to be successful. Probes released on such an orbit impact the asteroid with a relatively low velocity that is approximately normal to the local asteroid surface, such that the reaction restitution impulse dissipates a maximum amount of energy. Through simulation, it was demonstrated that this simple, analytical model may successfully predict trends in the global motion of a probe. Although higher periapses $r_P \rightarrow 1.0R$ may theoretically further improve landing success, release uncertainties render such deployments infeasible from a practical point of view. The presence of rocks, which affects the local surface inclination, has a detrimental effect on deployment success, as it prevents probes from normally impacting the surface following release on a high-periaipse orbit. Our results show that while naïve release strategies may fail to successfully deliver probes to the small-body surface, with careful design it is possible to establish reliable, low-risk strategies that are far more successful. Our results reveal high sensitivities to deployment and environmental parameters, and underscore the need for high-fidelity modeling.

4.2 Complex landers

The above section has provided extensive insight into the motion of spherical probes deployed to small-body surfaces, both in favorable and more challenging environments. However, as discussed in Chapter 1, actual small-body surface exploration payloads can and do have more complex shapes. It is therefore highly relevant to extend the above studies to probes with other, non-spherical shapes. This is the focus of the current section; in which the motion of the five platonic solids will be compared and contrasted, as well as equivalents of the MINERVA-II-1, MINERVA-II-2, and MASCOT lander/rover probes onboard Hayabusa2. We also investigate how the surface interaction coefficients and internal probe mass distribution affect the probe dynamics and compare that with

the behavior observed for spherical probes. These results will then be combined to establish some simple recommendations for the physical design of small-body exploration probes.

It is, however, necessary to state a major caveat for the results discussed in this section. Whereas the motion of spherical objects is understood relatively well, with a reasonable basis of experimental work, this is not true for the bouncing motion of complex shapes. For such object, there is a significant number of parameters that affect the motion, several of which are not well understood. Most importantly, the results presented here pivot on the assumption of constant coefficients of restitution and friction. Lacking experimental evidence of rigid-body impacts in micro-gravity, this is a fair assumption to make. Nevertheless, the results of this section should be read with this assumption in mind, and a series of experiments to test the assumption will be provided further on. These experiments, in combination with further simulations in geometrically simpler environments (*e.g.* a flat plane) must be used to confirm or disprove the trends observed here.

4.2.1 Nominal deployment scenario

The nominal scenario used to analyze the motion of complex probes mimics the low-altitude Itokawa deployment scenario applied in Section 4.1, with probes being released approximately 70 m (~ 0.35 body radius) above Itokawa's northern pole. The applied release velocity and corresponding uncertainty can be found in Table 4.6. The probes are given a random initial attitude in each simulation. One hundred thousand simulations of a cube are performed, with mass $m = 5$ kg, uniform density $\rho = 500$ kg/m³, and coefficients of restitution and friction of $e = f = 0.5$. The parallelized GPU simulation framework with SDF shape model is applied to perform such high number of simulations in an appreciably short time. A single typical bouncing trajectory is illustrated in Fig. 4.32. Although a variety of relevant statistics may be studied, *e.g.*, the number of bounces, the energy dissipation per bounce, the velocity change in each bounce, *etc.*, we are primarily interested in the most pragmatic statistics: the settling time and the surface dispersion. These two statistics will now be examined.

Table 4.6: Nominal release parameters of complex lander deployment to Itokawa.

Parameter	Value
Release position \mathbf{x} [m]	[0, 0, 200]
Position 3σ uncertainty $\Delta\mathbf{x}$ [m]	[15, 15, 15]
Release velocity \mathbf{v} [m/s]	[0, 0, -0.05]
Velocity 3σ uncertainty $\Delta\mathbf{v}$ [m/s]	[0.01, 0.01, 0.01]




Figure 4.32: Typical deployment of a cube in the low-altitude Itokawa scenario.

Settling time The first impact time t_1 and final settling time t_f are easily extracted from the simulation data. These values are plotted for all 100,000 simulations in probability-normalized histograms on the left side of Fig. 4.33. Both distributions appear as Gaussian, such that the mean and standard deviation of the settling time provide a good overview of the full statistics. The corresponding probability distribution functions (PDFs) are also included on the right side of the figure. These were constructed using Matlab’s kernel density smoother `ksdensity` with default bandwidth. In our parametric investigations, both the PDF and mean/standard deviation of the settling time distributions will be used to compare and contrast various simulation batches.

Surface dispersion Fig. 4.34 visualizes the surface dispersion of the nominal deployment scenario. On the left side, the release, first impact, and final settling locations are marked with respectively red, green, and blue points. One sample trajectory is plotted with a red line. To

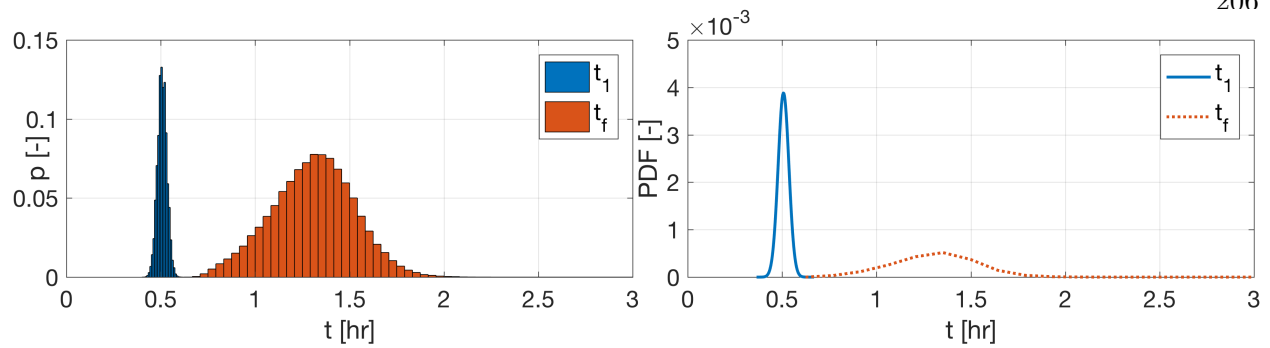


Figure 4.33: Settling time statistics of the nominal Itokawa deployment scenarios.

prevent the settling locations from saturating the surface, only 10% of the simulations are shown. The surface of asteroid Itokawa is shaded with respect to local slope, where dark areas have low slopes and bright areas have steep slopes.




Figure 4.34: Dispersion plots of the nominal cube deployment to Itokawa.

This figure illustrates the chaotic nature of the motion of a ballistically-deployed spacecraft bouncing on the surface of a small body. Although the settling positions remain fairly clustered around the first-impact region, notable scattering does exist; this scattering may increase significantly when the surface interaction coefficients or probe shape are modified. This chaotic scattering makes it difficult for the surface dispersion of a given scenario to be properly represented with a simple model such as a landing ellipse. This is especially true when multiple, distinct basins of attraction exist (as will be the case in some of the future simulations, see for example Fig. 4.38). A simple three-dimensional plot such as that on the left side of Fig. 4.34 also fails to quantify the degree to which the landing positions are clustered, in particular when many points are plotted.

We therefore follow the approach of [44] and use Matlab’s *multivariate* kernel density estimator `mvksdensity` to generate a three-dimensional PDF that captures the distribution of settling positions. Silverman’s rule of thumb is used for bandwidth selection. This PDF can be evaluated across Itokawa’s surface, as shown in the intensity plot in the middle of Fig. 4.34. Although this provides an excellent visual overview of the surface dispersion, it does not quantify the actual area across which landers disperse. In order to obtain such a quantification, an asteroid shape model with facets of approximately equal size is used. The PDF of the settling point distribution is evaluated at the center of each of its facets; each density value is multiplied with the area of the corresponding facet. We then accumulate facets, in order of descending density, until we have accounted for 68.27% of all simulations. This defines the 1σ dispersion area. The same process is used for the 2σ (95.45% of simulations) and the 3σ (99.73% of simulations) dispersion areas. For an easy global view of this settling position density, Fig. 4.35 also plots the dispersion areas on Mollweide plots. The Mollweide projection is an equal-area projection that is better suited to visualize global distributions across a surface than classical latitude-longitude maps that do not preserve area. Asteroid Itokawa has a strongly non-convex neck area that is poorly captured by the Mollweide plot, but since this area is high-sloped, landers do not settle in this region. Note that the projection cannot be used for bodies that are significantly more non-convex, such as comet 67P/C-G; in this case, it is more beneficial to stick with the three-dimensional plots of Fig. 4.34.




Figure 4.35: Mollweide projection of the dispersion plots of the nominal Itokawa deployment.

4.2.2 Effect of probe shape

The parallelized simulation methodology presented in Chapter 2 allows us to repeat a given scenario of many deployments of different probe shapes. When selecting a representative set of shapes to consider, it is natural to include the three lander/rover payloads included on the Hayabusa2 mission, since these are real-world examples of what a useful small-body surface exploration probe may look like. In addition to these, the five platonic solids are included: tetrahedron, octahedron, cube, dodecahedron, and icosahedron. These five shapes are regular, convex polyhedra with spherically symmetric inertia and a different number of faces (respectively 4, 6, 8, 12, and 20). They provide a reasonably continuous variation in geometry and mass distribution, from the nearly-spherical dodecahedron to the strongly warped tetrahedron. This yields a total of eight different shapes: five platonic solids with spherically symmetric inertias and three real-world analogues with a more ‘irregular’ shape. The actual mass and density of the Hayabusa2 payloads are ignored; instead, all eight shapes are given the same density $\rho = 500 \text{ g/cm}^3$ and mass $m = 5 \text{ kg}$; this fixes their sizes. Effectively, each of the considered shapes thus carries the same amount of ‘malleable’ payload. The eight shapes are shown at the same scale in Fig. 4.36.




Figure 4.36: The eight considered probe shapes.

The platonic solids can be defined in terms of their vertex radius r and side length s ; the latter is marked with a green line in Fig. 4.36. The ratio $l = r^2/s^2$ is tabulated in Table 4.7. Together with the shape parameter $v = V/s^3$, which can be thought of as a packing efficiency, it can be used to determine the shape size corresponding to some combination of m and ρ . Given these two parameters, the corresponding side length s can be computed as:

$$s = \sqrt[3]{\frac{m}{v\rho}} \quad (4.8)$$

By combining the computed s and the corresponding l from Table 4.7, the vertex radius r needed to set the proper shape size can be computed. In order to perform simulations, knowledge of the corresponding inertia $[\mathbf{I}]$ is also required. Given that the platonic solids have spherically symmetric inertias, they can be expressed in the following form:

$$[\mathbf{I}] = jms^2 \cdot \mathbf{I}_3 \quad (4.9)$$

in which $0 \leq j$ is the dimensionless *reduced inertia* of the shape, as introduced for the sphere in Section 4.1. Table 4.7 includes j values for both the solid (uniform mass distribution) and shell (with all mass concentrated at the faces) version of all five shapes. The l , v , and j values were obtained from [99].

Given that the shapes have different side lengths s , an examination of j by itself is not instructive when comparing their mass distributions. Therefore, Table 4.7 also includes the js^2 term that yields the inertia tensor in Eq. 4.9; this is the mass-normalized inertia. Comparing these values, notice that the dodecahedron and icosahedron have very similar js^2 values. The cube and octahedron also have similar mass-normalized inertias, that are approximately 7% greater than those of the dodecahedron/icosahedron. Finally, the tetrahedron has even higher js^2 ; it is approximately 35% greater than that of the dodecahedron/icosahedron. This shows that the dodecahedron/icosahedron have very similar mass distributions, as do the cube/octahedron. In contrast, the tetrahedron's mass is distributed considerably more outwards from its center of figure.

Finally, Table 4.7 also lists the *vertex defect* δ of the five solids. This parameter is obtained as the sum of all *interior angles* β_i at a vertex, subtracted from 2π :

$$\delta = 2\pi - \sum_i \beta_i \quad (4.10)$$

We provide an example calculation of the vertex defect of a cube in Fig. 4.37. For the platonic solids, all interior angles β_i have the same value. In the case of a cube, it holds that $\beta_1 = \beta_2 = \beta_3 = 90$ deg. From Eq. 4.10, it is thus found that the vertex defect $\delta = 90$ deg. This defect is directly related to the curvature of a shape at its vertices (see Chapter 3) and will be a useful concept in the examination of the simulation results further on.




Figure 4.37: The interior angles β_i used to determine the vertex defect δ .

Table 4.7: Relevant properties of tested symmetric lander shapes.

Shape	$n_{\mathcal{F}}$	$n_{\mathcal{V}}$	l [-]	v [-]	r [m]	Shell j [-]	Solid j [-]	Core j [-]	Solid $j s^2$ [m ²]	δ [deg]
Sphere	∞	∞	1.0000	4.1888	0.1337	0.6667	0.4000	0.0133	0.0071450	0
Dodecah.	12	20	1.9635	7.6631	0.1531	1.0123	0.6074	0.2025	0.0072528	36
Icosah.	20	12	0.9045	2.1817	0.1580	0.4363	0.2618	0.0873	0.0072240	60
Cube	6	8	0.7500	1.0000	0.1866	0.2778	0.1667	0.0556	0.0077360	90
Octah.	8	6	0.5000	0.4714	0.1957	0.1667	0.1000	0.0333	0.0076630	120
Tetrah.	4	4	0.3750	0.1179	0.2330	0.0833	0.0500	0.0167	0.0096548	180

Surface dispersion To investigate if the eight probe shapes behave differently, the nominal deployment scenario presented above is repeated. This results in 100,000 simulations with the same initial conditions, for each of the eight shapes. In Fig. 4.38, the surface dispersions of the eight simulation batches are plotted; corresponding Mollweide plots can be found in Fig. 4.39. The $1/2/3\sigma$ dispersion areas are also provided in Mollweide projection in Fig. 4.40 and are quantified

in Fig. 4.41. Of the five platonic solids, the tetrahedron has the smallest dispersion area, followed by the octahedron and the cube. The icosahedron and dodecahedron have the largest dispersions, though their values are similar. Both the MASCOT and MINERVA-II-2 probes display dispersions highly similar to that of the cube, while MINERVA-II-1 shows the greatest dispersion of all tested shapes. These trends will be further discussed below.

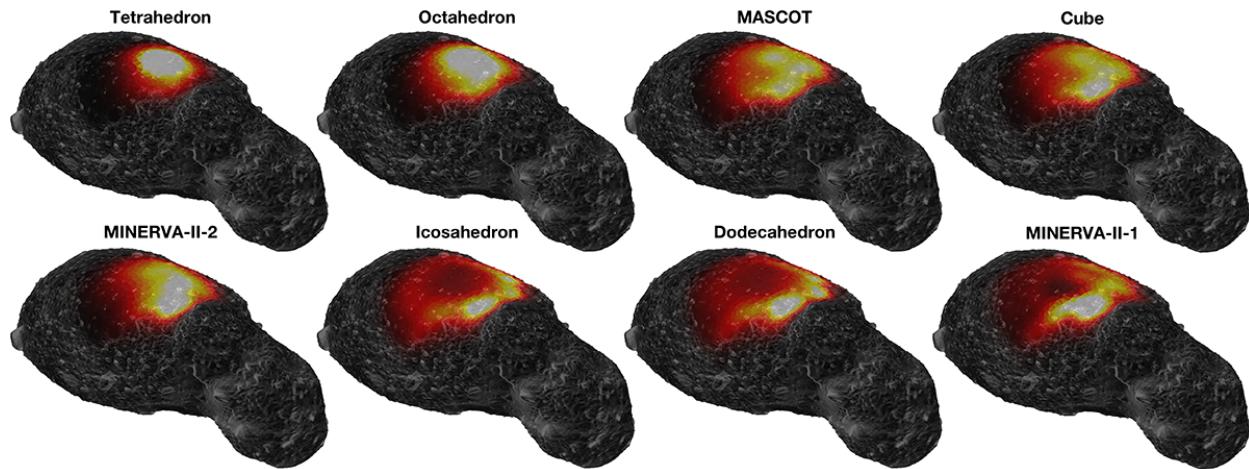


Figure 4.38: Settling position density of the considered shapes.

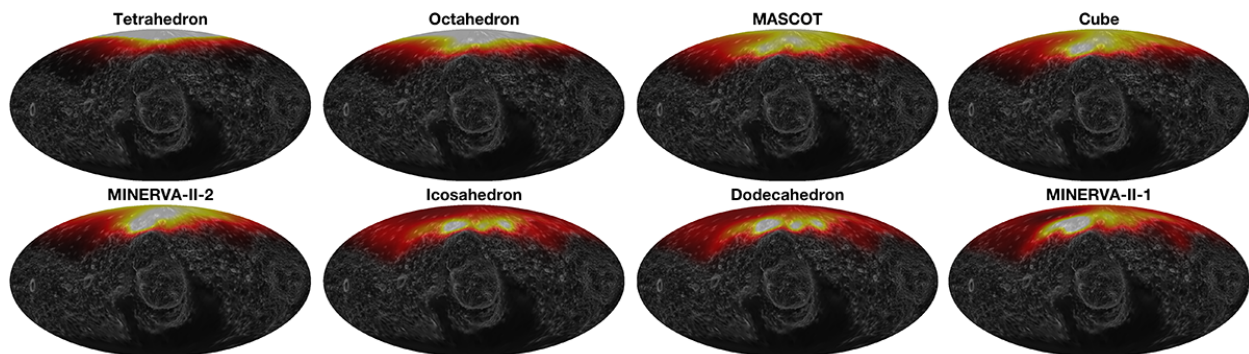


Figure 4.39: Mollweide projection of the settling position density of the considered shapes.




Figure 4.40: Mollweide projection of the dispersion areas of the considered shapes.




Figure 4.41: Dispersion statistics of the considered shapes.

Settling time The settling time PDFs of the eight deployments are provided in Fig. 4.42. At the top of the figure, all eight shapes are included in a single plot. The bottom left shows only the five platonic solids. The bottom right shows the three Hayabusa2 lander/rover payloads. Reviewing the distribution of the platonic solids, it is found that the tetrahedron settles the quickest, followed by the octahedron, cube, icosahedron, and finally the dodecahedron with the longest t_f . In contrast to the surface dispersion, the icosahedron and dodecahedron have distinctly different settling time distributions. Note that all five platonic t_f distributions appear mostly Gaussian, though the tetrahedron does show some positive skewness.

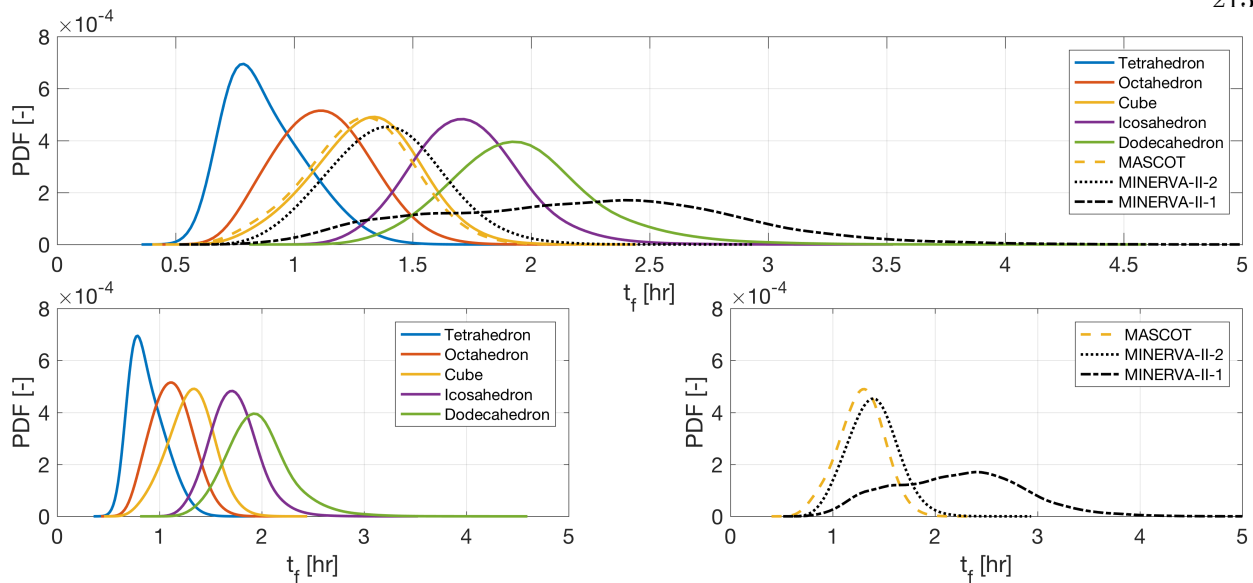


Figure 4.42: Settling time PDFs of the considered shapes.

Inspecting the t_f curves of the three Hayabusa2 payloads, it is found that the MASCOT lander has a Gaussian t_f distribution that is very similar to that of the cube; this makes intuitive sense given MASCOT's nearly cubical shape. The MINERVA-II-2 rover t_f distribution also appears as Gaussian; the rover settles only slightly slower than the cube. Finally, the MINERVA-II-1 rover is found to have the most peculiar t_f distribution: its PDF appears to reach a critical point at a t_f value close to the mean t_f of the cube. The curve continues to climb beyond this value and slopes out relatively gently, indicating that many simulations have higher t_f .

Discussion: The effect of shape It is obvious from the deployment statistics shown above that the shape of a small-body exploration probe strongly influences its dynamics. Our results suggest that a more outwards mass distribution (the tetrahedron) results in shorter t_f and smaller A_d than a more concentrated mass distribution (the dodecahedron). However, when comparing the js^2 values in Table 4.7 with the trends of Fig. 4.42, it is clear that inertia alone does *not* determine the settling behavior of a shape. This is clear from the observation that the dodecahedron/icosahedron and the cube/octahedron have very similar inertias, but display distinctly different settling behaviors. Furthermore, it is found that the t_f trends persist when the simulations are repeated while giving the five platonic solids identical inertia. This indicates

that the physical shape of a probe, through which it interacts with the small-body surface, has an effect on deployment dynamics that is distinct from the effect of the shape's corresponding mass distribution. To understand precisely how the probe shape affects its deployment dynamics, the collision geometry in a simple, two-dimensional setup is first investigated, as shown in Fig. 4.43.




Figure 4.43: The minimum and maximum aspect angle α of various regular, convex polygons.

This figure shows a triangle, cube, and pentagon (regular, convex polygons) in contact with a flat surface. For such convex shapes, collision with a flat surface must always occur at one of the corner vertices. For a given impact, the *aspect angle* α is defined as the angle between the surface normal $\hat{\mathbf{N}}$ and the opposite of the collision vector, $-\mathbf{r}$. This angle cannot vary arbitrarily but rather takes values between $-\alpha_{max} \leq \alpha \leq \alpha_{max}$. The angle is defined *ad hoc* to be positive when the shape is angled *towards* its linear velocity, *i.e.*, when $\mathbf{r} \cdot \mathbf{V} < 0$. Since the geometry is more complex for three-dimensional shapes, the positive/negative distinction is dropped for those shapes; α is always defined as positive for three-dimensional shapes.

When \mathbf{r} is parallel to $\hat{\mathbf{N}}$, it holds that $\alpha = 0$. This occurs when the polygon appears to ‘balance’ on one of its vertices, as illustrated on the left side of Fig. 4.43. The maximum α occurs when one of the polygon’s faces is parallel to the flat surface, as illustrated on the right side of Fig. 4.43. This can happen either with the shape rotated to the right (as shown in the figure) or to the left. It is clear from the figure that α_{max} decreases as the number of vertices n of a polygon increases. It holds that a triangle ($n = 3$) has $\alpha_{max} = 60$ deg, a square ($n = 4$) has $\alpha_{max} = 45$ deg, and a pentagon ($n = 5$) has $\alpha_{max} = 36$ deg. In the limit, it holds that a circle ($n = \infty$) has

$\alpha_{max} = 0$. For an arbitrary convex, regular polygon, the geometry illustrated in Fig. 4.44 may be used to find that $\alpha_{max} = \pi/n$.




Figure 4.44: Geometry of a regular, convex polygon.

The area of a regular, convex polygon can be found as $A = \frac{1}{4}ns^2 \tan \beta$. When given a unit area $A = 1$, the corresponding radius r can be found as:

$$r = \sqrt{\frac{2}{n \sin\left(\frac{2\pi}{n}\right)}} \quad (4.11)$$

When one of these polygons is dropped vertically onto a flat plane, the aspect angle at impact will vary *uniformly* between $-\alpha_{max}$ and α_{max} . Upon impact, the orientation of the collision vector \mathbf{r} determines how the contact point velocity \mathbf{v} decomposes along the normal and tangential dimensions. This is illustrated in Fig. 4.45, which shows the PDF of the two \mathbf{r} components for polygons with increasing n . These curves were produced by randomly generating $N = 10^6$ impact geometries and computing the corresponding $r \cos \alpha$ and $r \sin \alpha$. Four corresponding sample points of the triangle are visualized in Fig. 4.46.

Before discussing the trends of this figure, it is important to note that the friction force only dissipates energy whenever $s \neq 0$. Given that work is equal to the product of force and displacement, friction cannot perform work (and thus, it cannot dissipate energy) when $s = 0$, as in that case no displacement of the contact point occurs. In light of this fact, one may draw the following conclusions from Fig. 4.45:

- (1) The $r \cos(\alpha)$ and $r \sin(\alpha)$ curves show that polygons with fewer vertices n experience greater variation in the collision vector \mathbf{r} for random impact attitudes. As a result, they




Figure 4.45: Distribution of collision vector composition for random 2D polygon impacts.




Figure 4.46: Selected geometries as marked in Fig. 4.45.

also experience greater variation in the contact point velocity $\mathbf{v} = \mathbf{v} + \boldsymbol{\omega} \times \mathbf{r}$. Therefore, the following is true about shapes with low n : even if friction is sufficiently strong to drive $s \rightarrow 0$, the ‘random’ velocity variation that results from its shape is likely to yield some $s \neq 0$ at the next impact. Thus, friction is able to again drive $s \rightarrow 0$, dissipating energy in the process. This repeats at each successive impact, allowing friction to repeatedly dissipate energy in each of the impacts. The opposite is true for shapes with high n : they experience little variation in contact point velocity between successive impacts, so even if friction drives $s \rightarrow 0$ in some impact, it is unlikely that a large s will be present at the next impact. Friction is therefore unable to dissipate much energy. In conclusion, polygons with low n experience a greater energy dissipation rate than polygons with high n , due to the greater ‘natural’ variation in contact point velocity that low- n shapes experience in successive impacts.

- (2) As was mentioned at the start of our contact model development, it holds for non-central collisions (which have $\alpha \neq 0$) that the restitution and friction impulses are coupled. Mathematically, this is represented by a non-diagonal inverse mass matrix $[\mathbf{M}]$ in Eq. 2.54. As a result, the application of a normal impulse leads to a change in the tangential contact point velocity, and the application of a friction impulse leads to a change in the normal contact point velocity. The restitution and friction impulses may therefore either ‘help’ one another, or counteract one another. If friction is not sufficiently strong, the restitution impulse can therefore create some net $\dot{s} > 0$. This is illustrated in Fig. 4.47, where switching the direction of the angular velocity ω results in a switch in the direction of the friction impulse \mathbf{P}_F . On the left side of the figure, this results in (initial) alignment between the torques from the normal and friction impulses. On the right side of the figure, the torques are instead opposed. Effectively, this ‘tug of war’ between the normal and friction forces results in the same behavior mentioned above: a residual sliding velocity s continues to exist in successive impact, allowing friction to dissipate energy in most impacts, resulting in a higher energy dissipation rate for shapes with lower n .




Figure 4.47: Sample impacts of a tetrahedron with opposing friction force direction.

This analysis can be extended from the two-dimensional polygons of Fig. 4.43 to the three-dimensional shapes of Fig. 4.36. To substantiate that argument, $N = 10^6$ impact geometries of the eight shapes previously used in simulation are realized, and the distribution of the corresponding

aspect angles computed. These distributions are plotted in Fig. 4.48. Unlike the two-dimensional polygons, the distribution of aspect angles is *not* uniform for these three-dimensional shapes. Furthermore, any rotational velocity of the considered shapes will restrict the range of feasible α . Therefore, the provided curves only hold for initial impacts in which the angular velocity is negligible.

These distributions further indicate that the tetrahedron experiences the greatest variation in α of the five platonic solids, followed by the octahedron, cube, icosahedron, and finally the dodecahedron with the smallest variation in α for random impact geometries. By analogy of our discussion of friction dynamics for two-dimensional polygons, the energy dissipation rate of the tetrahedron will therefore be higher than that of the octahedron, etcetera. This agrees with the trends observed in our small-body deployment simulations: the tetrahedron settles the quickest, followed by the octahedron, etcetera. Dome sample impact geometries of the tetrahedron are illustrated in Fig.4.49.




Figure 4.48: Distribution of aspect angles for the eight probe shapes, for random impact geometries.




Figure 4.49: Example aspect angles of the tetrahedron.

The distribution of aspect angles shown in Fig. 4.48 further helps explain the deployment statistics observed for the three Hayabusa2 payloads. Given the similarity in aspect angle distribution between MASCOT and the cube, it makes sense for the two shapes to display similar dynamics; this is expected from their similarity in shape. MINERVA-II-2 experiences slightly less variation in α than the cube, its mean settling time is therefore slightly longer than the cube. Finally, although MINERVA-II-1 does experience high α values between 30 and 60 deg, Fig. 4.48 shows that it is most likely to experience far lower values, between 0 and 30 deg. This predicts that the rover will experience some trajectories in which it settles quickly, but far more in which it settles more slowly. This agrees with the simulation trends of Fig. 4.42.

Finally, it is noted that the range of aspect angles of the platonic solids is related to their vertex defect δ , that was introduced in Eq. 4.10. It holds that shapes with a greater defect have greater curvature, and thus a smaller *radius* of curvature¹. As a result, they are able to rotate around a contact vertex over a greater range of α angles before touching the surface with some other vertex. For the platonic solids, the vertex defect is therefore a good predictor of the settling behavior; the δ values given in Table 4.7 match the order of the shapes' mean settling times. The concept of vertex defect and curvature is more difficult to apply to the Hayabusa2 payloads, since

¹ The reader is referred to Section 3.2 for an in-depth discussion of curvature.

their shapes are not spherically symmetric. They therefore exhibit non-uniform curvature that changes depending on the considered ‘direction’ of curvature. For these probes, the aspect angle distributions of Fig. 4.48 are far more instructive when ‘predicting’ their deployment dynamics.

The behavior of MINERVA-II-1 As mentioned in the discussion above, the MINERVA-II-1 rover exhibits peculiar settling behavior: some simulations settling relatively quickly (at approximately the same t_f as the cube) but most have far longer settling times (with t_f exceeding even that of the dodecahedron). This behavior can be explained using the aspect angle statistics of Fig. 4.48. Resulting from the disk-like shape of MINERVA-II-1, the rover experiences two distinctly different modes of motion, which are illustrated in Fig. 4.50.




Figure 4.50: Illustration of the two modes of motion of the MINERVA-II-1 rover.

The rover can rotate about either its major axis of inertia ω_1 or about one of the other two axes ω_2 . In the first case, the rover effectively acts as a circle that easily rolls across the small-body surface. In this mode, the rover experiences relatively low aspect angles, because the \mathbf{r} vector is nearly parallel to the $\hat{\mathbf{N}}$ vector. In the second case, the rover acts as a rectangle and experiences far greater α angles. Following the above discussion that greater α variation results in faster settling, it

makes sense that the dichotomy of α values of MINERVA-II-1 yields a corresponding dichotomy in deployment statistics. The two modes of motion have been marked on the bottom right of Fig. 4.48. Two sample trajectories that display each mode are illustrated in Figs. 4.51 and 4.52. Of course, it is possible and likely for a given simulation to exhibit both modes of motion; it is the balance between the two that determines the simulation's settling behavior.




Figure 4.51: Sample MINERVA-II-1 simulation in which Mode 2 dominates.




Figure 4.52: Sample MINERVA-II-1 simulation in which Mode 1 dominates.

On a flat surface, a sphere always experiences the same normal and tangential directions; the energy dissipation rate from friction is therefore low. When bouncing on a rough surface, the surface normal is effectively randomized at each successive impact, allowing friction to dissipate far more energy. This results in shorter settling times and smaller surface dispersions. This is illustrated in Fig. 4.53. This phenomenon is also expected to occur when surface features are

included in simulations of the MINERVA-II-1 rover. Precisely how the inclusion of surface features affects the motion of the other considered probe shapes will be discussed in Section 4.3. Given that the tetrahedron naturally experiences a large ‘random’ variation in α , the inclusion of small surface features is predicted to have a relatively small additional effect. In contrast, given that the icosahedron and dodecahedron naturally experience very little α variation, the inclusion of small surface features is predicted to artificially increase the range of α values, resulting in shorter settling times. In this, the precise balance between the natural α variation (resulting from the probe shape) and the artificial α variation (resulting from small terrain features) is an interesting parameter to be investigated.




Figure 4.53: Bouncing of a sphere on (left) flat and (right) irregular terrain.

Duration of contact In the modeling of the contact interactions between a spacecraft and the small-body surface, a distinction between impulsive collisions and continued contact motion was made. The latter begins when a collision starts with a small normal contact point velocity $v_N < -\epsilon$, with ϵ some small value. When this happens, one final collision with only a compression phase is evaluated, such that $v_N = 0$ at the end of the collision. If the normal acceleration of the contact point, as computed with Eq. 2.54 with $\mathbf{F}_c = \mathbf{0}$, is negative, this implies that a continuous normal force is necessary to prevent the probe from penetrating the surface. Continued contact has thus begun, during which Eq. 2.106 is used to propagate the spacecraft motion. If the normal contact point acceleration is positive, the probe is simply returned to another ballistic arc.

The propagation of contact motion is numerically intensive, much more so than the propagation of ‘bouncing’ arcs. This is because a high number of low-velocity impacts occur before a probe finally comes to rest, following the onset of contact motion. Given that shapes with large variation in the aspect angle α experience large variations in contact point velocity, they must perform many collisions before v_N reaches below the ϵ limit and continued contact motion begins. By the time this happens, those shapes have already dissipated most of their initial energy and are close to coming to rest on the small-body surface. In contrast, shapes with smaller α variation are able to enter contact motion much sooner before coming to rest.

We are therefore curious to quantify the relative importance of continued contact motion for the tested shapes. In this, the icosahedron, dodecahedron, and MINERVA-II-1 are similarly expected to display significant contact motion, given their limited range of natural α variation. To investigate this, consider the results shown in Fig. 4.54. In it, solid lines indicate the t_f distribution of the complete trajectories that include continued contact motion. The dotted lines indicate the corresponding statistics when those same trajectories are terminated at the start of the first continued contact. Similarly, Fig. 4.55 plots the cumulative distribution function of the relative increase in settling time when contact motion is included. That figure also plots the additional distance that is covered between the start of the first contact motion and the final settling location.




Figure 4.54: Settling time PDFs of considered shapes, with and without continued contact.

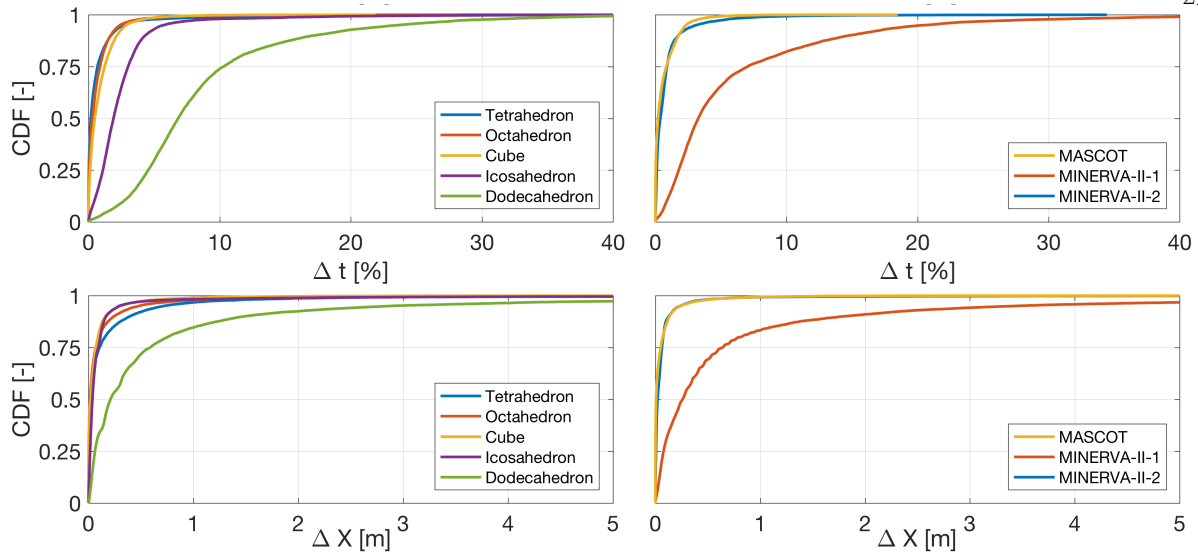


Figure 4.55: Changes in settling position and time when ignoring continued contact.

Reviewing Figs. 4.54 and 4.55, their results are found to support the predictions made based on the shapes' α variation: only the icosahedron, dodecahedron, and MINERVA-II-1 show any significant contact motion. All other shapes are virtually unaffected by the inclusion of contact motion, with 90% of simulations experiencing less than 0.25 m of additional movement and less than 2% increase in settling time after the first continued contact. For these shapes, it is thus appropriate to terminate simulations at the start of continued contact. The corresponding simulations will still provide the proper dispersion statistics. However, they may not provide the correct spacecraft settling *attitude*, since the spacecraft often still complete a full rotation after the start of contact motion. Nevertheless, this allows for a notable increase in the effective simulation speed.

A typical deployment of the cube was previously illustrated in Fig. 4.32. In it, the cube is shown at several time steps along its bouncing trajectory. This trajectory is a succession of multiple bouncing arcs, after which the cube eventually settles on one of its faces. In this trajectory, as is typical for the cube, the contact phases are extremely short and occur at the bottom right of the figure as the cube performs micro-impacts on its four bottom vertices before coming to a rest. To contrast this typical simulation, a rare example in which the inclusion of contact motion *is* important is also illustrated, in Fig. 4.56.




Figure 4.56: Atypical deployment of a cube, in which the inclusion of contact motion results in significant additional motion.

In this case, the motion is such that the cube skirts the side of a hill, allowing for its contact point to reach a zero normal velocity unusually early on in the simulation. The cube then continues to slide down the hill while spinning and eventually takes off again to perform more bounces before eventually settling. This kind of behavior, in which contact motion is important, is extremely rare for the cube. This does not necessarily hold for the other three shapes. The icosahedron shows more notable increase in t_f , though its additional movement is equally insignificant as that of the shapes mentioned above. The dodecahedron dynamics are significantly affected by the inclusion of contact motion, with 25% of simulations exhibiting an increase in t_f of more than 10%. Nevertheless, their additional surface motion remains small. Very similar results are found for MINERVA-II-1.

Suggested experiments Our simulations above revealed the significant effects of a probe's shape on its deployment dynamics, which could be explained from the properties of our contact model. In this, it is critically important to stress that these observations and trends only exist when the model's assumptions hold. More specifically, our analysis pivots on the assumption that the coefficient of restitution and friction are constant, *i.e.*, that they are independent of the impact attitude and velocity of an object. Despite the large body of research that exists on impact

dynamics and the application of Stronge's model, this assumption appears (to the best of our knowledge) to remain untested. Ahmad, Ismail, and Mat performed a literature review of various impact models and coefficient of restitution formulations [3] in which they provide three different definitions of the coefficient of restitution: one kinematic (Newton), one kinetic (Poisson), and one energetic (Stronge). They claim that the "*energetic COR that was introduced by Stronge is the most consistent and applicable in a wider application.*" Although a large body of research exists on the experimental measurement of the coefficient of restitution of spheres (see for example [32] and [25]), very little work on other shapes has been performed. Indeed, Ahmad, Ismail, and Mat note that "*The experimental works to determine COR in oblique and repeated impacts are still given less attention.*" Although Stronge's work is widely cited, the available literature appears to take the assumption of fixed e and f as a given. The basic research that validates this assumption remains an open area of research.

Given the importance of the fixed e/f assumption for the simulation of small-body surface exploration, two sets of experiments are therefore proposed. Both experiments avoid the complexity of three-dimensional impacts and restrict the geometry to two dimensions. In the first experiment, illustrated on the left side of Fig. 4.57, two plexiglass plates are placed close to one another. Various polygons (*e.g.*, triangle, square, pentagon) are then cut out of additional plexiglass material, and dropped between the two plates. The bouncing motion of the polygons is observed with a camera at high framerates, not unlike previous experiments carried out to measure the coefficient of rolling resistance [130]. The video material is then analyzed in order to estimate the position, attitude, and velocities of the polygons in each frame. This results in estimates of the linear and angular velocity change at each collision, which can be compared with Stronge's model to (in)validate the assumption of constant restitution and friction coefficients. It is likely necessary to develop some sort of clamping mechanism that is able to release the polygons with a fixed initial attitude and velocity, in order to make the experiment repeatable. Finally, it may be relevant to repeat this experiment with various surface materials on which the polygons bounce (shown in dark blue in Fig. 4.57), to reveal dependence on material properties. This experiment has a low cost.




Figure 4.57: Illustration of proposed (left) plexiglass and (right) air-bearing table experiments.

The first experiment is performed in Earth gravity, even though our small-body application takes place in zero-gravity. The first set of experiments is therefore proposed to be followed with a second set in which an air-bearing table is used. The same polygon shapes can then be ‘launched’ at low velocity across the table and impact a wall. The air flow cancels the gravitational acceleration, effectively creating an artificial micro-gravity environment. By observing the impact from above with a high-framerate camera, the velocity changes that occur over the course of an impact can again be estimated. By varying the initial attitude and velocities and comparing the results with simulations using Stronge’s model, one can validate if the constant e/f assumption persists in micro-gravity environments. The table can also be slightly tilted in order to simulate a milli-gravity environment. This second set of experiments carries a slightly higher cost than the first, but can still be done at relatively low cost, in particular if an air-bearing table is available for use.

4.2.3 Effect of surface interaction coefficients

Our contact model is governed by the coefficients of restitution e and friction f . The values of these coefficients determine a probe’s energy dissipation rate and thus also its settling statistics. Section 4.1 performed a parametric study investigating the effects of variations in these coefficients on the deployment of *spherical* lander probes. The restitution coefficient was found to be the major determinant of the settling behavior. No effects from the coefficient of friction were observed when it was varied within a realistic range; the statistics only changed when the coefficient reached

particularly low values of $f < 0.20$. This was also witnessed by Tardivel et al. [119].

Here, it is briefly investigate whether the effects of restitution and friction persist for probes with varying shapes. We repeat the nominal simulations of the tetrahedron, cube, and dodecahedron while independently varying the coefficients of restitution and friction. Fig. 4.58 scatter-plots the mean settling time \bar{t}_f as a function of the two coefficients. The plotted lines are function fits that are discussed in further detail below. Note that different vertical axes are used on the left and right sides of the figure.

Coefficient of restitution In agreement with the results for spherical probes, the restitution coefficient is found to have a significant effect on the settling time. As argued in Section 4.1, the settling time of a particle bouncing on a flat plane with coefficient of restitution e , in a uniform gravity field, can be written as:

$$t_f(e) = a_0 + a_1(1 - e)^{-1} \quad (4.12)$$

in which a_0 and a_1 are coefficients to be fit. We fit this expression to the $\bar{t}_f(e)$ curves obtained for the considered shapes, using a least-squares method with equal weights. These fitted curves are the solid lines included on the left side of Fig. 4.58. As can be seen from the figure, the fitted curve closely matches the scattered simulation data. This is surprising, given the complex dynamics that occur in our simulations. Despite the complex probe shape, asteroid shape, and asteroid gravity, the effect of a change in e is accurately captured by the simple scaling with $(1 - e)^{-1}$.




Figure 4.58: Mean deployment t_f for varying (left) restitution and (right) friction coefficient.

Coefficient of friction Varying the coefficient of friction, it is found that an increase in f increases the mean t_f for all shapes, though the effect appears much weaker for the dodecahedron than for the cube and tetrahedron. A saturation \bar{f} appears, beyond which further increases in the value of f no longer effect a change in the mean t_f . This saturation limit is easily explained from the properties of our friction model: when f is sufficiently high, the contact point reaches stick during *each* collision. If f is increased further, this behavior does not change. If f is decreased, the contact point no longer achieves stick during each collision, such that the settling time changes. As a result, the $\bar{t}_f(f)$ data visually appears as a sigmoid curve. We therefore attempt to fit a general sigmoid:

$$\bar{t}_f(f) = b_0 + \frac{b_1}{1 + \exp(-b_2(f - b_3))} \quad (4.13)$$

to the three sets of data, in which b_0 through b_3 are coefficients to be fit. The solid lines plotted on the right side of Fig. 4.58 are the resulting fitted sigmoid curves. Once again, a surprisingly close agreement between the scattered and data and the fitted curve is found.

Qualitatively, these results match the trends found for the sphere in Section 4.1: a saturation \bar{f} appears beyond which \bar{t}_f no longer changes. This limit was $\bar{f} \sim 0.2$ for a sphere. The saturation \bar{f} of the three tested shapes is much higher and shape-dependent: $\bar{f} \sim 0.6$ for the dodecahedron, $\bar{f} \sim 1.0$ for the cube, and $\bar{f} \sim 1.4$ for the tetrahedron. These differences can be explained from our analysis of the aspect angle variation for the different shapes. Given that a tetrahedron experiences the greatest variations in α , and therefore also in its contact point velocity, it requires a greater friction coefficient to ensure stick synchronization at each impact. Given that a more spherical shape experiences far smaller variations in α and in its contact point velocities, a lower friction coefficient is sufficient to achieve stick. Even if stick is not reached on the first impact, the tangential velocity will have changed very little by the second impact, such that it can easily be achieved on the second impact. The reason why an increase in friction coefficient somewhat counter-intuitively results in *longer* settling times will be detailed furthered in our analysis of the effects of the probe mass distribution below.

4.2.4 Effect of probe mass distribution

Above, the effect that a spacecraft's shape has on its deployment dynamics was demonstrated. More specifically, it was shown that changes in the shape have two effects: first, they modify the *interface* along which spacecraft-target interactions occur and, second, they modify the *inertia* of the spacecraft. The former affects the collision dynamics through its control over the range of aspect angles α (and thus, over the range of contact point velocities) that the spacecraft encounters across all feasible impact geometries. The effect of the latter has not yet been identified, but is the focus of the current section.

In order to investigate the effect of spacecraft inertia on deployment dynamics, the nominal simulations are repeated for three *variants* of each of the platonic solids. These variants have the same physical shape and mass but different internal mass distributions; they have different inertias. Fig. 4.59 illustrates different mass distributions of a triangle, where the 'shell' variant (with all mass concentrated on the surface), the 'solid' variant (with uniform density), and the 'core' variant are marked. This figure is highly similar to Fig. 4.1 that illustrated different mass distributions of a spherical probe. The reduced inertia j of the 'shell' variant is defined as:

$$j_{core} = 2j_{solid} - j_{shell} \quad (4.14)$$




Figure 4.59: Illustration of different mass configurations of a triangle.

The j_{solid} and j_{shell} values were previously listed in Table 4.7. We repeat the nominal simulations for the three variants of each platonic solid and plot the corresponding settling time distributions in Fig. 4.60. Reviewing these results, it is found for all shapes that lower inertias

(the ‘core’ variants) experience lower settling times than higher inertias (the ‘shell’ variants). This trend is clearly seen in the tetrahedron, octahedron, and cube. Strangely, it is found that the differences between the solid and shell variants of the icosahedron and dodecahedron are almost negligible; both curves nearly overlap. For the dodecahedron, this is true of its core variant as well. These trends are curious: how can it be that the internal mass distribution affects the motion of some given shapes, but not some other shapes? A clue to understanding this behavior lies in the behavior of the icosahedron and dodecahedron; it seems as if some *saturation* inertia exists, beyond which further increases in the inertia no longer affect the motion. This suggests that friction may be responsible for this behavior, given that it was previously observed saturation when varying the coefficient of friction in Fig. 4.58.

To further investigate how friction relates to the effect of inertia on probe deployment dynamics, the coefficient of friction f was varied for the three inertia variants of the tetrahedron, cube, and dodecahedron. The resulting mean settling time values are collected in Fig. 4.61. Note that different vertical axes are used in the three subplots. As with the f variation of the ‘solid’ variant previously shown in Fig. 4.58, a saturation limit \bar{f} appears, beyond which further increases in f cease to have an effect. In agreement with the results of Fig. 4.58, it is found that value of \bar{f} changes with the considered probe shape. However, it is found that \bar{f} also changes with the *mass distribution* of a considered probe. More specifically, \bar{f} decreases when the mass is moved outwards (increased inertia) within a given shape. For example, the tetrahedron shell variant has $\bar{f} \sim 1.1$, while its solid variant has $\bar{f} \sim 1.4$ and its core variant has some $\bar{f} > 1.5$. The effects of changing the spacecraft shape and changing its internal mass distribution appear to counteract one another:

- (1) Shapes with fewer vertices (that naturally have higher inertia) settle faster than shapes with more vertices (that naturally have lower inertia).
- (2) Shape variants with a more outwards mass distribution (that have higher inertia) settle slower than shape variants with a more concentrated mass distribution (that have lower inertia).




Figure 4.60: Settling time distribution with varied spacecraft inertias.




Figure 4.61: Mean settling time for varying f and j of the tetrahedron, cube, and dodecahedron.

To understand these seemingly contradictory observations, we continue to investigate the collision dynamics that occur during deployment. Since the results from Fig. 4.61 suggest that friction is responsible for the observed trends, we investigate the sliding velocity changes that occur in the simulations of the cube. To this end, Fig. 4.62 shows the CDF of three relevant velocities of the cube variants. First, the tangential contact point velocity after the first impact, s_1 . Second, the initial tangential contact point velocity \bar{s}_0 of all successive impacts (*i.e.*, all impacts excluding the first). Finally, the tangential contact point velocity \bar{s}_1 at the end of each successive impact. In this, *micro-impacts* that occur within five seconds of each other are grouped together into a single *macro-impact*. These distributions are shown for simulations performed with different friction coefficient $0.40 < f < 1.20$.




Figure 4.62: Tangential contact point velocity statistics for cube variants.

As might be expected, an increase in the friction coefficient f reduces the post-impact s , both after the first and after all successive impacts. This makes intuitive sense: when f is increased, there is a stronger friction force opposing the sliding velocity of the contact point. However, an

additional trend can be seen that is not expected: for a given f , the core variant (with lower inertia) shows *greater* pre- and post-impact s than the shell variant (with greater inertia). Given that lower inertia implies that greater changes in the angular velocity can be made with the same friction force/torque, one would expect the core variant to achieve stick ($s = 0$) *more* frequently than the shell variant. To understand why the opposite is observed (*i.e.*, why lower inertia achieves stick *less* frequently), we return to Eqs. 2.59 and 2.72 that express the change in contact point velocity \mathbf{v} and sliding velocity s :

$$\begin{cases} \frac{d\mathbf{v}}{dp} = \begin{bmatrix} \frac{dv_1}{dp} & \frac{dv_2}{dp} & \frac{dv_3}{dp} \end{bmatrix}^T = [\mathbf{M}] \frac{d\mathbf{P}}{dp} \\ \frac{ds}{dp} = \cos \phi \frac{dv_1}{dp} + \sin \phi \frac{dv_2}{dp} \end{cases} \quad (4.15)$$

These equations can be used to evaluate the change in sliding velocity s in response to some impulse $d\mathbf{P}$. That impulse combines both a normal component (dP_3) and a frictional component (dP_1 and dP_2). Since the mass matrix $[\mathbf{M}]$ is not diagonal for an arbitrary spacecraft shape and collision geometry, it holds that the normal (restitution) and frictional impulses each change both the normal *and* tangential contact point velocity. The resulting coupling between restitution and friction is the reason for the unexpected behavior seen in Fig. 4.62: in shapes with lower inertia, the two impulses perturb one another more strongly, such that greater f is necessary to achieve stick. To illustrate this point, consider the results shown in Fig. 4.63. This figure plots the change in sliding velocity magnitude ds at the start of some sample impact of the cube with a flat plane. We vary the coefficient of friction f and direction of initial slip ϕ , and compute the resulting ds of the three cube variants.

Reviewing this figure, it is seen that the core variant (lowest inertia) experiences the greatest range of initial $ds = [-4, 1] \text{ m/s}^2$. In contrast, the shell variant (with highest inertia) only sees changes between $ds = [-1, 0.5] \text{ m/s}^2$. This difference in ds magnitude is expected: lower inertia implies that the contact point velocity changes more rapidly in response to a given force. More importantly, it is seen that the core variant has a relatively wide range of (ϕ, f) that have $ds > 0$. Although the solid and shell variants also exhibit such a region, it is much smaller.




Figure 4.63: Tangential contact point rate for sample cube collision.

The fact that ds can be positive indicates that, for certain combinations of f and ϕ , friction is not sufficiently strong to prevent the restitution from increasing the sliding velocity s . When the inertia is reduced, this balance appears to shift further in favor of the restitution impulse, with a wider range of f and ϕ resulting in $ds > 0$. Shape variants with lower inertia therefore experience greater residual s , allowing friction to dissipate energy during more impacts. This leads to low-inertia variants experiences greater energy dissipation rates than high-inertia variants and, correspondingly, to low-inertia variants having shorter settling times. This is only true when the friction coefficient is smaller than the respective saturation value, *i.e.*, when $f < \bar{f}$. If $f \geq \bar{f}$ for all considered variants of a shape, the resulting behavior will be approximately equal. Although a change in shape from a ‘concentrated’ configuration (the dodecahedron) to a more ‘outwards’ configuration (the tetrahedron) inherently increases the inertia, the net effect on the shape’s deployment mirrors that of a *decrease* in inertia. This indicates that a probe’s shape has a more determinant effect on its dynamics than the probe’s inertia does. This implies that, when one designs a small-body exploration probe to settle as quickly as possible, the selection of the probe shape is more important than its internal configuration. This creates some freedom when placing scientific instruments and designing thermal control systems; their placement within the probe is not as important as the shape of that shell. Additionally, these results suggest that the placement of spikes on the probe corners will further restrict the settling time and surface dispersion.

Discussion: Deployment of complex probes Using a sample scenario of 100,000 low-altitude deployments to asteroid Itokawa, our parallelized simulation framework was used to investigate how the shape of a probe affects its deployment dynamics. In this, the five platonic solids were considered, as well as the three deployed payloads included on Hayabusa2. The tetrahedron was found to have the shortest settling time and smallest surface dispersion, followed by the octahedron, cube, icosahedron, and dodecahedron. The MASCOT and MINERVA-II-2 rovers showed similar behavior to the cube, but the MINERVA-II-1 rover had a considerably longer settling time. These trends result from the variation in the aspect angle that relates to the net contact point velocity in spacecraft-target impacts. This not only determines the frequency with which friction is able to dissipate energy, but also the degree to which the normal and friction forces are coupled. The net effect of these properties is that shapes more distorted from a sphere (which experience greater aspect angle variation) settle faster and disperse over a small surface area compared to shapes that are more spherical (which experience smaller aspect angle variation). We also showed that the inclusion of continued contact motion is important only for shapes with little variation in that aspect angle, *i.e.*, the icosahedron, dodecahedron, and MINERVA-II-1.

Next, the coefficients of restitution e and friction f that determine the magnitude of the spacecraft-target surface interactions were varied. In agreement with previous results obtained for spherical probes, the coefficient of restitution e was found to govern the deployment dynamics. Higher e values result in longer settling times and greater surface dispersions, with the mean settling time scaling as $\bar{t}_f \propto (1 - e)^{-1}$. Despite the complex behavior encountered in a single trajectory, the emergent behavior in a batch of simulations fits this simple relation surprisingly well. The same is true for the effect of the coefficient of friction, which scales the mean settling time as a sigmoid curve with $\bar{t}_f \propto (1 + \exp(f))^{-1}$. In this, higher friction coefficients result in longer settling times. A saturation \bar{f} occurs, beyond which further increases in f cease to affect the deployment dynamics. This \bar{f} changes for different shapes and is higher for less spherical shapes and shape variants with a greater inertia.

Finally, the internal mass distribution of the considered shapes was varied. When $f < \bar{f}$, it was found that a given shape with a more outwards mass distribution (*i.e.*, higher inertia) settles faster than that same shape with a more inwards mass distribution (*i.e.*, lower inertia). We showed that this occurs because lower-inertia variants of a shape require greater f in order to overcome the sliding velocity changes imparted by the normal force. Friction is therefore able to dissipate energy during more collisions for lower-inertia variants than for higher-inertia variants of that shape.

The effect of moving mass outwards within a given shape and of ‘molding’ a shape outwards are opposite. Although changes in shape imply changes in inertia, it was found that the outwards ‘molding’ has a greater effect. This implies that, if a probe is to be designed with minimal settling time and surface dispersion, its shape should be moved outwards while its mass is moved inwards. In the limit, the minimal settling time will be achieved by a point mass with long spikes. Either way, it allows engineers more liberty with the internal placement of components within a probe, as long as the probe shell within which these components are contained is properly designed. In the following section, it will be investigated how the inclusion of surface details such as rocks and boulders, that are not included in global small-body shape models, further affect these trends.

4.3 Surface topography

In most of the parametric studies discussed above, the considered small-body surface was relatively smooth. Although the applied shape models (in either polyhedron or SDF form) account for surface features such as boulders and hills down to about 5 m in size, as per our discussion in Section 2.3, they cannot resolve smaller features. Section 4.1 investigated the effect of polyhedral rock presence on the motion of spherical probes and found it to significantly influence the probes’ motion. In the current section, this analysis is extended to the complex probes of Section 4.2. This is done using the various feature generation methods discussed in Section 2.3. These allow for generation of statistical features at much lower computational cost than the polyhedral rocks mentioned before. Finally, the effects of regolith presence on the small body are also investigated, using the restitution modulation scheme that was presented in Section 2.6.

4.3.1 Effect of surface noise

The inclusion of statistical surface features, which effectively ‘perturbs’ the otherwise smooth small-body surface, has two implications for the motion of a surface exploration probe. Both of these relate to the geometry of probe-target collisions and were hinted at in our discussion of Fig. 2.25, which illustrated how spacecraft of different resolution (with respect to the distribution of vertices across their shell) resolve impacts with small surface features differently. This is now reviewed in detail. Consider the impact geometry of a tessellated rectangle impacting a surface protrusion, as illustrated in Fig. 4.64.




Figure 4.64: Definition of the α and β angles.

This illustration defines two angles present in a collision, that will be relevant to our discussion of the effects of surface topography. First is the *aspect angle* α , which was previously used in Section 4.2 to explain the effects of a probe’s shape on its bouncing behavior. In that case, the surface was a simple flat plane. Here, the concept of the aspect angle to arbitrary surfaces is extended and defined as the angle between the collision vector \mathbf{r} and the collision normal:

$$\alpha = \arccos \left(\frac{|\hat{\mathbf{N}} \cdot \mathbf{r}|}{\|\mathbf{r}\|} \right) \quad (4.16)$$

Second, the *surface inclination* angle β is defined as the angle between the collision normal $\hat{\mathbf{N}}$ and the local net acceleration vector \mathbf{a} , which combines gravitational and centrifugal surface accelera-

tions, as:

$$\beta = \arccos \left(\frac{\hat{\mathbf{N}} \cdot \mathbf{a}}{\|\mathbf{a}\|} \right) \quad (4.17)$$

The variation of the α and β angles across all feasible impact geometries of a probe is not arbitrary; both angles are subject to certain constraints that predict some of the effects of the presence of statistical features. For this, consider Fig. 4.65, which illustrates the variation in impact geometries for a rectangular probe impacting both a large and a small surface feature. In both cases, the impact point is fixed, but the probe is rotated to impact at different attitudes.




Figure 4.65: Variation of α and β in impacts on (left) a large and (right) a small surface feature.

Reviewing this figure, the following trends are observed: for varying impact attitudes on a large surface feature (the left side of the figure), the collision normal $\hat{\mathbf{N}}$ remains fixed. As a result, the incidence angle β is invariant with respect to the probe impact attitude. The impact angle α *does* vary when impacting a large feature. For impacts on a small surface feature, the situation is opposite: the impact angle α remains fixed while the incidence angle β varies.

In Section 4.2, which examined the effects of probe shape on deployment dynamics, we previously investigated how the statistics of the encountered impact angles α determine the rate at which energy is dissipated in collisions. We showed that shapes with a greater range of α

generally experience greater rates of energy dissipation, and therefore settle faster and with a smaller dispersion area. A similar theoretical argument can be made for the effect of the surface inclination β : when the range of β is high, the surface normal with respect to which collisions occur is greater. This implies that the normal and tangential directions may change significantly between successive impacts, allowing the normal and friction forces to dissipate a greater amount of energy than they would if the surface normal remains more constant.

Fig. 4.48 of Section 4.2 investigated the distribution of the aspect angles for the five platonic solids, dropped on a flat surface. This investigation is now extended by dropping the same probes above a noisy surface generated by including fBm noise onto a flat surface, as previously demonstrated in Fig. 2.22. This is done with an accumulation of noise octaves. The five spacecraft shapes are the platonic solids used in Section 4.2, again with a mass of $m = 5$ kg and uniform density of $\rho = 500$ g/cm³. The five shapes are dropped $N = 10^5$ times above the noisy surface; this is done for both the untessellated shapes and for finely tessellated variants of the same shape. In each collision, the aspect angle α and the surface incidence angle β are registered. The distributions of these angles are shown in Fig. 4.66, in which the untessellated shapes are shown with a solid line and the tessellated shapes are shown with a dashed line. We now review the observed trends in α and β of this figure.

Incidence angle. Immediately, it is seen (on the right side of Fig. 4.66) that the distributions are identical for all five shapes. This result makes sense, given that the five shapes have access to the same surface, and given that the probe attitude at impact does not affect the surface incidence angle β . When no noise is included onto the surface, the β distribution is a dirac function at $\beta = 0$ deg, though the kernel density estimator applied to generate the angle distributions smooths this into a steep, continuous curve. As increasingly smaller noise levels are included onto the surface, the β distribution begins to spread out towards higher values. The curves appear to converge on a fairly uniform distribution between $\beta = 0$ deg and $\beta = 60$ deg, with a maximum around $\beta = 25$ deg. In other words, the addition of smaller and smaller surface features has a smaller and smaller effect on the surface incidence angles experienced by the probes.




Figure 4.66: Variation in the α and β angles for impact with different noise levels.

This effect is understood from geometry. As previously mentioned, impacts on very large surface features must always occur at the corner vertices of a shape. Impacts on small surface features are much more likely to occur at some vertex in the middle of a facet of the considered probe, as illustrated in Fig. 4.64. Given that the surface normal in is equal to the probe facet normal those cases, the β distribution for surfaces with features smaller than the probe shape is effectively a distribution of the attitude of the probe itself, and is not directly related to the surface. This also explains why the β distributions of the tessellated and untessellated shapes are virtually indistinguishable. Only when the smallest surface features are included is there a very slight deviation, that is due to the mismatched normal selection when only detecting collisions at the corner vertices, as previously illustrated in Fig. 2.25.

Impact angle We now change our attention to the aspect angle distributions on the left side of Fig. 4.66. Here, significant differences are found when comparing the α distributions of the five shapes. This is expected given that even on a flat surface, different shapes experience different impact angles, as discussed in Section 4.2. Despite these differences, some common trends are observed. First of all, it can be seen that the two largest noise levels (with $A = 8$ m and $1/f = 16$ m, and with $A = 2$ m and $1/f = 4$ m) have a negligible effect on the α distributions. This makes sense: in the presence of features larger than the considered probe size, the probes continue to only experience impacts on their corner vertices, as was the case with impacts on a flat surface. This also explains why both the tessellated and untessellated variants experience the same α in the presence of large features: the additional vertices which the tessellated variants have spread across their vertices never touch the surface. As smaller noise levels are included, that generate statistical features smaller than the considered probe sizes, changes in the α distribution finally begin to appear. Unlike the β distributions, these changes are different for the tessellated and untessellated shape variants; they are less dramatic in the tessellated variants, the remaining discussion is focused on them. For the tetrahedron and octahedron, the changes in α are small; the peak of their distributions shifts ever so slightly towards lower α . For the cube, the differences are negligible. Only the icosahedron and dodecahedron display notable changes in their α distribution;

the curves become more spread out for both shapes. This implies that they experience greater variation in the impact angle α . Following the observation made in Section 4.2 that shapes with greater α variation settle faster and with smaller dispersion, the icosahedron and dodecahedron settling time and surface dispersion are therefore expected to reduce in the presence of small surface features.

These results are aggregated into Fig. 4.67. The left side of this figure shows the change in α distributions between a noise-less and a noisy surface, for untessellated shapes. The right side shows the same curves for the finely tessellated shape variants. This again underscores the importance of sufficiently tessellating spacecraft shapes when including small surface features into simulations. The right side of this figure indeed shows that the dodecahedron and icosahedron have significantly affected α distribution; the central peaks at respectively 30 deg and 20 deg have shrunk and the distributions have become more spread out. The α curves of the tetrahedron, octahedron, and cube have change very little; only a slight shift towards lower α values can be seen. At best, one could therefore expect their settling time and surface dispersion to increase slightly when including statistical surface features.

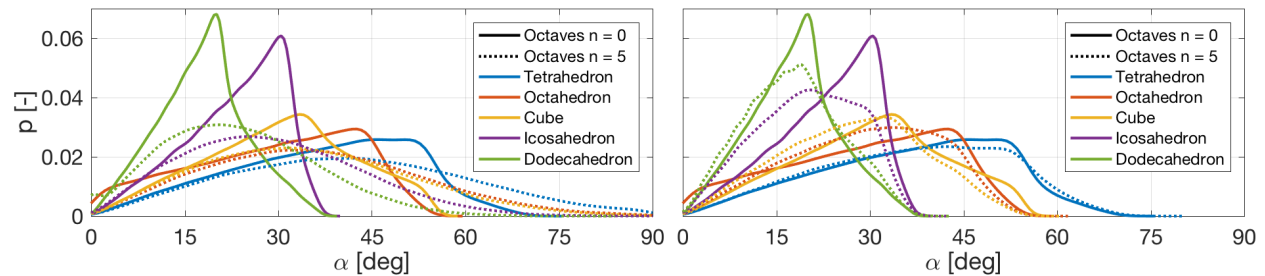


Figure 4.67: Impact angle distributions for flat and noisy surfaces, (left) for untessellated and (right) finely tessellated spacecraft shapes.

Simulations Having concluded the above analysis on the aspect and incidence angle distribution, it is now examined if the predictions made therein agree with simulations to a noisy surface. For this the nominal low-altitude deployment scenario to Sagami-hara region is repeated, as used before in Section 4.2. This is done for the finely tessellated variants of the five platonic

solids, with an increasing number of noise octaves. The same noise levels are used as applied above in the analysis of the α and β distributions. The settling time PDFs of these simulation scenarios are shown in Fig. 4.68.

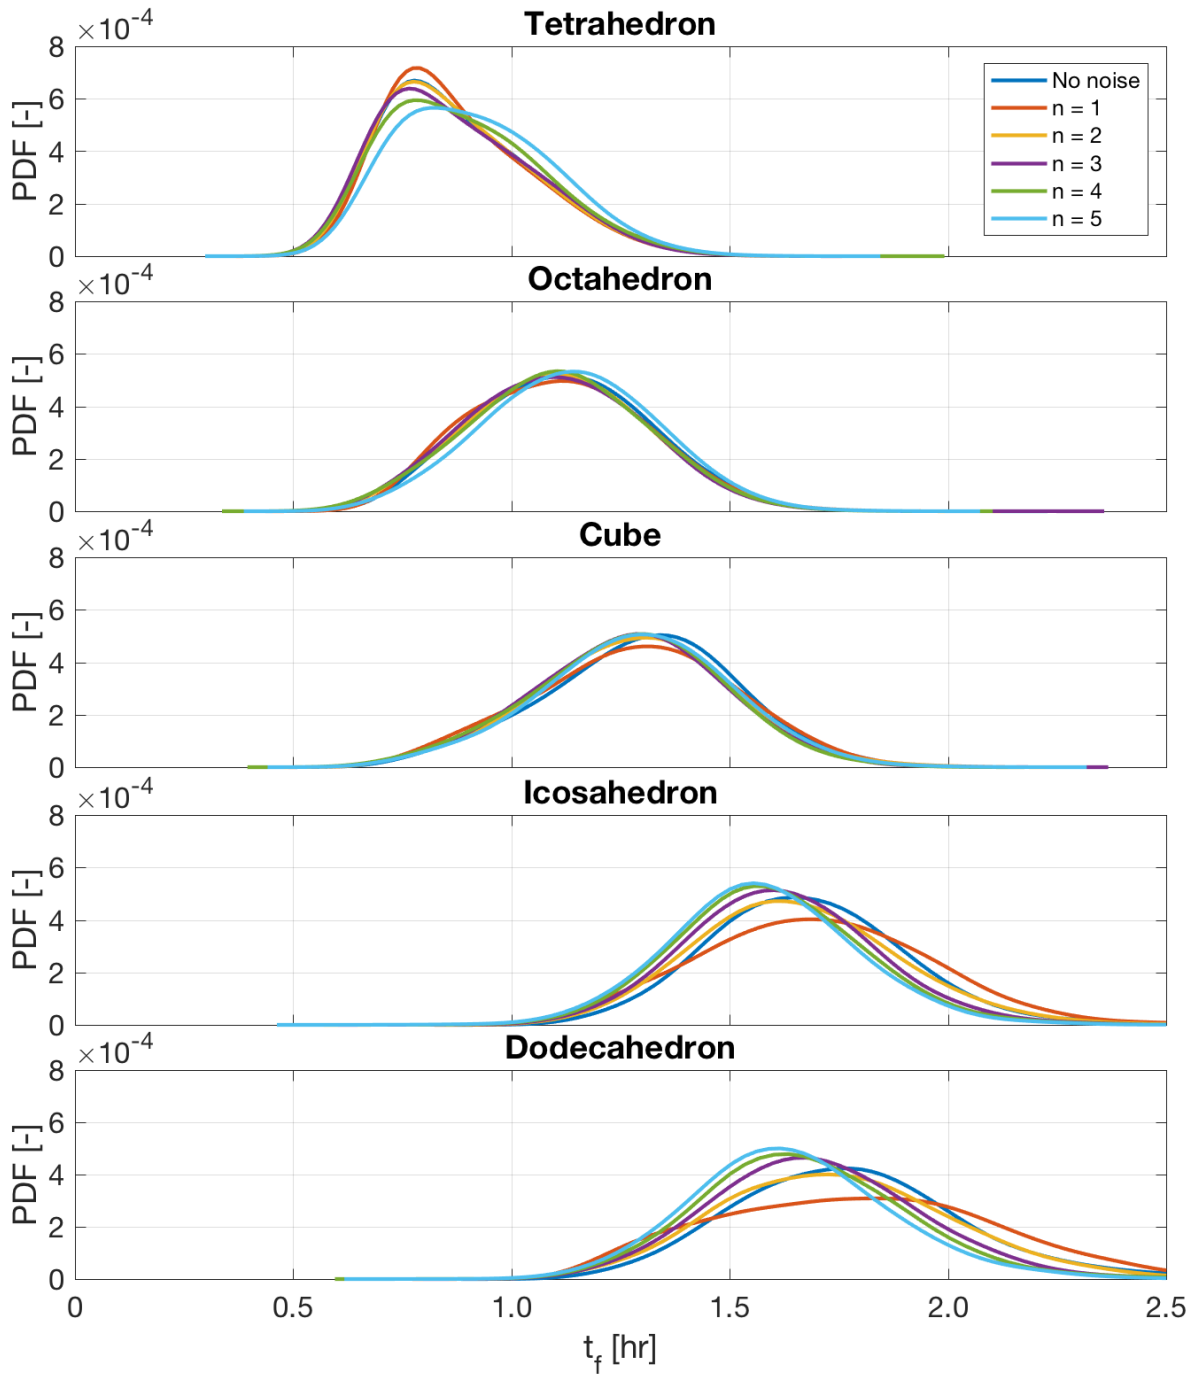


Figure 4.68: Settling time PDFs of deployments to various Itokawa surfaces.

We examine the trends in this Fig. 4.68 for the five considered shapes:

- (1) *Tetrahedron*. When including the largest three noise levels, the tetrahedron t_f distribution is affected very little. The central peak is preserved, with the mean t_f barely shifting. Only when the smallest features are included do the curves display notable change by flattening out and slightly moving towards higher t_f . This indicates that the presence of small surface features increases the settling time of the tetrahedron.
- (2) *Octahedron*. The observed changes in the octahedron t_f are barely noticeable at this scale, and probably not statistically significant. Although the curves shift ever so slightly towards higher t_f when the smallest noise features are included, the differences appear negligible.
- (3) *Cube*. The behavior of the cube is highly similar to that of the octahedron, though opposite in direction. When small surface features are included, the cube t_f is very slightly decreased.
- (4) *Icosahedron*. Compared to the three shapes discussed above, the icosahedron is affected far more strongly by the presence of surface features. When only the largest surface features are included, the t_f curve is simply seen to flatten out. However, when smaller features are accounted for, the t_f curves shift notably towards lower t_f values. This indicates that the presence of statistical surface features reduces the settling time of the icosahedron.
- (5) *Dodecahedron*. The behavior of the dodecahedron matches that of the icosahedron; its settling time is reduced notably when surface features are included in simulation.

Given that the observed t_f changes in the tetrahedron, octahedron, and cube are very small and difficult to see at the scale of Fig. 4.68, plots of the mean t_f of the considered simulations are provided in Fig. 4.69.

Overall, the results of Fig. 4.69 agree with those listed above for the different shapes. We see that the mean t_f of the cube remains virtually unaffected. The tetrahedron and octahedron t_f are increased slightly, while that of the icosahedron and dodecahedron are reduced more significantly.




Figure 4.69: Mean settling time trends of deployments to various Itokawa surfaces.

We provide one final visualization of these results, by plotting the t_f distribution of the the five shapes on respectively the smooth and $n = 5$ noisy Itokawa surface, in Fig. 4.70.




Figure 4.70: Settling time PDFs of deployments to smooth and noisy Itokawa.

The results shown in this figure agree with those mentioned above. In conclusion, they suggest that the inclusion of statistical surface features removes some of the bias in settling statistics that results from the shape of a deployed probe. On a smooth surface, the settling time distributions were fairly spread out. When surface features are included, the curves shift closer to that of the cube. In other words, the tetrahedron and octahedron (which are more ‘deformed’ than the cube, relative to a sphere) experience increased settling times. The icosahedron and dodecahedron (which are more spherical than the cube) experience decreased settling times. The cube itself is barely affected by the presence of surface features. This underscores the importance of accounting for statistical surface features when performing simulations of probe deployment to the small-body surface. This is of particular importance when the considered probe shape is nearly spherical in at least some of its principal axes. For example, the presence of rocks should be accounted for when simulating an icosahedron, dodecahedron, sphere, or the MINERVA-II-1 rover. It is of lesser

importance when simulating the tetrahedron or octahedron, and can be ignored when simulating a cube or shapes that behave like a cube, such as the MASCOT or MINERVA-II-2 spacecraft.

4.3.2 Effect of procedural rocks

The quantification of the effects of statistical surface features given above is relevant, since such features are typically not included in global small-body shape models. However, as discussed in Section 2.3.3, fractional Brownian motion is somewhat limited in the type of features that are produced. More specifically, it is difficult to tune the surface features to match observations of a small-body surface. Furthermore, fBm distorts the entire small-body surface, whereas the presence of rocks typically leaves parts of the base surface unaffected. Thankfully, global populations of rocks can be generated efficiently using the procedurally seeded rocks (PSRs) algorithm discussed in Section 2.3.3. Here, the effects of such PSRs on the motion of small-body surface exploration craft are investigated.

For this, the nominal deployment scenario used in the probe shape variation of Section 4.2 is repeated, with various populations of PSRs. More specifically, the Itokawa rock population observed by the Hayabusa spacecraft is used, as discussed in [68]. This population has a power index of $\alpha = 2.1$ and can be nominally initialized using the observation that Itokawa has $k_0 = 2.05 \times 10^{-3}$ rocks larger than $d_0 = 5$ m per unit surface area. We generate rocks between 2 m and 5 cm in diameter, using Eqs. 2.22 and 2.23. In order to examine the effects of rock populations with a different spatial density of rocks, the k_0 parameter is varied. Effectively, this will generate rocks with a fixed size distribution, but with varying numbers of rocks per unit surface area. Fig. 4.71 shows some sample rock populations generated with varying k_0 .

The same 10,000 simulations are now repeated for a range of k_0 values, and the resulting deployment statistics collected, for the eight considered shapes. The mean settling time values can be found in Fig. 4.72; the $1/2/3\sigma$ dispersion areas are quantified in Fig. 4.73.




Figure 4.71: Sample rock populations with varying k_0 .




Figure 4.72: Mean settling time of various shapes in the presence of PSRs.

Reviewing the settling time trends, it is obvious that the presence of rocks on the small-body surface has an effect on the probe dynamics, though the particular effect differs for different shapes. Consider the tetrahedron: its mean settling time is virtually unaffected by the presence of rocks. If anything, it is very slightly increased. A more noticeable effect is observed for the octahedron and dodecahedron, which experience a decrease in mean settling time of approximately 10% between a

smooth surface and a surface covered entirely with rocks. The Mascot and MINERVA-II-2 rovers display the same behavior, in agreement with the results from Section 4.2.




Figure 4.73: Surface dispersion of various shapes in the presence of PSRs.

The PSR inclusion appears to have a far more dramatic effect on the motion of the icosahedron, dodecahedron, and MINERVA-II-1 rover. The icosahedron and dodecahedron display very similar trends, with a reduction in mean settling time of approximately 40% between the smooth and rock-covered surfaces. A clear saturation limit appears: beyond $k_0 \sim 3 \times 10^{-3}$ rocks/ m^2 , a further increase in k_0 ceases to significantly affect the mean settling time. We also find that most of the change in t_f occurs at low k_0 . In other words, even a light-to-moderately sparse rock population is able to restrict the motion of the icosahedron and dodecahedron. The most dramatic trends are observed for the MINERVA-II-1 rover; even the inclusion of a very sparse population of rocks is sufficient to reduce the settling time by over 25%. A dense rock population even reduces the mean settling time of this rover to below that of the icosahedron and dodecahedron, despite the rover's

natural ability to roll around its axis of major inertia. This is notably different from those shapes' behavior on a smooth surface. Again, it is found that the inclusion of additional rocks beyond $k_0 \sim 3 \times 10^{-3}$ ceases to affect the settling time statistics. The results for Minerva-II-1 agree with our prediction made in Section 4.2 and illustrate that the inclusion of surface rocks in small-body modeling can be critically important. If not, the settling time of shapes such as the icosahedron, dodecahedron, and MINERVA-II-1, that are able to 'roll' on a smooth surface, may be greatly overestimated. This inclusion of rocks is far less important for the Mascot and MINERVA-II-2 rovers, since they already naturally resist rolling motion.

Corresponding results are found in the surface dispersion plots of Fig. 4.73. Again, it is found that the tetrahedron is barely affected by the presence of surface rocks. The surface dispersion of the octahedron and cube do appear to be more affected by the rocks than their settling time is. Again, the icosahedron, dodecahedron, and MINERVA-II-1 appear to be strongly affected by the presence of rocks, which limits the area over which these probes are able to disperse.

As a final investigation, the deployments of the tetrahedron, cube, and icosahedron are repeated for various spatial rock densities, while also varying the incidence angle γ . This angle will be defined in the subsequent section, and expresses whether the first impact of a deployment is normal or grazing. An incidence angle of $\gamma = 0$ deg indicates a fully normal impact, while an angle of $\gamma = 90$ deg corresponds to a fully grazing impact. Reviewing the results of Fig. 4.74, it is found that the effects of rocks hold for both normal and more grazing impacts. Although the settling time is dependent on the incidence angle γ , as will be detailed in the subsequent section, the effect of varying k_0 appears to be analogous for all γ values. In conclusion, these simulations indicate that it can be important to include the presence of rocks into surface modeling of small bodies. This is particularly true for shapes that are naturally able to roll across a smooth surface, such as the icosahedron, dodecahedron, and MINERVA-II-1.




Figure 4.74: Mean settling time for various impact angles and rock densities.

4.3.3 Effect of surface regolith

Our applied contact model, as detailed in Section 2.5, assumes that the coefficients of restitution and friction are constant and independent of the spacecraft velocities and attitude at impact. Although this is an assumption frequently made in literature [3], it is one of questionable accuracy. The verification or rejection of this assumption would require extensive experimental investigation, such as the experiments proposed for hard surface layers in Section 4.2. The assumption appears even more tenuous for soft regolith layers. Indeed, as discussed in Section 2.6, Nishida et al. observed that the coefficient of restitution e experienced spherical particles impacting a granular surface is strongly dependent on their impact angle θ . More specifically, they observed that normal impacts occur at notably lower e than tangential impacts [80]. Experiments of complex-shape impact into granular surfaces in micro-gravity are under preparation [77], but have not been published yet. Nevertheless, it is reasonable to expect similar behavior for non-spherical spacecraft. For that purpose, the observations by Nishida et al. are extrapolated and modulate the coefficient of restitution based on the effective impact angle, as discussed in Section 2.6.

Regolith model In order to apply the restitution-modulation model as presented in Eq. 2.112, knowledge of the density and diameter of both the target small body, as well as the impacting spacecraft, are required. We continue to use Itokawa as a sample body, which has a

bulk density of approximately $\rho_t = 1.9 \text{ g/cm}^3$. Miyamoto et al. note about the asteroid that “*Grain sizes observed in close-up images range from centimeters to several tens of meters. The finest particles are centimeter-sized pebbles, whose concentrations are found in the Muses-C smooth terrain.*” [73]. Given this observation, two regolith models are tested: one with relatively small regolith grains of $D_t = 1.0 \text{ cm}$ and one with relatively large regolith grains of $D_t = 2.5 \text{ cm}$. We will investigate the motion of three spacecraft shapes on this regolith surface: the tetrahedron, cube, and dodecahedron, as previously applied in Section 4.2. These probes have the same density $\rho_p = 500 \text{ g/cm}^3$, but slightly different radii as listed in Table 4.7 with $D_p = 2r$. Substituting these values into Eq. 2.112 for the critical θ angle, the values listed in Table 4.8 are obtained. To provide a broader intuition of how the critical impact angle of these three spacecraft shapes would vary on different regolith surfaces, consider the contour plots of Fig. 4.75.

Table 4.8: Critical impact angles θ_E for the tested spacecraft shapes and regolith sizes.

θ_E [deg]	Tetrahedron	Cube	Dodecahedron
$D_t = 1.0 \text{ cm}$	31.6	35.5	41.8
$D_t = 2.5 \text{ cm}$	58.2	65.4	77.0

These plot the value of θ_E as a function of the regolith diameter and density, assuming fixed spacecraft properties as given above. The figure also marks the two ‘nominal’ regolith surfaces whose θ_E values were highlighted in Table 4.8. The trends observed in this figure are very similar for the three shapes: θ_E varies only slightly with changing regolith bulk density ρ_t , but is much more strongly dependent on the regolith grain size D_t . Larger particles have higher θ_E , such that the range of θ values over which $e = e_0$ is relatively small. This is better illustrated in Fig. 4.76, which plots the coefficient of restitution, modulated by the effective impact angle, following Eq. 2.113. This is done for the three tested spacecraft shapes and for both regolith surfaces.

This figure illustrates that, with the current model, the coefficient of restitution is, on average, higher for larger regolith particles. This model makes some intuitive sense: when the regolith particles are larger, they are more resistant to being moved around by an impacting spacecraft. As a result, fewer grain-grain interactions occur, such that the amount of energy dissipated by




Figure 4.75: Critical impact angle variation for the three tested shapes.




Figure 4.76: Restitution modulation in the two tested regolith surfaces.

the regolith is smaller. We now use our simulation framework to investigate regolith presence, as mimicked by this simple e -modulation model, affects the dynamics of a deployed spacecraft.

Probe dynamics on regolith Prior to simulating spacecraft motion on a regolith-covered surface, using the restitution modulation of Eq. 2.113, some relevant statistics about the motion on a hard surface layer are provided. In these, e is set to be a constant. For this, several scenarios of 2,000 simulations are initialized just above Itokawa's northern pole. In each simulation, the considered probe is given a random attitude and a negligible angular velocity ω . This is representative of a realistic release scenario: although probes are released by a mothership with some known attitude and negligible angular velocity, the time to first impact is often on the order of tens of minutes. This allows even a small angular velocity to result in a seemingly random attitude at impact. The nominal linear velocity \mathbf{V} is set to some fixed magnitude, though multiple scenarios are generated

in which the direction of \mathbf{V} varies. This is shown in Fig. 4.77, which illustrates the *incidence angle* γ of the probe velocity relative to the local surface normal $\hat{\mathbf{N}}$. In other words, the incidence angle is defined as:

$$\gamma = \arccos \left(\frac{-\mathbf{V} \cdot \hat{\mathbf{N}}}{\|\mathbf{V}\|} \right) \quad (4.18)$$

Note that this incidence angle is different from the surface inclination angle, β , defined before between the local surface acceleration and the local surface normal. While the incidence angle depends on the linear probe velocity \mathbf{V} , the impact angle θ instead depends on the net contact point velocity \mathbf{v} , as defined in Section 2.6:

$$\theta = \arcsin \left(\frac{-\mathbf{v} \cdot \hat{\mathbf{N}}}{\|\mathbf{v}\|} \right) \quad (4.19)$$




Figure 4.77: Definition of the incidence angle γ .

The reason for generating scenarios with different γ angles is simple: given that the initial angular velocity of the probes is small, the impact angle θ will be nearly equal to the complement of γ . As proof for this claim, consider the results of Fig. 4.78, which plots the distribution of the initial impact angle θ_1 for different incidence angles γ , for the tetrahedron, cube, and dodecahedron. Not only do we see the clear effect of γ on θ_1 in these scenarios, but we also observe that the θ_1 distributions are nearly identical for the three shapes. This makes intuitive sense, given that the impact angle θ only depends on the probe shape through the inclusion of the collision vector \mathbf{r} , since $\mathbf{v} = \mathbf{V} + \boldsymbol{\omega} \times \mathbf{r}$. However, given that the angular velocity $\boldsymbol{\omega}$ is small, it holds that θ is almost entirely determined by the linear velocity \mathbf{V} . Given that the three spacecraft shapes are simulated with the same initial conditions, it makes sense for them to also have the same θ_1 distribution.




Figure 4.78: Initial impact angles for different shapes and incidence angles.

More relevant to the discussion are the statistics of the impact angle at all successive impacts, denoted θ_2 . These are included in Fig. 4.79. In this figure, the critical θ_e of the two regolith models are marked with dashed lines, corresponding to the values of Table 4.8. Reviewing this figure, some relevant trends are found. First of all, the θ_2 statistics of the dodecahedron indicate that most of its collisions occur with an impact angle of $\theta > 80$ deg, in other words, that the dodecahedron tends to collide nearly normal to the surface. The majority of impacts occur with $\theta > \theta_{e,1.0}$; virtually all impacts occur with $\theta > \theta_{e,1.0}$. As a result, the following is expected when including the e -modulation of the regolith model: given that most impacts occur at $\theta > \theta_e$, they will occur with $e = e_0$. The resulting simulations are thus expected to be highly similar to simulations performed with a fixed $e = e_0$.

For the cube and tetrahedron, the situation is more complex. The θ_2 statistics of the cube continue to display a significant number of impacts with $\theta_2 > 80$ deg, though not as dramatically as the dodecahedron. Not only does it display more impacts with $\theta < 80$ deg, a second peak in the impact angle distribution is also observed at $\theta_2 < 10$ deg. This indicates that the cube

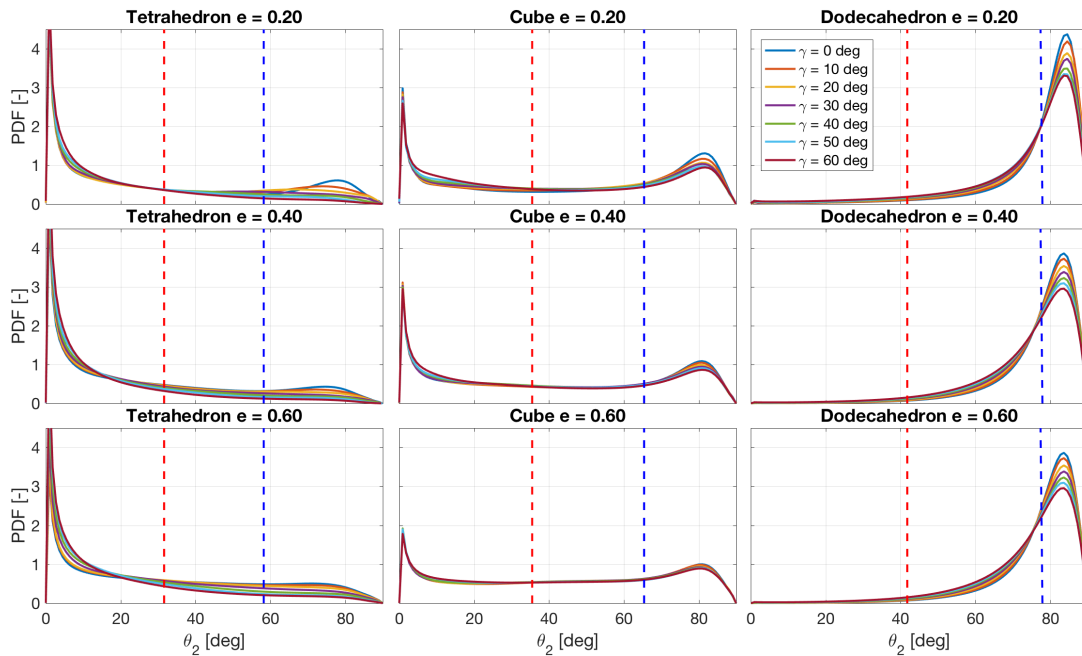


Figure 4.79: Successive impact angles for different shapes and γ , for deployment with fixed e .

also experiences a notable number of nearly tangential impacts with the surface. As a result, the application of the regolith e -modulation is expected to strongly bias simulations towards shorter settling times, given that most impacts will occur with $e = e_0$ as given by Eq. 2.113. This effect is anticipated to be stronger for the tetrahedron than for the cube. Indeed, reviewing the θ_2 distribution of the tetrahedron, it is found that many of its impacts occur with $\theta < \theta_e$. For low incidence angles, a secondary peak does occur at $\theta \sim 80$ deg, though it is relatively small. We thus expect the tetrahedron dynamics to be more significantly affected by the presence of regolith than the cube.

We now proceed with simulations that include the regolith e -modulation model. We repeat the various γ scenarios of the three shapes, while using $e_0 = 0.2$ and $e_1 = 0.5$, for both regolith models. The critical θ_e angles are taken from Table 4.8. The resulting mean settling time values μ_{t_f} are shown in Fig. 4.80. This figure also includes the settling times of the three fixed- e scenarios from which the incidence angle statistics of Figs. 4.78 and 4.79 were obtained. As expected, the tetrahedron settles with shorter t_f than the cube, which in turn settles with shorter t_f than the

dodecahedron. This matches the results previously shown in Section 4.2. The figure also allows us to quantify the effect of the incidence angle γ . For both the cube and the dodecahedron, it is found that t_f increases with increasing γ , regardless of the applied coefficient of restitution, fixed or modulated. The effect appears to be stronger for the dodecahedron than it is for the cube. The situation is different for the tetrahedron, which appears to be only very slightly affected by γ ; the variations in t_f are barely noticeable at this scale. Curiously, it is found that the settling time actually *increases* with increasing γ with fixed $e = 0.60$.




Figure 4.80: Settling time statistics for relevant deployments to hard and regolith surfaces.

We now turn our eye from the fixed- e simulations to those with modulated e . In the case of the dodecahedron, it is found that the t_f statistics of both regolith surfaces very closely matches that of the scenario with fixed $e = 0.20$. This matches the prediction made based on the θ statistics: given that most of the dodecahedron impacts occur at $\theta > \theta_e$, it must hold that they occur with $e = e_0$. This is precisely what is observed in Fig. 4.80: the two modulated- e curves almost exactly match the $e = e_0 = 0.2$ curve, for all incidence angles γ .

For the cube, it is found that the regolith curves lie between the $e = 0.2$ and $e = 0.4$ curves at low γ , but that they approach the $e = 0.4$ curve at high γ . This also agrees with our predictions: given that the cube experiences both high and low θ , it is expected to experience an effect from the regolith presence. Given that high θ_1 tends to occur when γ is low, and that the first impact often dissipates the most energy of any of the impacts encountered along a trajectory, it also makes sense that more tangential γ results in longer settling times. Finally, it is found that tetrahedron is most strongly affected by the presence of regolith. Its modulated- e curves coincide with the $e = 0.4$

curve at low γ and grow towards the $e = 0.6$ curve at high γ . This indicates that, regardless of the initial incidence angle γ , the tetrahedron experiences a relatively large number of collisions in which $e > e_0$. This agrees with the θ trends of Fig. 4.79 which indicated that many of the tetrahedron impacts occur with $\theta < \theta_e$ and are therefore quite shallow. On a final note, it is found for all three shapes that the results from the two different regolith models are very similar, with both the 1.0 cm and 2.5 cm curves nearly overlapping.

Effect of friction We note that the above simulations were performed with intermediate friction coefficient $f = 0.50$. In Section 4.2, it was shown that all three (solid variants of the) shapes are *unsaturated* with respect to friction at this f value. In other words, at $f = 0.50$, many of the impacts experienced by the three shapes have some residual sliding velocity at their conclusion. To investigate whether their behavior in the presence of regolith changes with higher f , the simulations with $f = 1.0$ are repeated. At this friction coefficient, both the cube and dodecahedron collisions are saturated; all collisions conclude in stick with $s = 0$. Only the tetrahedron is able to have residual sliding velocity at this value of f . We plot the corresponding θ_2 statistics in Fig. 4.81.

Comparing the θ statistics of the $f = 1.0$ scenarios with those shown previously in Fig. 4.79 for the $f = 0.5$ scenarios, the following is observed. The θ statistics of the dodecahedron are nearly identical in the two scenarios, suggesting that it will continue to experience mostly normal impacts even with higher f . For the cube, the situation is different: while it experienced both tangential and normal impacts under $f = 0.5$, it is found from Fig. 4.81 that it mostly experiences normal impacts with $f = 1.0$. In other words, its θ statistics have become very similar to that of the dodecahedron. We therefore expect the cube with $f = 1.0$ to display deployment statistics that closely approach those of fixed $e = e_0$ simulations. Finally, it is found that the θ distribution of the tetrahedron with $f = 1.0$ has become very similar to that of the cube with $f = 0.5$. We therefore expect its t_f statistics to match the cube with $f = 0.5$.

Reviewing the settling time statistics with $f = 1.0$ in Fig. 4.82, it is found that its results match our predictions based on the θ distribution. All three shapes are now heavily affected by the presence of the regolith, since most impacts occur nearly normal. As a result, the t_f curves match those of the $e = e_0 = 0.2$ scenarios.

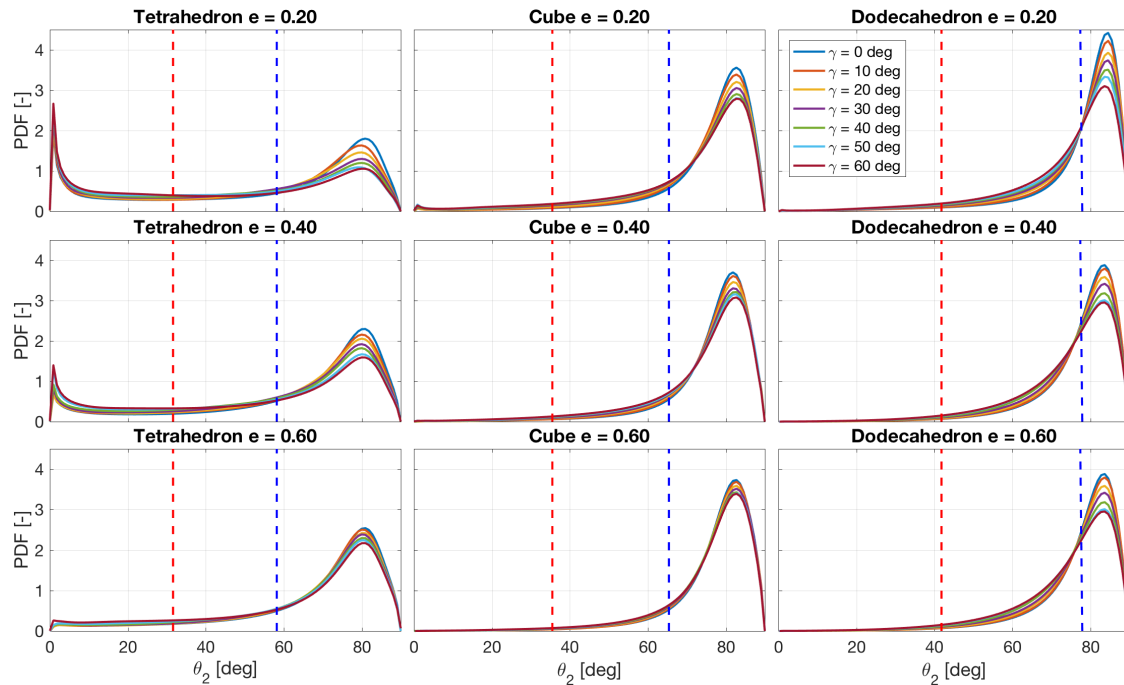


Figure 4.81: Successive impact angles for different shapes and γ , deployment with fixed e .




Figure 4.82: Settling time statistics for relevant deployments to hard and regolith surfaces.

Summary We account for the presence of regolith on the small-body surface by following the model by Nishida et al. [80]. They observed that normal impacts of spherical particles into a granular matter occur at lower coefficients of restitution than tangential impacts. They derived an

expression for the impact angle at which the transition occurs, as previously given in Eq. 2.112. Using this, a simple modulation of the coefficient of restitution is implemented, that applies some low e_0 for normal impacts and transitions to some higher e_1 at more tangential impacts. We applied this model to test two different regolith models on asteroid Itokawa: one with 1.0 cm regolith grains and one with 2.5 cm grains. We performed simulations of a tetrahedron, cube, and dodecahedron deployed to these regolith-covered surface with modulated e , and compared them with simulations performed to a hard surface layer with fixed e .

At high (saturation) friction coefficients, it is found the deployment dynamics of all three shapes to match those of simulations performed with a fixed lower boundary e_0 . This occurs because the majority of the probe collisions occur at nearly-normal impact angles for high friction coefficients. All three shapes are thus affected equally by the presence of regolith; they all experience the same (low) restitution coefficient. At lower friction coefficients, in which the contact point sliding velocity is not reduced to zero, the behavior is different. Although spherical shapes (the dodecahedron) continue to exhibit mostly normal impacts, less spherical shapes (the tetrahedron and cube) experience more tangential impacts. As a result, the restitution coefficients experienced by the three shapes are different. Spherical shapes continue to experience particularly low restitution, while less spherical shapes experience a greater range of high-to-low restitution values.

Although these results are strongly model-dependent, one general conclusion persists from this first-attempt at accounting for the presence of regolith on the small-body surface. If effectively normal impacts are more strongly damping than effectively tangential impacts, it holds that shapes that are more spherical will experience lower effective restitution coefficients than shapes that are less spherical. This would limit the ‘disadvantage’ that spherical shapes have compared to more distorted shapes, with the settling time of the former shrinking and more closely approaching that of the latter.

Chapter 5

Case study: Hayabusa2 rovers

The Japanese Hayabusa2 asteroid sample return mission was launched in December of 2014. The spacecraft is currently in good health and is expected to arrive at its target, asteroid 162173 Ryugu, some time between June and July of 2018 [124]. Following arrival, the spacecraft will perform approximately one and a half year of proximity science operations at the asteroid. During this time, Hayabusa2 will deploy three rovers: the Japanese MINERVA-II-1A, MINERVA-II-1B, and MINERVA-II-2; as well as the European MASCOT lander. All rovers are equipped with internal momentum exchange mechanisms, which allow them to hop across Ryugu's surface and perform scientific investigations at several sites [43, 54]. The spacecraft will also descend to Ryugu's surface three times to obtain samples that will be returned to Earth for further study, in November or December of 2020. One of the surface samples will be obtained from a crater created by the spacecraft's explosive Small Carry-on Impactor (SCI), allowing for the acquisition of freshly exposed sub-surface material [95].

During descent to the surface, the Hayabusa2 spacecraft performs autonomous surface-relative navigation using its lidar, laser range finder, and target marker tracking. The target markers are small retroreflective passive devices that reflect the light from a flash lamp (FLASH) onboard the mothership. Given that the first rover deployment takes place *before* the first sampling operation, there is a risk of the rover presence interfering with FLASH. The deployment should therefore be designed in a way that provides confidence the rover will avoid the exclusion zone around the Hayabusa2 sampling site(s). This in turn requires high-fidelity simulation of the rover

motion. An accurate prediction of the motion may also help plan the camera scanning program used to track the rover descent and aid the scheduling of rover-spacecraft communications. The dynamics of rovers operating in the small-body environment are complex, due to the bodies' irregular gravity, shape, and surface interactions. Rosetta's Philae lander demonstrated that uncontrolled probes are prone to bounce several times in such an environment, during which they may cover a significant distance across the small-body surface [11]. Simulations of rover deployment to a small body must be performed at sufficient fidelity in order to account for these various complexities. Unfortunately, relatively little information is currently known about the shape, mass, and rotation of asteroid Ryugu, which makes it difficult to produce accurate deployment predictions and relevant deployment analysis prior to the arrival of Hayabusa2. Nevertheless, it is important to perform a *pre-arrival* deployment analysis, in order to reveal challenges and opportunities in the rover deployment. In this chapter, our simulation framework is applied to perform this pre-arrival analysis. This will serve as a baseline for rigorous comparison with the 'real' deployment analysis to be performed following arrival at Ryugu. The results have also adjusted the nominal deployment operations for the Hayabusa2 mission.

Our analysis consists of two parts: in the first part, deployment to Ryugu's reference sphere is simulated to obtain order-of-magnitude estimates of the deployment statistics, and to examine the effects of relevant environmental parameters. In the second part, deployment simulations are performed using Ryugu's training model that was developed in order to rehearse the Hayabusa2 mission operations. The training model is constructed from a low-resolution lightcurve shape model, with artificially added rocks, boulders, and other topography that is realistically expected on Ryugu's actual surface [68, 71]. The chapter is structured as follows. Section 5.1 presents the characteristics of the asteroid, lander, and release sequence that are relevant to the rover deployment. The simulated deployment to Ryugu's reference sphere is then reviewed in Section 5.2, which concludes with a discussion on the suggested inclusion of a horizontal pre-release maneuver. The lessons learned from these relatively simple simulations are then applied to the training model in Section 5.3, which examines how lander release at two epochs and above different terrain types

yields varying deployment statistics. Section 5.4 then investigates the effects of variations in the asteroid mass on the rover deployment. Finally, Section 5.5 concludes the pre-arrival analysis by summarizing the challenges and opportunities of selecting the release scenario to be carried out. This summary also underscores how the completed analysis has adjusted the nominal deployment operations of Hayabusa2.

5.1 The Hayabusa2 mission

5.1.1 Asteroid Ryugu

Asteroid 162173 Ryugu orbits the Sun at a semi-major axis of $a \approx 1.2$ AU, with eccentricity $e \approx 0.19$ and inclination $i \approx 5.88^\circ$ [40]. Due to the low amplitude variations in its visual lightcurves, an Earth-based determination of Ryugu’s shape and spin properties is challenging. Müller et al. combined various data types and sets with thermophysical models to produce best estimates for these asteroid properties. They estimate Ryugu’s effective diameter at $R = 850 - 880$ m, with a rotation period of $P = 7.63107$ hr and an ecliptic rotation axis orientation of $(\lambda, \beta)_{ecl} = (340^\circ, -40^\circ)$. The corresponding best-fit shape model is shown in Fig. 5.1 [76].




Figure 5.1: The formally best-fit shape model of Ryugu, taken from Müller et al. [76].

This shape model is of low resolution and lacks detailed surface features such as boulders, craters, and cliffs that are likely to exist on the ‘real’ Ryugu. As such, the best-fit shape model has limited applicability for simulation and mission rehearsal purposes. The Japan Aerospace

Exploration Agency (JAXA) has therefore developed a *training model*, by adding realistic surface detail to the best-fit shape model, based on observations of asteroid Itokawa. The resulting model, shown in Fig. 5.2, is being used at JAXA to carry out a broad rehearsal for the Hayabusa2 proximity operations at Ryugu. Please note that the training model axis system differs from that shown in Fig. 5.1. The training model has an approximate volume of $V = 0.316 \text{ km}^3$, which corresponds to a reference sphere radius of $R = 423 \text{ m}$. All other relevant asteroid properties have been fixed at realistic values to enable full mission rehearsal. The training model is in uniform rotation around its Z -axis with a period of $P = 7.631 \text{ hr}$; the rotation pole is oriented at $(\lambda, \beta)_{J2000EQ} = (13.67807^\circ, -34.77563^\circ)$. Given the uncertainties in Ryugu's true size and density, the training model gravitational parameter μ is assumed to vary uniformly as $\mu = \mathcal{U}(11, 92) \text{ m}^3/\text{s}^2$, with a nominal value of $\mu = 32 \text{ m}^3/\text{s}^2$.




Figure 5.2: The Ryugu training model.

5.1.2 Mothership and rovers

The Hayabusa2 spacecraft measures approximately $2.0 \times 1.6 \times 1.3 \text{ m}$ and has a wet mass of 600 kg. Upon arrival at Ryugu in June or July of 2018, the spacecraft will move into *home position* (HP), located 20 km above the asteroid and along the sub-Earth line. It will hover at this position for most of the 1.5-year proximity operation phase; Hayabusa2 will nominally *not* enter orbit about Ryugu. Throughout proximity operation, the spacecraft maintains a fixed attitude in which its Z_{SC} axis is aligned with the Earth-Ryugu line. The X_{SC} axis is oriented such that the

Sun-Ryugu vector lies within the $X_{SC} - Z_{SC}$ plane. Both this axis alignment and the spacecraft frame definition are illustrated in Fig. 5.3.




Figure 5.3: (left) Rear and (right) side view illustrations of the Hayabusa2 spacecraft at home position.

The mothership carries three rovers developed by a consortium of Japanese universities: the MINERVA-II-1A and -1B rovers are encased in a single protective fairing and deployed together along the $+Y_{SC}$ axis. The MINERVA-II-2 rover is separately released along the $-Y_{SC}$ axis. The initial positions of these rovers are marked with the numbers ‘1’ and ‘2’ in Fig. 5.3. The spacecraft also deploys the European lander MASCOT (marked with the letters ‘MT’), though this chapter is only concerned with the deployment of the MINERVA-II rovers. The deployment of MASCOT is under investigation by the European centers CNES and DLR [64]. The rovers are released at different initial velocities V_0 and angles γ . Note the large uncertainty on the release velocity of MINERVA-II-2, which varies between 5 and 20 cm/s. In comparison, the release velocity of MINERVA-II-1A/B is fixed at 20 cm/s. Table 5.1 specifies these values, as well as the respective rover masses m , dimensions, shapes, and inertias $I_{x,y,z}$. Fig. 5.4 illustrates the polygonal shape models used to represent these deployable payloads at the same scale. The rover inertias were computed from these shape models by assuming a uniform density [29].




Figure 5.4: Applied shape models of the MINERVA-II rovers.

Table 5.1: Physical properties of the MINERVA-II rovers and their release.

Rover	Size [mm]	m [kg]	Shape	V_0 [cm/s]	γ [deg]	$I_x = I_y$ [kg·m ²]	I_z [kg·m ²]
MINERVA-II-1A	225 × 225 × 103	1.2	Hexadecagonal	20	5	0.0047614	0.0074011
MINERVA-II-1B	225 × 225 × 103	1.1	Hexadecagonal	20	5	0.0043646	0.0067843
MINERVA-II-2	175 × 175 × 200	0.88	Octagonal	$\mathcal{U}(5 - 20)$	5	0.0080968	0.0055270

5.1.3 Rover release

As mentioned above, the Hayabusa2 spacecraft maintains a fixed position at HP, 20 km above the sub-Earth point on Ryugu, with the asteroid rotating underneath. When viewed in an asteroid-fixed frame over the course of the 7.631 hr rotation period, the spacecraft thus maintains some latitude ϕ and sweeps out all longitudes λ . Since Earth and Ryugu have notably different orbits, the HP latitude ϕ slowly changes over the course of the proximity operation phase. By combining the (true) orbital ephemerides of the Sun, Earth, and Ryugu with the assumed rotation pole orientation, the variation of the sub-solar and sub-Earth latitudes on Ryugu can be computed, as plotted in Fig. 5.5.

The three MINERVA-II rovers are deployed during two release events that are carried out at different epochs during the 1.5-year proximity operations. MINERVA-II-1A and MINERVA-II-1B are released together; MINERVA-II-2 is deployed in a separate event. The order in which these releases will be performed has not yet been decided. Furthermore, the actual rover deployment




Figure 5.5: Variation of the sub-solar and sub-Earth latitudes on the Ryugu training model.

schedule is subject to the asteroid environment. The release epochs used in this analysis are assumptions that enable simulation, with the first release epoch defined as September 20, 2018 and the second epoch defined as July 1, 2019. Both of the epochs are marked in Fig. 5.5 together with the corresponding sub-solar and sub-Earth latitudes. These latitude values can be used to plot the quasi-orbit that Hayabusa2 sweeps out over the course of one Ryugu rotation period. This is shown in Fig. 5.6 for both release epochs. As this figure shows, the spacecraft has access to Ryugu’s northern hemisphere during the first release epoch, and to the southern hemisphere during the second epoch. Although the particular value of these latitudes is dependent on Ryugu’s real pole orientation, this dichotomy in hemisphere accessibility must also occur during the actual proximity operations.

At the start of the release sequence, Hayabusa2 descends from HP along the Earth-Ryugu line at a constant descent velocity of 10 cm/s. This continues until the spacecraft lidar registers an altitude of 60 m above Ryugu’s surface, upon which a deceleration maneuver of $\Delta V = 7$ cm/s is triggered to reduce the spacecraft descent velocity to 3 cm/s along the Earth-Ryugu line. As the release sequence can be started at any longitude λ around Ryugu, there exists a line of possible trigger altitudes that lies 60 m above the surface, as illustrated in Fig. 5.5, together with its ground track projected along the Earth-Ryugu line. Once the deceleration maneuver has been triggered, Hayabusa2 coasts for 140 s while waiting for its attitude control to converge. The MINERVA-II rover(s) is/are then released with their respective release velocities as given in Table 5.1. The




Figure 5.6: Hayabusa2 quasi-orbit about the Ryugu training model, at both release epochs.

rover release is triggered through the activation of a pyro wire cutter that opens the cover of the release mechanism. The rover is then released using a spring, with an initial velocity equivalent to the elastic energy of the spring. As MINERVA-II-1A/B and MINERVA-II-2 have different shapes, masses, and release mechanisms, they have different initial velocity errors. Following release, the mothership continues to coast for another 60 s before performing an ascent maneuver, in order to prevent plume contamination of the rover(s). Table 5.2 repeats these values and includes the corresponding 3σ errors.

Ryugu's gravitational field is captured using the constant-density polyhedron model. A volume-conserving quadratic edge collapse technique can be applied to reduce the high-resolution 1.5-million facet shape model into a low-resolution 1,000-facet model. Both models are shown in Fig. 5.7. This resolution is chosen such that the resulting model has a maximum error of $< 1\%$ in the norm of its gravitational acceleration, compared to the high-resolution model. Additionally, gravity field linearization is applied to further reduce the cost of gravity field evaluations, as discussed in Section 2.4.




Figure 5.7: (left) High-resolution and (right) reduced-resolution training model used for gravity evaluations.

Table 5.2: Release sequence parameters and 3σ errors.

Parameter	Nominal value	3σ error
Horizontal H2 position	0 m	28.4 m
Vertical H2 position	(determined by lidar)	(determined by lidar)
Horizontal H2 velocity	0 cm/s	1.4 cm/s
Vertical H2 velocity	-10 cm/s	1.7 cm/s
H2 Attitude error	0 deg	0.1 deg per axis
H2 Attitude rate error	0 deg/s	0.01 deg/s per axis
Lidar trigger altitude	60 m	3 m
Decelerating ΔV	+7 cm/s	1.05 cm/s
M-II-1A/B release velocity	20 cm/s	Unknown, assumed zero
M-II-2 release velocity	12.5 cm/s	Uniform ± 7.5 cm/s
M-II-1A/B release direction	-5 deg from $+Y_{SC}$	Unknown, assumed zero
M-II-2 release direction	-5 deg from $-Y_{SC}$	Unknown, assumed zero

5.2 Deployment to spherical Ryugu

Before investigating the MINERVA-II deployment dynamics on the training model, it is useful to better understand the sensitivities of deployment to the surface interaction coefficients and the release geometry. More specifically, it will be investigated how the coefficients of restitution and friction, as well as the release latitude and azimuth, affect the observed deployment statistics. In order to perform this *sensitivity analysis* in a simple but relevant environment, rover deployment to the Ryugu reference sphere is simulated. This is the sphere whose volume is equivalent to that of the training model. As mentioned before, the training model has an approximate volume of

$V = 0.316 \text{ km}^3$; this corresponds to a sphere of radius $R = 423 \text{ m}$. The simulations are performed with a nominal release latitude of $\phi = 45 \text{ deg}$ and a nominal azimuth of $\alpha = 90 \text{ deg}$. The coefficients of restitution and friction are fixed at respectively $e = 0.50$ and $f = 0.50$. The lander release velocity is fixed at $V_0 = 20 \text{ cm/s}$. An overview of these nominal parameters is provided in Table 5.3.

Table 5.3: Nominal parameters used in Ryugu reference sphere simulations.

Parameter	e [-]	f [-]	ϕ [deg]	α [deg]	V_0 [cm/s]	μ [m^3/s^2]	R [m]
Value	0.50	0.50	45	90	20	32	423

The reference sphere is illustrated on the left side of Fig. 5.8, together with results from 1,000 sample simulations. These simulations are performed following most of the steps of the release strategy discussed in Section 5.1.3, using the error model of Table 5.2 and the MINERVA-II-2 shape model shown in Fig. 5.4. However, the constraints on the mothership position and attitude are removed, such that one is free to choose the mothership latitude ϕ and the rover release azimuth α . The latter parameter expresses the direction in which a rover is released, with an azimuthal direction of $\alpha = 0 \text{ deg}$ indicating a release with the $+Y_{SC}$ axis pointing due north, $\alpha = 90 \text{ deg}$ indicating a release with the axis pointing due east, etcetera.




Figure 5.8: Sample simulation batch to the Ryugu reference sphere, and corresponding settling time statistics.

In Fig. 5.8, the release, first impact, and settling positions are marked in respectively red, green, and blue. Also shown are the three trajectories that have respectively the shortest, median,

and maximum settling time. This sample simulation batch is reviewed in order to provide details on our process for analyzing deployment simulations, in which we are primarily interested in the settling time t_f and surface dispersion A_d statistics. The settling time statistics are easily found from the final epoch of each simulation and are unimodal in form. This can be seen in the settling time histogram on the right side of Fig. 5.8. Due to the structure of this distribution, it can be more represented more compactly using a simple boxplot, as shown on the bottom of the figure. This will be done for all other simulation batches to be analyzed in this chapter.

Although the surface dispersion can be visually inferred from plots such as that on the left side of Fig. 5.8, a quantification of the dispersion area A_d is not trivial. For this, the reference sphere is tessellated into facets of near-equal area; it is then counted how many simulations settled on each facet. To obtain the 1σ , 2σ , and 3σ dispersions, we add together the areas of the smallest set of facets that together contain respectively 68%, 95%, and 99% of the settling locations. The dispersion of the sample simulation batch is illustrated in Fig. 5.9, in which the two-dimensional plot is generated using the equal-area Mollweide projection. Note that the dispersion area values generally provide a *lower bound*, as is clear from the plot on the right side of Fig. 5.9. In this, it is seen that the 2σ and 3σ areas contain some empty facets, on which one likely would see some rover settle if additional simulations were performed.




Figure 5.9: Surface dispersion of the sample simulation batch.

5.2.1 Effect of surface interaction coefficients

Section 4.2 showed how the coefficients of restitution e and friction f affect the motion of probes with different shapes. More specifically, it was found that the coefficient of restitution governs the motion, with variations in e scaling both the settling time and the surface dispersion. Although it was found that a probe shaped like MINERVA-II-2 had deployment statistics similar to that of the cube when deployed to asteroid Itokawa, we want to investigate whether these trends hold when simulating deployment to Ryugu instead, which is about twice the size of Itokawa. Furthermore, by quantifying the effects of e and f on the rover motion, the accuracy to which the values of these coefficients should be determined for asteroid Ryugu can be determined.

Coefficient of restitution e First, the effects of variations in the coefficient of restitution e are investigated. For this, the nominal simulation batch shown before is repeated, for different values of e , while maintaining all other environment parameters at their nominal values as given in Table 5.3. The resulting settling time and surface dispersion statistics are shown in Fig. 5.10.

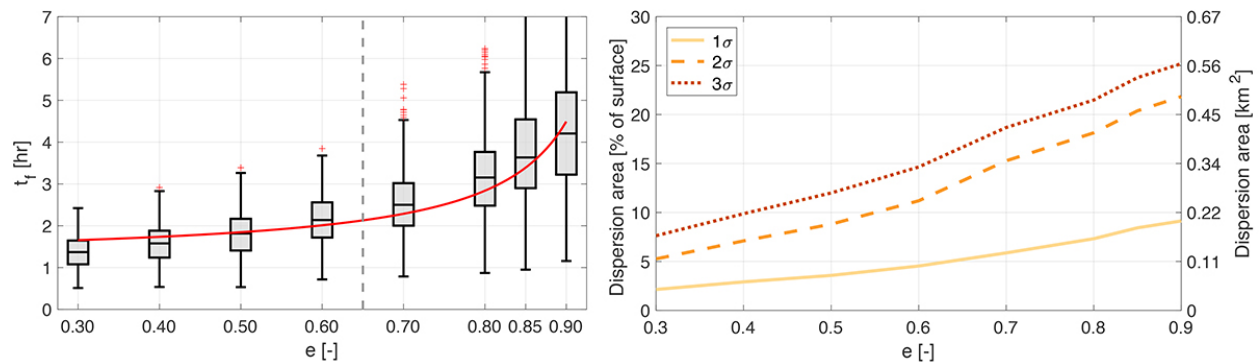


Figure 5.10: Deployment statistics for e variation on Ryugu reference sphere.

In agreement with the results from Section 4.2, the coefficient of restitution strongly affects the rover settling time as seen on the left side of Fig. 5.10, with lower e values resulting in shorter settling times. The changes in t_f are relatively benign when $e < 0.65$, with the mean t_f varying between approximately 1 and 2 hours. When $e > 0.65$, the increase in t_f becomes much stronger, with the mean t_f increasing to beyond 4 hours at $e = 0.90$. The standard deviation of t_f also grows

significantly when $e > 0.65$, such that many simulations have a settling time much longer than 4 hours. This again agrees with the trends revealed in Section 4.2, that indicated that the settling time of an object scales with $t_f \propto (1 - e)^{-1}$. The best-fit of this proportionality to the observed mean t_f is shown in red on the left side of Fig. 5.10. The trends in this curve roughly match the observed behavior, where the settling time increases sharply beyond $e > 0.65$. Reviewing the effects on the surface dispersion on the right side of Fig. 5.10, similar scaling effects are found where higher e values result in larger surface dispersions. Although the slope of the dispersion curves does increase beyond $e > 0.65$, the increase is not as dramatic as for the settling time. These results predict that the rover landing ellipse and settling time remain appreciably small when $e < 0.65$. In order to estimate what range of e values may be expected on asteroid Ryugu, relevant results from previous small-body exploration missions are reviewed, as well as from simulation and experiment.

The Hayabusa spacecraft touched down on the surface of asteroid Itokawa in November 2015. The coefficient of restitution of Itokawa's surface was estimated using Hayabusa altitude measurements, which in turn were derived from laser rangefinder data. Yano et al. used this to estimate a coefficient of restitution of $e \approx 0.84$ by fitting a free-fall parabola to the altitude measurements, although there is a notable measurement gap of approximately 3 minutes centered at the collision epoch [145]. They were also unable to establish an accurate estimate of the horizontal velocity of the spacecraft. This may help explain the relatively low quality of the parabolic fit after the impact; see also Fig. 5.11. This figure was taken from Yano et al. and edited to indicate the measurement gap and the low-quality parabolic fit on ascent after the first impact. The accuracy of the relatively high value of $e \approx 0.84$ is therefore questionable.

The Philae lander onboard the Rosetta spacecraft touched down on the surface of comet 67P/C-G in November 2014. Due to failures in its landing systems, the lander performed several bounces on the comet surface. Using an energy balance analysis, Biele et al. determined the cometary surface to have a coefficient of restitution of $0.20 < e < 0.50$ [13]; notably lower than the value estimated for Hayabusa. Finally, granular mechanics simulations of the MASCOT touchdown using *pkdgrav* code estimate a coefficient of restitution of $0.1 < e < 0.6$ [64], while experimental




Figure 5.11: Altitude profile of the Hayabusa spacecraft during touchdown [145].

results for this impact in a gravity-offloading setup measured values of $0 < e < 0.6$ [12]. All of the mentioned estimates are visually combined in Fig. 5.12. Excluding the questionable e value of Hayabusa, this data suggests that the coefficient of restitution on Ryugu may indeed be expected to be $e < 0.65$. This provides some careful optimism about the MINERVA-II dispersion area.




Figure 5.12: Relevant estimated coefficient of restitution values in literature [64, 145, 13, 12].

Coefficient of friction f To examine the effects of the coefficient of friction on the rover deployment, the same strategy applied above is followed, by repeating the simulations to the Ryugu reference sphere for varying values of $0.30 < f < 1.00$, while maintaining all other parameters at their nominal values as given in Table 5.3. The resulting settling time and surface dispersion statistics are shown in Fig. 5.13.

Reviewing this figure, it can be seen (on the left side) that the mean settling time increases




Figure 5.13: Deployment statistics for f variation on Ryugu reference sphere.

with higher f values, though the effect is of much smaller magnitude than previously observed for the coefficient of restitution. The variation in mean settling time between the lowest and highest tested f values is less than one hour. The standard deviation in t_f remains approximately constant between all tested values. A saturation limit appears at $\bar{f} \sim 0.80$, beyond which further increases in f do not have a notable effect on the settling time. This behavior again agrees with that observed in Section 4.2 when examining the deployment dynamics of probes with a complex shape. Again, it is found that greater f results in longer settling times and larger surface dispersions.

Although the coefficient of restitution has been estimated in a few relevant environments, very little information is available on expected values for the coefficient of friction. For the touchdown of the Hayabusa spacecraft, it was estimated as $f > 0.80$ [145]. The precise value of f for the Philae lander is unknown; Roll and Witte used a range of $0.2 < f < 1.0$ in their attempts to reconstruct the lander's trajectory [94]. In their simulation work with MASCOT, Maurel et al. set the rover's coefficient of friction equal to that of the applied grain-grain friction coefficient of $f = 1.0$ [64]. Lunar regolith models suggest grain-grain friction coefficient values of $0.36 < f < 1.20$ [141]. It must be noted that these grain-grain friction coefficient values relate to distinctly different behavior than the object-surface friction coefficient used in our contact models. Thankfully, since our results show that f has a minor effect on the deployment statistics, an approximate estimate of its value suffices.

Combined effects of e and f The separate effects of the coefficients of restitution and friction were revealed above. To investigate whether any strong coupling effects between e and f exist, which may be obscured in a single-parameter variation, the reference sphere simulations are repeated while mutually varying $0.30 < e < 0.80$ and $0.30 < f < 0.60$. The mean settling time and 3σ surface dispersion values are plotted in Fig. 5.14. These results match the previously observed single-parameter trends (as well as those discussed in Section 4.2) and confirm that the coefficient of restitution remains the primary determinant of the rover deployment statistics, though friction does play a non-negligible role.

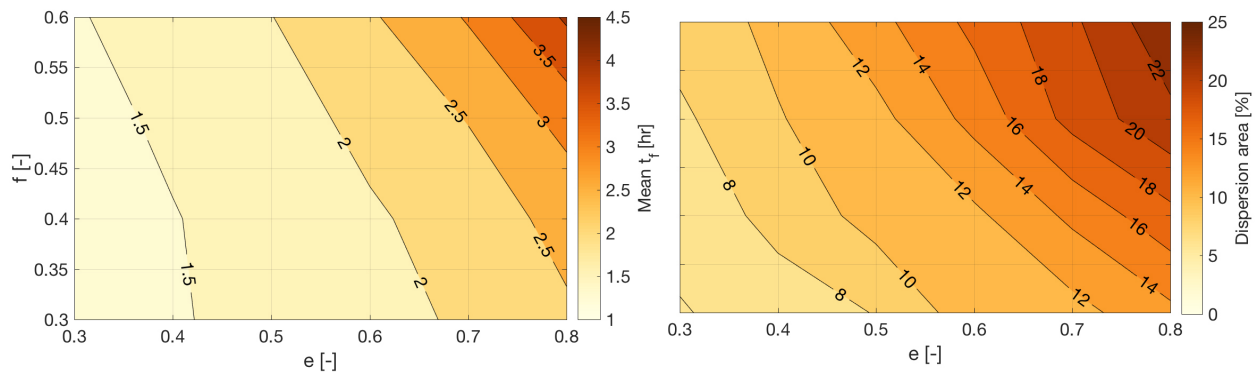


Figure 5.14: (left) Mean t_f and (right) $3\sigma A_d$ for (e, f) variation on the Ryugu reference sphere.

5.2.2 Effect of release geometry

In the nominal scenario, rovers are released at a latitude of $\phi = 45$ deg above the reference sphere with an initial velocity directed due east, *i.e.*, with an azimuth of $\alpha = 90$ deg. Given that Hayabusa2 maintains a fixed attitude relative to Sun and Earth, the values of ϕ and α are determined by the epoch of release. In this, it holds that the α values of MINERVA-II-1A/B and MINERVA-II-2 differ by 180 deg, since they are released in opposite directions along the Y_{SC} axis. To help inform the decision of which rover(s) should be released at which epoch, the effects of variations in ϕ and α on the resulting rover deployment statistics are investigated, in analogy with the coefficient variation discussed above. The corresponding settling time statistics are shown in Fig. 5.15.




Figure 5.15: Settling time statistics for (left) ϕ and (right) α variation on the Ryugu reference sphere.

Reviewing this figure, it can be seen that the release latitude ϕ has an almost negligible effect on the deployment statistics, with only a very slight decrease in mean t_f at very high latitudes. The release azimuth α , on the other hand, strongly affects the settling time. The minimum t_f occurs for rover release due east at $\alpha = 90$ deg; the maximum occurs when releasing due west at $\alpha = 270$ deg. The difference in mean t_f between the extrema is quite dramatic, with the maximum t_f more than twice the minimum t_f . Fig. 5.16 also plots the surface dispersions of the $\alpha = 90$ deg and $\alpha = 270$ deg cases. It can be seen that the release due west results in an extremely large dispersion that nearly covers the entire surface of the reference sphere.




Figure 5.16: Surface dispersion of rover releases due east and due west.

For context, Table 5.4 contains the azimuth values for both rovers at the two release epochs; these values were found using the training model pole orientation, similar to how Fig. 5.6 was created. Out of the four release azimuths listed in Table 5.4, only the Epoch 2 release of the MINERVA-II-2 rover at $\alpha = 45$ deg results in reasonably small settling times. If this deployment

were to be performed, MINERVA-II-1A/B would be released with $\alpha = 331$ deg. Such a deployment would result in large surface dispersion that makes it impossible to accurately aim for some surface region. Given this challenge, the source of the strong effects of α is further examined, which can be explained from the MINERVA-II velocity at first impact.

Due to the applied release sequence, the rovers always impact Ryugu's reference sphere with the same *normal* impact velocity v_N . However, the *tangential* impact velocity v_T is strongly dependent on the release azimuth α . More specifically, due to Ryugu's rotation, a release due east results in a lower surface-relative tangential rover velocity than a release due west. This is illustrated in Fig. 5.17, which plots the rover tangential impact velocity for varying release latitude ϕ and azimuth α .

Table 5.4: Release azimuth values of the MINERVA-II rovers at both release epochs.

Rover	Epoch 1	Epoch 2
MINERVA-II-1A/B	$\alpha = 151$ deg	$\alpha = 45$ deg
MINERVA-II-2	$\alpha = 331$ deg	$\alpha = 225$ deg




Figure 5.17: Variation of the tangential rover impact velocity v_T for $V_0 = 20$ cm/s.

This figure further demonstrates how the rover tangential impact velocity is minimized when releasing due east and is maximized when releasing due west. When the tangential velocity is higher, rovers can move across the asteroid surface for a longer amount of time, and can therefore cover greater distances. Similar results were also found for the MASCOT lander [112]. Even when ignoring the release azimuth, the rover’s fixed release velocity of $V_0 = 20$ cm/s is notably high, given that Ryugu’s nominal surface circular velocity is $\sqrt{\mu/R} = 27.5$ cm/s. Since the MINERVA-II rovers are released via a spring mechanism whose impulse cannot be changed in-flight, this high release velocity cannot be adjusted. An alternative method for reducing the rover tangential impact velocity is therefore investigated; we propose the inclusion of an additional pre-release maneuver with the Hayabusa2 mothership spacecraft.

5.2.3 Pre-release maneuver

As mentioned in Section 5.1, the Hayabusa2 spacecraft maintains a fixed vertical velocity of -10 cm/s during descent, until its lidar altimeter registers an altitude of 60 m above Ryugu’s surface. The spacecraft then performs a vertical ΔV of +7 cm/s to reduce its descent velocity to -3 cm/s. In order to limit the MINERVA-II tangential velocity at first impact, we propose combining this vertical maneuver with a horizontal *pre-release maneuver* (PRM) designed to cancel the tangential rover velocity at first impact. A single-shooting algorithm is used to compute the magnitude and direction of the PRM; these are plotted in Fig. 5.18

As expected, the maneuver is largest when releasing a rover due west, in which case the ΔV is applied towards the east, such that the net surface-relative lander velocity matches the surface rotational velocity. To verify that the use of the PRM indeed results in smaller surface dispersions, the simulations are repeated with release due east and due west, both with and without the PRM. The resulting surface dispersions are plotted in Fig. 5.19.




Figure 5.18: Magnitude and direction of the horizontal pre-release maneuver for $V_0 = 20$ cm/s.

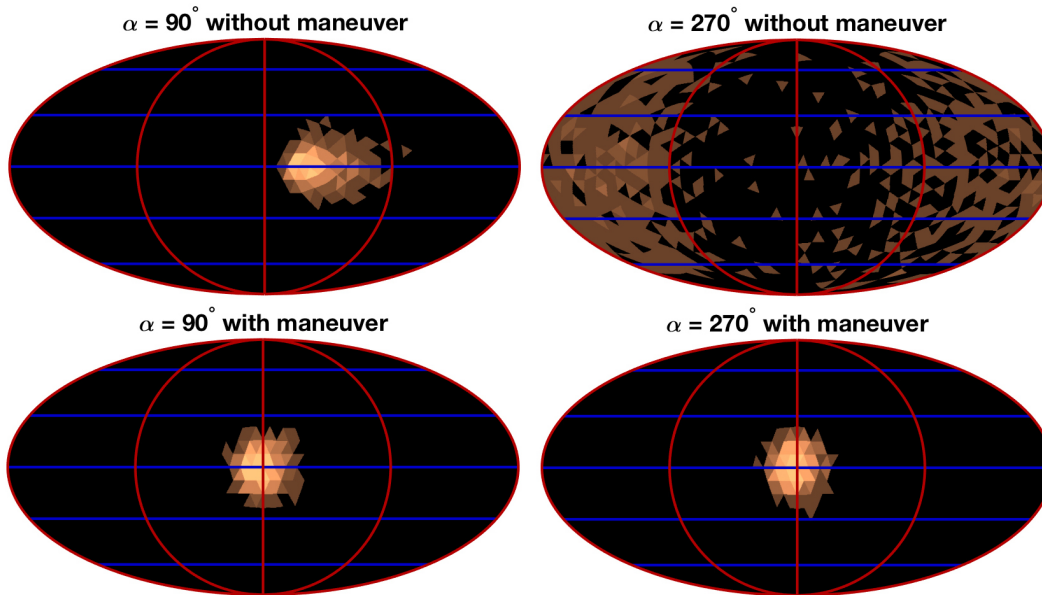


Figure 5.19: Surface dispersion of east/west release scenarios, with and without PRM.




Figure 5.20: Settling time of east/west release scenarios, with and without PRM.

This figure shows how the inclusion of a horizontal PRM renders the dispersion of a release due west acceptable, when the nominal release had an unacceptably large dispersion. Furthermore, even though the nominal release due east already performed well, Fig. 5.19 shows a slight but notable reduction in this scenario's surface dispersion as well. These results are also reflected in the corresponding settling times plotted in Fig. 5.20. The east/west releases have become equivalent in terms of both their settling time and surface dispersion statistics, while they differed greatly when the PRM was not included.

The above PRM analysis was performed for a fixed rover release velocity of $V_0 = 20$ cm/s, as holds for MINERVA-II-1A/B. However, the velocity error of the MINERVA-II-2 release mechanism is much larger, with release velocity varying uniformly as $V_0 = \mathcal{U}(5, 20)$ cm/s. In this case, the PRM should be designed to target the mean release velocity of $\bar{V}_0 = 12.5$ cm/s. The residual rover release velocity will then vary up to ± 7.5 cm/s from the \bar{V}_0 for which the PRM was designed. As problematic surface dispersions were observed when the net rover velocity was high, it must be verified whether this residual velocity still maintains acceptable levels of dispersion. Deployment is therefore simulated with a 'true' release velocity of $V_0 = [5.0, 12.5, 20.0]$ cm/s, with a PRM designed for $\bar{V}_0 = 12.5$ cm/s. This is done for releases due east, south, and west above the equator of Ryugu's reference sphere, as well as for a deployment above the northern pole. The resulting surface dispersions are shown in Fig. 5.21.

Reviewing these results, it can be seen that the surface dispersion remains appreciably small in all of the tested cases. As the PRM is designed around $\bar{V}_0 = 12.5$ cm/s, the simulations with release velocity $V_0 = 12.5$ cm/s settle quite evenly around the impact point at $\lambda = \phi = 0$ deg. The simulations with $V_0 = 5.0$ cm/s and $V_0 = 20.0$ cm/s have a net release velocity of ± 7.5 cm/s, which is reflected by a directional bias in the corresponding surface dispersions of Fig. 5.19. Nevertheless, the dispersion area remains small enough to feasibly avoid some sample site exclusion zone.




Figure 5.21: Surface dispersion for various release α and V_0 when including the nominal PRM.

5.2.4 Discussion: Deployment to the reference sphere

The simulation of rover deployment to the Ryugu reference sphere, whose total volume matches that of the training model, allows for a systematic investigation of the effects of relevant environmental and release parameters on the resulting deployment statistics. More specifically, it was revealed how the coefficient of restitution e governs both the rover settling time and surface dispersion. Although it is therefore important to use a realistic estimate of e during simulations, our results show that the surface dispersion remains limited if $e < 0.65$. The estimates of the coefficient of restitution in small-body environments available in literature predict that e will indeed remain below this limit. Given that a more accurate estimate is not possible at this time, simulations to the Ryugu training model will be performed with a uniformly varying coefficient of restitution $e = \mathcal{U}(0.1, 0.6)$, such that each rover-surface collision will be performed using a different and random e value. The resulting set of simulated trajectories will then include realistic ranges for the estimated settling time and surface dispersion. Given that the coefficient of friction f has a smaller effect on the deployment statistics, its value is fixed at $f = 0.60$, as an average value between those encountered in literature.

Most importantly, it was shown how the MINERVA-II release velocity is relatively high, resulting in high tangential impact velocities. In turn, this renders the rover surface dispersion area unacceptable large unless they are released in an easterly direction. Given that the fixed Hayabusa2 attitude limits the choice of release azimuth, an alternative strategy must be considered in order to reduce the net MINERVA-II release velocity. The inclusion of a horizontal pre-release maneuver was shown to be successful, effectively rendering all release azimuths equivalent, even when the full range of release velocity uncertainty is applied. The magnitude and direction of the maneuver will need to be adjusted to account for Ryugu's topography. The need for PRM inclusion is a critically important result that addresses dispersion concerns visible in even the simple spherical Ryugu model, and has since been included in the nominal plan for the upcoming rover deployment.

5.3 Deployment to the Ryugu training model

The simplified analysis of rover deployment to Ryugu's reference sphere revealed how the surface interaction coefficients and release geometry affect the deployment dynamics. Although these simulation results are insightful, they were obtained using a smooth sphere that lacks topographical details such as the hills, craters, and large boulders that are present on the Ryugu training model and are expected to exist on the 'real' Ryugu. This topography may strongly influence the deployment dynamics of a lander and create basins of attraction or rejection (as demonstrated in Chapter 4); the full topography must thus be taken into account. The pre-deployment analysis is therefore continued using the Ryugu training model.

For this, the previously mentioned release sequence with error model of Table 5.2 is again used, in order to initialize batches of 1,000 simulations. All simulations include a horizontal PRM that is designed using the nominal descent trajectory. Fig. 5.22 illustrates the effect of this maneuver inclusion when deploying to Ryugu's training model. In it, the Hayabusa2 spacecraft is shown (not to scale) at the lidar trigger point. The mothership coast arc is plotted as a black line, with the release event marked with a black dot and the ascent maneuver marked with a blue dot. The MINERVA-II-2 descent trajectory is drawn until first impact with a red line. The Ryugu surface

1, in Fig. 5.23. One thousands simulations are performed for each release; it is repeated that the restitution coefficient is randomly drawn from a uniform distribution as $e = \mathcal{U}(0.1, 0.6)$ for each surface collision; the friction coefficient is fixed at $f = 0.60$. Finally, note that the rover is given a random attitude at the start of each simulation. The results obtained at both epochs are now reviewed.




Figure 5.23: Selected release longitudes λ for Epoch 1 simulations, with nominal descent trajectory.

5.3.1 Epoch 1 deployment

We show overviews of all 12 Epoch 1 release scenarios in Fig. 5.24, in which the release positions are marked in red, the first-impact positions are marked in green, and the final settling positions are marked in blue. The asteroid is rotated in each subfigure such that the nominal release positions are always located at the center of the respective subfigures. Inspecting the 12 result sets, notice that the surface dispersions are all sufficiently small to fit on the single hemisphere that is visible. This is already a promising result with regards to the goal of avoiding some region on the surface. Beyond this blanket result, marked differences between the surface dispersions of the 12 scenarios are observed. This difference is perhaps the most significant when comparing the $\lambda = 30$ deg and $\lambda = 150$ deg scenarios. Detailed views of these are provided in Figs. 5.25 and 5.26 to better examine the difference. In these figures, the trajectories that have respectively the shortest, median, and longest settling time are plotted with bright, intermediate, and dark lines.




Figure 5.24: Release, first impact, and final settling positions of Epoch 1 deployment simulations to the Ryugu training model.

Fig. 5.25 shows how the $\lambda = 30$ deg first impacts occur over a relatively flat and smooth region of Ryugu's surface. As a result, the direction of the post-impact rover velocity is chiefly determined by the direction of its pre-impact velocity. Given that the MINERVA-II-2 release velocity varies as $V_0 = \mathcal{U}(5, 20)$ cm/s, and that the applied PRM targets the mean release velocity of $\bar{V}_0 = 12.5$ cm/s, there exists notable variation (roughly ± 7.5 cm/s) in the first-impact velocity of the 1,000 simulations. Because of the flatness of the first-impact region, rovers are effectively free to 'bounce' in whichever random direction results from their impact attitude, similar to the erratic motion of a die thrown on a table. The surface dispersion of the $\lambda = 30$ deg scenario is therefore relatively high.

The deployment dynamics of the $\lambda = 150$ deg release, as shown in Fig. 5.26, are quite different. This release scenario is such that rovers impact the wall of a large crater. This wall 'slopes up' towards the right of the figure; the orientation of this crater wall affects the post-impact rover velocity and directs them towards the center of the crater. Although the rovers continue to bounce around, the circular crater wall governs the geometry of individual bounces and contains the rover motion within the crater. As a result, the surface dispersion of this scenario remains relatively small, save for a handful of trajectories that manage to escape the crater, such as the longest-settling-time trajectory shown on the right of the figure.

These differences in surface dispersion are quantified in Fig. 5.27, which shows that the dispersion of the plains-deployment is approximately twice as high as that of the crater-deployment. This plot shows the marked differences in surface dispersion between the 12 release scenarios. Interestingly, these differences are much less pronounced in the settling time statistics of Fig. 5.28. Note that the settling time t_f includes the free-fall time after release and before the first impact. The mean settling time of all 12 scenarios lies between 1 and 1.5 hr, with the plains- and crater-deployments having almost identical time statistics. It is noted that the 3σ outliers mainly consist of higher-than-average t_f values. Due to the symmetric plotting of Fig. 5.28, it is therefore possible for the lower half of the 3σ region to (incorrectly) reach below $t_f = 0$. To properly examine the $3\sigma t_f$ statistics, the full distribution of values should be considered instead. Nevertheless, Fig. 5.28

provides a good visual comparison of the different scenarios.

These results indicate that the topography of the first-impact region has only a weak effect on the rate at which the rover energy is dissipated, but has a very strong effect on the surface area to which the corresponding bounces are contained. In summary, these deployment results suggest that the minimum rover dispersion (and, correspondingly, the maximum confidence in avoiding some surface exclusion zone) can be achieved by deploying rovers to a sufficiently large crater.

5.3.2 Epoch 2 deployment

We repeat the above strategy and perform 12 sets of deployment simulations with $\Delta\lambda = 30$ deg at release Epoch 2. The release, first-impact, and final settling positions of the respective cases are plotted in Fig. 5.29, where Ryugu is now viewed from the southern hemisphere. The dispersion trends seen in this picture match those of the Epoch 1 deployment: notable differences exist between the individual release scenarios, but all scenarios have a sufficiently small dispersion such that it should be possible to avoid some sampling exclusion zone. Two extreme cases are again highlighted: deployment to plains at $\lambda = 240$ deg and deployment to a crater at $\lambda = 90$ deg, in respectively Figs. 5.30 and 5.31.

These two highlighted scenarios display behavior that matches Epoch 1. When releasing rovers above a region that is mostly flat, rovers are free to bounce around and have a relatively large dispersion. When deploying above a large crater, rovers are contained by the crater and have a relatively small dispersion. In this case, the crater deployment does have one disadvantage: the large crater contains two smaller, secondary craters, as marked with white circles in Fig. 5.31. These secondary craters have a further focusing effect that ‘attracts’ a notable number of simulations. Given the small size of these craters, their walls may obstruct the line-of-sight between rover and mothership and severely limit the available communication window. This reduces the advantage of this release scenario.




Figure 5.25: Detailed view of the $\lambda = 30$ deg Epoch 1 deployment to plains.




Figure 5.26: Detailed view of the $\lambda = 150$ deg Epoch 1 deployment to a crater.




Figure 5.27: Surface dispersion statistics of Epoch 1 deployment to the Ryugu training model.




Figure 5.28: Settling time statistics of Epoch 1 deployment to the Ryugu training model.

5.3.3 Discussion: Deployment to the training model

The simulation results for deployment to the Ryugu training model predict that the rover dispersion area remains appreciably small when the horizontal pre-release maneuver is included. Notable effects of the release longitude λ do exist, with the surface dispersion varying by up to 100% between the different scenarios. Overall, the 3σ dispersion area varies between roughly 3% and 10% of the Ryugu surface. This is sufficiently small such that the exclusion zone around some targeted Hayabusa2 sampling site can be avoided. The dispersion area can be further reduced by choosing to deploy MINERVA-II-2 to a crater, assuming craters indeed exist on the real Ryugu.

The small dispersion of crater deployment scenarios has several advantages. The prediction that the rover trajectory will be contained by the targeted crater should improve the odds of the mothership successfully tracking the rover throughout its descent. If the rover position can be tracked both before and after an impact, one may estimate the effective surface interaction coefficients that act during the impact. This provides highly valuable scientific information about the asteroid surface. Furthermore, if the surface interaction coefficients can be estimated following the Epoch 1 deployment, the Epoch 2 (and MASCOT) deployment prediction can be performed with greater confidence. This may enable more precise aiming of rovers during successive deployments, such that they can better target regions of interest. In order to do this, it will be important to estimate the true MINERVA-II-2 release velocity, given the large release mechanism uncertain-




Figure 5.29: Release, first impact, and final settling positions of Epoch 2 deployment simulations to the Ryugu training model.




Figure 5.30: Detailed view of the $\lambda = 240$ deg Epoch 2 deployment to plains.




Figure 5.31: Detailed view of the $\lambda = 90$ deg Epoch 2 deployment to a crater.




Figure 5.32: Settling time statistics of Epoch 2 deployment to the Ryugu training model.




Figure 5.33: Surface dispersion statistics of Epoch 2 deployment to the Ryugu training model.

ties. This velocity can be estimated from observed changes in the Hayabusa2 velocity at release. Although the mission geometry renders these changes poorly observable in linear velocity measurements (through, for example, Doppler effects), a strong signal should be visible in the Hayabusa2 attitude motion. By combining the release velocity estimate with descent tracking and on-board measurements from the rovers, an accurate trajectory reconstruction should be possible, which has the potential to yield valuable scientific information about Ryugu’s surface.

5.4 Effect of the asteroid mass

Although the size of asteroid Ryugu has been established to a reasonable precision, no information on its density is available. As mentioned in Section 5.1, the training model is assumed to have a nominal gravitational parameter of $\mu = 32 \text{ m}^3/\text{s}^2$, corresponding to a mean density of $\rho = 1.5 \text{ g/cm}^3$ when using the volume of the training shape model. Given the density uncertainty, the mission rehearsal performed by JAXA takes into account a uniform error model for $\mu = \mathcal{U}(11, 92) \text{ m}^3/\text{s}^2$. These correspond to densities of $\rho = \mathcal{U}(0.4, 4.3) \text{ g/cm}^3$.

We note that Ryugu is a C-type (carbonaceous) asteroid; such bodies typically have a density of 1.3 to 1.4 g/cm^3 [53]. For context, these values are compared with relevant density measurements available in literature. The rubble-pile asteroid Itokawa was measured to have a density of $\rho =$

$1.90 \pm 0.13 \text{ g/cm}^3$ during the Hayabusa mission [34]. Asteroid Bennu, the target of the OSIRIS-REx mission, is estimated to have a density of $\rho = 1.26 \pm 0.07 \text{ g/cm}^3$ [20]. This estimate was obtained using a large set of orbit and radar data, from which models of the YORP effect were able to establish a value for ρ . The NEAR-Shoemaker spacecraft measured the density of asteroid Eros at $\rho = 2.67 \pm 0.03 \text{ g/cm}^3$ [146]. The density values of Itokawa, Bennu, and Eros are perhaps the most relevant, as they are of a similar size and (likely) composition as Ryugu. In addition, comet 67P/C-G was found to have a mean density of $\rho \approx 0.53 \text{ g/cm}^3$ [88]. This density is notably lower than the aforementioned asteroids, due to the comet's composition of low-density ices. Finally, note the densities of the Martian moons Phobos and Deimos, respectively $\rho = 1.88 \text{ g/cm}^3$ and $\rho = 1.47 \pm 0.17 \text{ g/cm}^3$, though these bodies are roughly an order of magnitude larger than Ryugu, and exist in a very different dynamical environment. These density values, together with the given range of the Ryugu density, are plotted in Fig. 5.34 for visual reference.




Figure 5.34: Comparison of Ryugu training model ρ with relevant small bodies [34, 20, 146, 88].

Comparing the density range of the Ryugu training model with the density values of the relevant bodies included in Fig. 5.34, note that the ρ_{min} value is typical of a cometary body. Given the spectral observation that classifies Ryugu as a C-type asteroid, its density is therefore unlikely to be this low. Similarly, the ρ_{max} value is particularly high and more typical of M-type (metallic) asteroids [53]. It has also been found in fast-spinning asteroids of various types, though in bodies of much smaller size ($< 100 \text{ m}$) than Ryugu [89]. Despite both extremes being unlikely for a C-type asteroid of Ryugu's size, it is still valuable to investigate how variations in the asteroid

density would affect rover deployment. The simulation of rover deployment to the Itokawa training model is thus repeated, using the extreme gravitational parameter values of $\mu_{min} = 11 \text{ m}^3/\text{s}^2$ and $\mu_{max} = 92 \text{ m}^3/\text{s}^2$. In the interest of conciseness, the discussion is restricted to an analysis of the best- and worst-case deployments that were discussed in the previous section, for both release epochs. Figs. 5.35 and 5.36 compare the release, first-impact, and final settling positions of the best- and worst-case deployments, for the three relevant values of μ . These are respectively the crater and plains deployment scenarios.




Figure 5.35: Worst- and best-case Epoch 1 training model deployments, with varying μ .

In all cases, it is found that the rover surface dispersion increases notably for μ_{min} and decreases slightly for μ_{max} . This trend is expected: when μ is smaller, Ryugu's gravitational pull is weaker, such that it takes rovers more time to return to the surface after an impact. With a fixed tangential velocity, rovers thus cover more distance across the asteroid surface when μ is lower. These differences in surface dispersion appear mostly in the scenarios that deploy to Ryugu's




Figure 5.36: Worst- and best-case Epoch 2 training model deployments, with varying μ .

plains. When deploying to craters, the differences are much smaller. These visually observed trends are also quantified in Fig. 5.37, which agrees with our discussions so far: the differences between the $\mu = 11 \text{ m}^3/\text{s}^2$ and $\mu = 32 \text{ m}^3/\text{s}^2$ scenarios are far more dramatic than those between the $\mu = 32 \text{ m}^3/\text{s}^2$ and $\mu = 92 \text{ m}^3/\text{s}^2$ scenarios. The settling time variations of Fig. 5.38 also show corresponding variations.

These results suggest that rover deployment becomes more challenging when the asteroid density is lower. Since deployment to a low-density asteroid has larger associated rover dispersions, it becomes more difficult to reliably guarantee the avoidance of some exclusion zone on the surface. This is also expected to render the rover descent tracking more challenging. Although deployment to craters still results in the lowest surface dispersion even at μ_{min} , the craters' effectiveness at containing rover trajectories notably decreases. As a result, many simulations do exit the craters and settle in the surrounding plains. In summary, this indicates that the selection of the rover




Figure 5.37: Surface dispersion statistics of deployment with varying μ .




Figure 5.38: Settling time statistics of deployment with varying μ .

release point requires greater care if the asteroid density ρ is low. In these cases, the inclusion of the horizontal pre-release maneuver is paramount in order to sufficiently restrict the surface dispersion area. At nominal or high asteroid densities, rover trajectories are naturally more spatially contained, resulting in more robustness to position and velocity errors at release. In conclusion, the combination of Ryugu's shape and mean density set the playing field that determines precisely how challenging the selection of the MINERVA-II release location is.

5.5 Summary: The deployment of MINERVA-II

The Japanese Hayabusa2 asteroid sample return mission will arrive at its target, asteroid Ryugu, at some time between June and July of 2018. Following several months of remote sensing operations, the spacecraft will deploy three MINERVA-II rovers to Ryugu's surface and obtain sev-

eral surface samples during touch-and-go maneuvers. In these maneuvers, Hayabusa2 autonomously descends to the asteroid surface using its lidar, as well as visual tracking of target markers. To avoid interfering with the Hayabusa2 descent guidance system, the MINERVA-II rovers must be deployed in a way that minimizes the rovers' chances of entering some exclusion zone around the targeted sampling sites. Given the complex dynamics of rovers in small-body environments, their release must be carefully designed using high-fidelity simulation. Although the shape and spin models of Ryugu will remain uncertain until arrival, a pre-arrival deployment analysis can be performed using the Ryugu training model, which was developed for broad mission rehearsal at the Japan Aerospace Exploration Agency (JAXA). The pre-arrival analysis, which was performed using the methodology outlined in this thesis, has helped plan for the actual deployment analysis to be carried out post-arrival by revealing challenges and opportunities in the rover deployment.

Using first the Ryugu reference sphere with equivalent volume as the training model, the effects of the surface interaction coefficients are investigated. This revealed that the coefficient of restitution e is the governing determinant of the rover deployment statistics, with the friction coefficient f playing a minor role. When the restitution coefficient $e < 0.65$, as can be expected from literature, the statistics scale roughly linearly with changes in e . Continuing to use the Ryugu reference sphere, it was shown that the MINERVA-II release velocity and associated uncertainties are relatively high, resulting in high tangential impact velocities and unacceptably large corresponding surface dispersions. In order to mitigate this issue, we propose the inclusion of a horizontal pre-release maneuver (PRM) that allows the Hayabusa2 mothership to cancel some of the rover's net velocity and ensure a first impact that occurs normal to the asteroid surface. The PRM is shown to indeed significantly reduce the rover surface dispersion and settling time and has since been included in the nominal operation plan for the rover deployment.

Using this insight into the general motion trends of MINERVA-II, simulations of deployment to the full Ryugu training model are performed at two candidate release epochs. Although the release latitude is fixed at both epochs due to the Ryugu-Earth-Sun geometry, the release longitude can be freely chosen using the particular time of deployment. Release is therefore simulated from

a range of longitudes, while the coefficient of restitution is randomly varied at each impact, with a fixed friction coefficient. With the included PRM, only minor variations in the settling time t_f are found between the various release λ , with a mean first-impact time of $t_1 \approx 11$ min and a mean settling time of $t_f \approx 75$ min. Far greater variation is observed in the surface dispersion area, which depends strongly on the topography of the first-impact region. Larger dispersions are observed when deploying to a flat region on Ryugu's surface and smaller dispersions when deploying to a crater. Sufficiently large craters are able to contain the rover descent trajectories even in the presence of uncertainties in the mothership navigation, rover release, and surface interaction properties. This suggests that rovers should be deployed to large craters on Ryugu's surface in order to minimize the surface dispersion and maximize the confidence of avoiding a sample site exclusion zone. Furthermore, a smaller surface dispersion increases the chances of successfully tracking a rover throughout its descent to and impacts on the surface of Ryugu. This increases the data available for estimation of the e and f coefficients, which provide scientific information on Ryugu's structure and composition. It is noted that the Hayabusa2 sampling is also likely to target a surface crater, given that exposed crater materials tend to be less affected by space weathering and contain greater amounts of scientifically-interesting hydrous minerals [144]. If multiple large craters exist on Ryugu, it is therefore likely for the Hayabusa2 sampling to target one such crater, while rovers are deployed to another crater.

This completed pre-arrival analysis has outlined the steps to be taken by the 'real' deployment analysis that must be performed with the accurate Ryugu models that will be available following the arrival of Hayabusa2. The results documented in this case study nonetheless establish trends that are expected to persist on the real body. Our work has revealed the need for PRM inclusion in order to limit the rover surface dispersion and stresses the need for an accurate estimate of the true rover release velocity in order to accurately reconstruct the rover descent trajectory. These results have adjusted the nominal Hayabusa2 operational plans.

Chapter 6

Conclusions

This thesis has presented a comprehensive methodology to simulate the motion of ballistic spacecraft probes, operating in the vicinity and on the surface of a small body such as an asteroid, comet, or small moon. The work is distinguished from that of previous authors primarily through its ability to handle arbitrary probe shapes. To the best of our knowledge, all previous efforts to simulate *global* motion across a small body are limited to either pointmass or spherical lander models, as detailed in Chapter 1. Such models are unable to reproduce the complex and chaotic bouncing behavior of non-spherical shapes, resulting in biased deployment dispersions. These models have been implemented in a GPU environment to allow for parallel evaluation of multiple simulations, which greatly increases the speed at which simulation can be performed, and enables broad studies of probe deployment. The resulting simulation framework has been applied to perform extensive parametric studies that reveal the effect of relevant environmental parameters, as well as to the real-world scenario of rover deployment on the Hayabusa2 asteroid exploration mission.

6.1 Simulation of small-body probes

The applied modeling methods were discussed in detail in Chapter 2. After a brief discussion of the general probe equations of motion, different shape modeling methods were presented. Most importantly, the use of the *implicit* signed distance field (SDF) was proposed as a replacement for the classically-used *explicit* polyhedron model. The SDF can be constructed prior to simulation to allow for fast online interpolation, which significantly increases the speed at which collision

detection between a probe and the target body can be performed. These global shape models are limited with regards to the smallest resolved surface features. Using procedural generation techniques, rock and boulders can be efficiently generated on the small-body surface. Both the generation of polyhedral rocks on the polyhedron shape model are demonstrated, as well as fractional Brownian motion and fully procedural rocks on the signed distance field. These can be tuned to match the rock distributions observed on actual small-body surfaces. Although we continue to use the classical (constant-density) polyhedron model for gravity field modeling, this field is also pre-computed; online interpolations of the perturbation with respect to a central gravity field are performed in order to minimize the required memory storage. The modeling of probe-target contact interactions distinguishes between impulsive collisions and continued contact motion. A full derivation of the contact model is provided, and includes (energetic) restitution, Coulomb friction, and rolling resistance forces and torques, which are governed by their respective coefficients. The model captures the interaction between a probe and a hard surface layer and is optimized for easy resolution of stick/slip behavior and implementation in a parallelized framework. Complex interactions with a regolith layer on top of the hard surface layer are abstracted through a modulation of the coefficient of restitution, in agreement with (limited) experiments available in literature. An efficient GPU implementation is achieved by distinguishing between flying, contact, and collision phases of motion, and executing these across separate GPU kernels. This allows for approximately 20 typical simulations to be performed each second, on a 2018 high-performance machine.

The *lift-off velocity* is the velocity that an object needs to cease contact with the small-body surface when moving (tangentially) across its surface, and is the focus of Chapter 3. A general derivation of the velocity is provided, and simplified the expression for the practical cases of lift-off from a curved surface, a sharp ridge, and a plane. The surface radius of curvature is found to dominate the expressions; methods to obtain the radius of curvature of an ellipsoid, spherical harmonic, and polyhedron shape are provided. This is then first applied to some simple shapes, sphere and ellipsoid, to understand some general properties of the lift-off velocity. Finally, the expressions are applied to generate lift-off velocity distributions across several solar system bodies.

The simulation framework was then applied to perform parametric studies of probe motion in the small-body environment in Chapter 4. First, the motion of spherical probes was investigated, for which the effects of the respective surface interaction coefficients were established. The coefficient of restitution e was found to govern the deployment statistics; the coefficient of friction f did not affect the motion unless it was very small. Studying the probe mass distribution, it was found that more concentrated mass distributions result in shorter settling times and smaller surface dispersions. Finally, the presence of polyhedral rocks was found to significantly affect the deployment statistics, with the inclusion of rocks reducing both the settling time and surface dispersion. A saturation limit appears when the spacing between rocks becomes smaller than the probe; the inclusion of additional rocks is found to no longer affect the deployment statistics.

Second, the motion of probes with more complex shapes was investigated. Varying the probe shape, those with a more spherical shape (such as the dodecahedron) were found to settle more slowly and with greater surface dispersion than less spherical shapes (such as the tetrahedron). These trends are explained through inspection of the aspect angle, which determines the feasible geometries under which probe-target impacts can occur. Varying the surface interaction coefficients, the restitution coefficient was found to govern the probe motion, much like it did for spherical probes. However, the effect of the friction coefficient on complex probes is notably different from the effect on spherical probes. It was found that increased f results in increased settling time and surface dispersion, up to some saturation value f^* beyond which further increases cease to have an effect. The value of this f^* is dependent on the probe shape: it is very small for spheres, but reaches values beyond 1 for very non-spherical shapes such as the tetrahedron. Finally, the effect of mass distribution on the motion of these complex probes are investigated. In general, it was found that variants of a given shape with a more outwards mass distribution settle more slowly than a variant of that same shape with a more concentrated mass distribution. This effect is almost negligible for nearly-spherical shapes, but becomes significant for less spherical shapes such as a cube or tetrahedron.

Third, the effects of small-body surface topography on the probe motion are examined. Simulation to surfaces with procedurally-generated rocks indicate that the presence of rocks reduces the settling time and surface dispersion of deployed probes. This effect is more significant for spherical probes than it is for more complex shapes. When including rocks of various sizes, only those larger than the considered probes are found to affect deployment dynamics. The presence of regolith, which experiments have shown to damp normal impacts more strongly than tangential impacts, is accounted for using a restitution-modulation model. This indicates that near-spherical probes experience effectively lower restitution on a regolith-covered surface than more distorted shapes. In conclusion, the presence of fine surface topography reduces the settling time and surface dispersion of near-spherical probes; they will act more like distorted, non-spherical shapes.

Finally, the simulation framework is applied to investigate rover deployment to asteroid Ryugu (formerly 1999 JU₃), in Chapter 5. More specifically, a pre-arrival deployment analysis was performed for the MINERVA-II rovers onboard the Hayabusa2 asteroid sample return mission. The rovers will be deployed a few months after arrival, and although current knowledge of the Ryugu system is limited, it is important to prepare for the rover deployment through a pre-arrival study. Simulation of deployment to the Ryugu reference sphere revealed a major challenge, namely, the relatively large magnitude of the MINERVA-II ejection velocity of up to 20 cm/s. This high velocity results in first impacts with high tangential velocity that result in unacceptably large surface dispersions. To alleviate this issue, our work has suggested the inclusion of a horizontal pre-release maneuver with the Hayabusa2 spacecraft, in order to reduce the effective rover velocity at release. This was shown to reduce the rover dispersion to an acceptable size. Including this maneuver, simulations to the detailed Ryugu training shape model were performed. This model was developed by JAXA for broad mission rehearsal purposes. The effects of rover deployment to different terrains on the asteroid were investigated; this showed that deployment to large craters minimizes the rover surface dispersion. We therefore suggest deployment to such a crater in order to improve the chances of successfully tracking the rover through its descent trajectory. The impact dynamics that can thus be estimated may provide information on Ryugu's (sub-)surface structure.

6.2 Future work

The simulation framework presents, to the best of our knowledge, the first attempt at capturing the complex dynamics between a small-body exploration probe with arbitrary shape and a complex target surface. The applied parallel implementation of this framework enables fast simulation of a large number of trajectories. Nevertheless, the framework can be extended in several directions. Here, some areas in which this work could be extended are discussed. Some preliminary guidelines for doing so are also provided, both with respect to modeling and with respect to applications.

- (1) *Post-arrival MINERVA-II deployment analysis and trajectory reconstruction.* The simplest extension of this thesis work is the continued application of the simulation framework to the deployment planning of the rovers onboard Hayabusa2. Following the spacecraft's arrival at Ryugu in June/July of 2018, its observations will be used to produce a high-resolution shape model of the asteroid, similar to the training shape model presented in Chapter 5. Hayabusa2 will also estimate the asteroid's mass and rotation pole. These updated models must be used to repeat the deployment analysis of Chapter 5. Although the pre-arrival analysis was performed independently, it is paramount that the post-arrival analysis be performed in conjunction with the Hayabusa2 and MINERVA-II science teams. This will provide the necessary constraints and requirements needed for the deployment analysis to select the optimal rover release site. Furthermore, once this deployment is performed near the end of September 2018, the bouncing MINERVA-II trajectory can be reconstructed using observations from the Hayabusa2 spacecraft. This reconstructed trajectory can be compared with simulations in order to estimate the effective restitution and friction coefficients that acted during the rover's impacts on the asteroid surface. This will provide scientific information on the (sub-)surface structure and will inform subsequent deployments to Ryugu and similar bodies.

- (2) *Advanced impulsive contact modeling and experiments.* Although this thesis has presented a successful effort at modeling the contact between a complex probe shape and small-body surface, the applied model remains somewhat limited in scope. More specifically, our impulsive model abstracts away much of the complex impact mechanics into effective coefficients of restitution and friction. Lacking experimental results on the likely effects of probe attitude and velocity at impact on the coefficient values, we believe that this is presently the best possible effort. However, there remains great utility in performing a broad study combining the analytics and experiments of rigid probes impacting onto a hard surface layer in microgravity. Some of these suggested experiments were detailed in Chapter 4. It may be possible to empirically fit some dependency of e and f on the impact conditions, though more advanced impact models employing elastoplastic theory will likely yield better and more substantiated results. Although this would be a significant project, it constitutes an area of basic research which, to the best of our knowledge, is currently lacking in the field of impact mechanics.
- (3) *Advanced regolith contact modeling.* Similar to the suggestion made above, it holds that the applied regolith interaction model significantly abstracts the complex interactions that are expected to occur on the small-body surface. Although our implementation can recreate the limited experimental work of low-velocity impacts into granular materials, it continues to be a highly simplified model. The extension of this model will require both extensive modeling using DEM-type simulation, as well as experimental work. An extension of the experiments suggested in Chapter 4 with granular materials would provide an excellent starting point, given the high degree of control that can be achieved in such experiments. This combination of experiment with detailed but *local* simulations could be used to empirically fit some effective collision parameters to the observed behavior. These parameters can then be incorporated into the *global* simulations detailed in this thesis in order to produce realistic

behavior without requiring computationally intensive DEM models in those simulations.

- (4) *Simulation of scientific probe instruments.* Although ballistically-deployed probes are now included on small-body exploration missions, they are still considered as secondary payloads. Although the European MASCOT lander has an extensive suite of scientific instruments, the Japanese MINERVA-II rovers remain quite simple technology demonstrators. As a result, although they are equipped with a few instruments, they do not carry instrumentation that can be used to extract scientific information directly from bouncing trajectories. The inclusion of accelerometers, for example, would greatly improve the accuracy of reconstructed trajectories and thus of the information returned about the small-body (sub-)surface. The implementation of instrument models into our simulation framework would allow for the generation of realistic measurement data. This would enable a broad investigation of how effectively surface properties can be extracted from the measurement data. It will also provide insight into how the internal placement of these instruments would affect their science return. Finally, it would allow for the development of some conceptual probe designs aimed at maximizing scientific return from bouncing trajectories.

Bibliography

- [1] Michael F A'Hearn, MJS Belton, WA Delamere, J Kissel, KP Klaasen, LA McFadden, KJ Meech, HJ Melosh, PH Schultz, JM Sunshine, et al. Deep impact: excavating comet tempel 1. science, 310(5746):258–264, 2005.
- [2] Michael F A'Hearn, MJS Belton, WA Delamere, J Kissel, KP Klaasen, LA McFadden, KJ Meech, HJ Melosh, PH Schultz, JM Sunshine, et al. Deep impact: excavating comet tempel 1. science, 310(5746):258–264, 2005.
- [3] M Ahmad, KA Ismail, and F Mat. Impact models and coefficient of restitution: A review. Journal of Engineering and Applied Sciences, 11:6549–6555, 2006.
- [4] Ergun Akleman and Jianer Chen. Generalized distance functions. In Shape Modeling and Applications, 1999. Proceedings. Shape Modeling International'99. International Conference on, pages 72–79. IEEE, 1999.
- [5] M. A. Ambroso, C. R. Santore, A. R. Abate, and D. J. Durian. Penetration depth for shallow impact cratering. Phys. Rev. E, 71:051305, May 2005.
- [6] Robert C Anderson, Daniel Scheeres, Steven Chesley, and BASiX Team. Binary asteroid in-situ explorer mission (basix): A mission concept to explore a binary near earth asteroid system. In Lunar and Planetary Science Conference, volume 45, page 1571, 2014.
- [7] Robert C Anderson, Daniel Scheeres, Steven Chesley, and BASiX Team. Binary asteroid in-situ explorer mission (basix): A mission concept to explore a binary near earth asteroid system. In Lunar and Planetary Science Conference, volume 45, page 1571, 2014.
- [8] Christopher Batty. Sdfgen, 2015.
- [9] MJS Belton, J Veverka, P Thomas, P Helfenstein, D Simonelli, C Chapman, ME Davies, R Greeley, R Greenberg, J Head, et al. Galileo encounter with 951 gaspra: First pictures of an asteroid. Science, 257(5077):1647–1652, 1992.
- [10] Vivek Bhatt and Jeff Koechling. Three-dimensional frictional rigid-body impact. TRANSACTIONS-AMERICAN SOCIETY OF MECHANICAL ENGINEERS JOURNAL OF APPLIED MECHANICS, 62:893–898, 1995.
- [11] J-P Bibring, H Rosenbauer, H Boehnhardt, S Ulamec, J Biele, S Espinasse, B Feuerbacher, P Gaudon, P Hemmerich, P Kletzkine, et al. The rosetta lander (philae) investigations. Space science reviews, 128(1-4):205–220, 2007.

- [12] Jens Biele, Lars Kessler, Christian D Grimm, Silvio Schröder, Olaf Mierheim, Michael Lange, and Tra-Mi Ho. Experimental determination of the structural coefficient of restitution of a bouncing asteroid lander. arXiv preprint arXiv:1705.00701, 2017.
- [13] Jens Biele, Stephan Ulamec, Michael Maibaum, Reinhard Roll, Lars Witte, Eric Jurado, Pablo Muñoz, Walter Arnold, Hans-Ulrich Auster, Carlos Casas, et al. The landing (s) of philae and inferences about comet surface mechanical properties. Science, 349(6247):aaa9816, 2015.
- [14] Paul Bourke. Interpolation methods. 2012.
- [15] Don Brownlee. The stardust mission: analyzing samples from the edge of the solar system. Annual Review of Earth and Planetary Sciences, 42:179–205, 2014.
- [16] Nazli Can. Space mining. In Recent Advances in Space Technologies (RAST), 2017 8th International Conference on, pages 495–499. IEEE, 2017.
- [17] Elisabet Canalias, Alejandro Blázquez, Eric Jurado, and Thierry Martin. Philae descent trajectory computation and landing site selection on comet churyumov-gerasimenko. In Proceedings of 23rd International Symposium on Spaceflight Dynamics, Pasadena, California, 2012.
- [18] F Capaccioni, A Coradini, G Filacchione, S Erard, Gabriele Arnold, P Drossart, MC De Sanctis, D Bockelee-Morvan, MT Capria, F Tosi, et al. The organic-rich surface of comet 67p/churyumov-gerasimenko as seen by virtis/rosetta. Science, 347(6220):aaa0628, 2015.
- [19] Andy F Cheng, J Atchison, Brian Kantsiper, Andrew S Rivkin, A Stickle, Cheryl Reed, Andres Galvez, Ian Carnelli, Patrick Michel, and S Ulamec. Asteroid impact and deflection assessment mission. Acta Astronautica, 115:262–269, 2015.
- [20] Steven R Chesley, Davide Farnocchia, Michael C Nolan, David Vokrouhlický, Paul W Chodas, Andrea Milani, Federica Spoto, Benjamin Rozitis, Lance AM Benner, William F Bottke, et al. Orbit and bulk density of the osiris-rex target asteroid (101955) bennu. Icarus, 235:5–22, 2014.
- [21] Christopher F Chyba, Paul J Thomas, and Kevin J Zahnle. The 1908 tunguska explosion: atmospheric disruption of a stony asteroid. Nature, 361(6407):40, 1993.
- [22] Paolo Cignoni, Marco Callieri, Massimiliano Corsini, Matteo Dellepiane, Fabio Ganovelli, and Guido Ranzuglia. MeshLab: an Open-Source Mesh Processing Tool. In Vittorio Scarano, Rosario De Chiara, and Ugo Erra, editors, Eurographics Italian Chapter Conference. The Eurographics Association, 2008.
- [23] J Colwell, J Brisset, A Dove, A Whizin, H Nagler, and N Brown. Prime: Studying low-velocity impacts in microgravity. In European Planetary Science Congress 2015, held 27 September-2 October, 2015 in Nantes, France, Online at <http://meetingorganizer.copernicus.org/EPSC2015>, id. EPSC2015-766, volume 10, 2015.
- [24] Joshua E Colwell. Low velocity impacts into dust: Results from the collide-2 microgravity experiment. Icarus, 164(1):188–196, 2003.
- [25] Rod Cross. Measurements of the horizontal coefficient of restitution for a superball and a tennis ball. American Journal of Physics, 70(5):482–489, 2002.

- [26] Fabio Vittorio De Blasio and May-Britt Saeter. Rolling friction on a granular medium. Phys. Rev. E, 79:022301, Feb 2009.
- [27] Fabio Vittorio De Blasio and May-Britt Saeter. Rolling friction on a granular medium. Physical Review E, 79(2):022301, 2009.
- [28] Mario Di Martino, A Carbognani, A Cellino, G De Sanctis, and V Zappalà. The asteroid hazard-evaluating and avoiding the threat of asteroid impacts. In ESA Special Publication, volume 1310, 2009.
- [29] Anthony R Dobrovolskis. Inertia of any polyhedron. Icarus, 124(2):698–704, 1996.
- [30] David S Ebert. Texturing & modeling: a procedural approach. Morgan Kaufmann, 2003.
- [31] Yuval Fisher, Michael McGuire, Richard F Voss, Michael F Barnsley, Robert L Devaney, and Benoit B Mandelbrot. The science of fractal images. Springer Science & Business Media, 2012.
- [32] Samuel F Foerster, Michel Y Louge, Hongder Chang, and Khedidja Allia. Measurements of the collision properties of small spheres. Physics of Fluids, 6(3):1108–1115, 1994.
- [33] Akira Fujiwara, J Kawaguchi, DK Yeomans, M Abe, T Mukai, T Okada, J Saito, H Yano, M Yoshikawa, DJ Scheeres, et al. The rubble-pile asteroid itokawa as observed by hayabusa. Science, 312(5778):1330–1334, 2006.
- [34] Akira Fujiwara, J Kawaguchi, DK Yeomans, M Abe, T Mukai, T Okada, J Saito, H Yano, M Yoshikawa, DJ Scheeres, et al. The rubble-pile asteroid itokawa as observed by hayabusa. Science, 312(5778):1330–1334, 2006.
- [35] Akira Fujiwara, J Kawaguchi, DK Yeomans, M Abe, T Mukai, T Okada, J Saito, H Yano, M Yoshikawa, DJ Scheeres, et al. The rubble-pile asteroid itokawa as observed by hayabusa. Science, 312(5778):1330–1334, 2006.
- [36] R. Gaskell, J. Saito, M. Ishiguro, T. Kubota, T. Hashimoto, N. Hirata, S. Abe, O. Barnouin-Jha, and D. Scheeres. Gaskell itokawa shape model v1.0. hay-a-amica-5-itokawashape-v1.0., 2008.
- [37] Daniel I. Goldman and Paul Umbanhowar. Scaling and dynamics of sphere and disk impact into granular media. Phys. Rev. E, 77:021308, Feb 2008.
- [38] Michael Grewing, Francoise Praderie, and Rüdiger Reinhard. Exploration of Halley's comet. Springer Science & Business Media, 2012.
- [39] V Guibout and Daniel J Scheeres. Stability of surface motion on a rotating ellipsoid. Celestial Mechanics and Dynamical Astronomy, 87(3):263–290, 2003.
- [40] Sunao Hasegawa, Thomas G Müller, Kyoko Kawakami, Toshihiro Kasuga, Takehiko Wada, Yoshifusa Ita, Naruhisa Takato, Hiroshi Terada, Takuya Fujiyoshi, and Masanao Abe. Albedo, size, and surface characteristics of hayabusa-2 sample-return target 162173 1999 ju3 from akari and subaru observations. Publications of the Astronomical Society of Japan, 60(sp2):S399–S405, 2008.

- [41] A Herique, J Biele, P Bousquet, V Ciarletti, TM Ho, JL Issler, W Kofman, P Michel, D Plettemeier, P Puget, et al. Fantina: Fathom asteroids now: Tomography and imagery of a near-payload for marco polo r cv3/esa mission. In EGU General Assembly Conference Abstracts, volume 14, page 9633, 2012.
- [42] HJ Herrmann and Stefan Luding. Modeling granular media on the computer. Continuum Mechanics and Thermodynamics, 10(4):189–231, 1998.
- [43] Tra-Mi Ho, Volodymyr Baturkin, Christian Grimm, Jan Thimo Grundmann, Catherin Hobbie, Eugen Ksenik, Caroline Lange, Kaname Sasaki, Markus Schlotterer, Maria Talapina, et al. Mascotthe mobile asteroid surface scout onboard the hayabusa2 mission. Space Science Reviews, 208(1-4):339–374, 2017.
- [44] B Hockman and M Pavone. Stochastic motion planning for hopping rovers on small solar system bodies. In Proceedings of ISRR, 2017.
- [45] Benjamin J Hockman, Andreas Frick, Robert G Reid, Issa AD Nesnas, and Marco Pavone. Design, control, and experimentation of internally-actuated rovers for the exploration of low-gravity planetary bodies. Journal of Field Robotics, 34(1):5–24, 2017.
- [46] R Scott Hudson and Steven J Ostro. Shape and non-principal axis spin state of asteroid 4179 toutatis. Science, 270(5233):84, 1995.
- [47] RS Hudson, Steven J Ostro, and DJ Scheeres. High-resolution model of asteroid 4179 toutatis. Icarus, 161(2):346–355, 2003.
- [48] Íñigo Quílez. Improved texture interpolation, 2009. Last accessed: 29 March 2018.
- [49] Íñigo Quílez. Texture repetition, 2015. Last accessed: 26 April 2018.
- [50] Seth A Jacobson and Daniel J Scheeres. Dynamics of rotationally fissioned asteroids: Source of observed small asteroid systems. Icarus, 214(1):161–178, 2011.
- [51] Mark W Jones, J Andreas Baerentzen, and Milos Sramek. 3d distance fields: A survey of techniques and applications. IEEE Transactions on visualization and Computer Graphics, 12(4):581–599, 2006.
- [52] Horst Uwe Keller, WA Delamere, WF Huebner, HJ Reitsema, HU Schmidt, FL Whipple, K Wilhelm, W Curdt, R Kramm, N Thomas, et al. Comet p/halleys nucleus and its activity. In Exploration of Halleys Comet, pages 807–823. Springer, 1988.
- [53] George A Krasinsky, Elena V Pitjeva, Mikhail V Vasilyev, and EI Yagudina. Hidden mass in the asteroid belt. Icarus, 158(1):98–105, 2002.
- [54] Takashi Kubota and Tetsuo Yoshimitsu. Intelligent rover with hopping mechanism for asteroid exploration. In Recent Advances in Space Technologies (RAST), 2013 6th International Conference on, pages 979–984. IEEE, 2013.
- [55] Caroline Lange, Claudia Dietze, Tra-Mi Ho, Olaf Kroemer, Michael Lange, and Susanne Wagenbach. Baseline design of a mobile asteroid surface scout (mascot) for the hayabusa-2 mission. In Intelligence for Space Robotics, San Antonio, Texas, USA, pages 115–135, 2006.

- [56] Dante S Lauretta and OSIRIS-Rex Team. An overview of the osiris-rex asteroid sample return mission. In Lunar and Planetary Science Conference, volume 43, 2012.
- [57] Dante S Lauretta and OSIRIS-Rex Team. An overview of the osiris-rex asteroid sample return mission. In Lunar and Planetary Science Conference, volume 43, page 2491, 2012.
- [58] HF Levison et al. Lucy: surveying the diversity of the trojan asteroids, the fossils of planet formation. In Lunar and Planetary Science Conference, volume 47, page 2061, 2016.
- [59] John S Lewis. Mining the sky: untold riches from the asteroids, comets, and planets. Reading, Mass.: Addison-Wesley Pub. Co., c1996., 1, 1996.
- [60] Yu I Lobanovsky. Parameters of chelyabinsk and tunguska objects and their explosion modes. arXiv preprint arXiv:1307.1967, 2013.
- [61] Adriano Lopes and Ken Brodlie. Improving the robustness and accuracy of the marching cubes algorithm for isosurfacing. IEEE Transactions on Visualization and Computer Graphics, 9(1):16–29, 2003.
- [62] Laurence Lorda, Elisabet Canalias, Thierry Martin, Romain GARMIER, and Jens BIELE. Mascot: Analyses of the descent and bouncing trajectories to support the landing site selection. In 26th International Space Symposium in Flight Dynamics, Matsuyama, Japan, 2017.
- [63] Benoit Mandelbrot. The fractal geometry of nature. WH Freeman, 1982.
- [64] Clara Maurel, Patrick Michel, Jens Biele, Ronald-Louis Ballouz, and Florian Thuillet. Numerical simulations of the contact between the lander mascot and a regolith-covered surface. Advances in Space Research, 2017.
- [65] Mark Meyer, Mathieu Desbrun, Peter Schröder, and Alan H Barr. Discrete differential-geometry operators for triangulated 2-manifolds. In Visualization and mathematics III, pages 35–57. Springer, 2003. DOI: [dx.doi.org/10.1007/978-3-662-05105-4-2](https://doi.org/10.1007/978-3-662-05105-4-2).
- [66] P. Michel, M. Kueppers, A. Cheng, Carnelli I., and the Hera team. The hera mission: European component of the asteroid impact and deflection assessment (aida) mission to a binary asteroid. In 49th Lunar and Planetary Science Conference. Lunar and Planetary Institute, 2018.
- [67] Patrick Michel, Willy Benz, Paolo Tanga, and Derek C Richardson. Collisions and gravitational reaccumulation: Forming asteroid families and satellites. Science, 294(5547):1696–1700, 2001.
- [68] Tatsuhiro Michikami, Akiko M Nakamura, Naru Hirata, Robert W Gaskell, Ryosuke Nakamura, Takayuki Honda, Chikatoshi Honda, Kensuke Hiraoka, Jun Saito, Hirohide Demura, et al. Size-frequency statistics of boulders on global surface of asteroid 25143 itokawa. Earth, planets and space, 60(1):13–20, 2008.
- [69] JK Miller, AS Konopliv, PG Antreasian, JJ Bordi, S Chesley, CE Helfrich, WM Owen, TC Wang, BG Williams, DK Yeomans, et al. Determination of shape, gravity, and rotational state of asteroid 433 eros. Icarus, 155(1):3–17, 2002. DOI: [dx.doi.org/10.1006/icar.2001.6753](https://doi.org/10.1006/icar.2001.6753).

- [70] David L Mitchell, R Scott Hudson, Steven J Ostro, and Keith D Rosema. Shape of asteroid 433 eros from inversion of goldstone radar doppler spectra. *Icarus*, 131(1):4–14, 1998. DOI: [dx.doi.org/10.1006/icar.1997.5815](https://doi.org/10.1006/icar.1997.5815).
- [71] Akira Miura, Tomohiro Yamaguchi, Rie Honda, Yasuhiro Yokota, Hiroki Senshy, Kohei Kitazato, Yukio Yamamoto, Takanao Saiki, and Satoshi Tanaka. Data production of the "ryuguoid" for operational training of hayabusa2 landing site selection. In Fall meeting. The Japanese Society for Planetary Sciences, 2017.
- [72] Hideaki Miyamoto, Hajime Yano, Daniel J Scheeres, Shinsuke Abe, Olivier Barnouin-Jha, Andrew F Cheng, Hirohide Demura, Robert W Gaskell, Naru Hirata, Masateru Ishiguro, et al. Regolith migration and sorting on asteroid itokawa. *Science*, 316(5827):1011–1014, 2007.
- [73] Hideaki Miyamoto, Hajime Yano, Daniel J Scheeres, Shinsuke Abe, Olivier Barnouin-Jha, Andrew F Cheng, Hirohide Demura, Robert W Gaskell, Naru Hirata, Masateru Ishiguro, et al. Regolith migration and sorting on asteroid itokawa. *Science*, 316(5827):1011–1014, 2007.
- [74] Hirdy Miyamoto. Unconsolidated gravels on asteroid itokawa, 2007.
- [75] Oliver Montenbruck and Eberhard Gill. Satellite orbits: models, methods and applications. Springer Science & Business Media, 2012.
- [76] TG Müller, J Ďurech, M Ishiguro, M Mueller, T Krühler, H Yang, M-J Kim, L ORourke, F Usui, C Kiss, et al. Hayabusa-2 mission target asteroid 162173 ryugu (1999 ju3): Searching for the objects spin-axis orientation. *Astronomy & Astrophysics*, 599:A103, 2017.
- [77] N Murdoch, G Nguyen, A Calandry, O Cherrier, and Y Gourinat. Studying the influence of target and projectile properties on low-velocity collisions. In European Planetary Science Congress, volume 11, 2017.
- [78] Naomi Murdoch, Iris Avila Martinez, Cecily Sunday, Emmanuel Zenou, Olivier Cherrier, Alexandre Cadu, and Yves Gourinat. An experimental study of low-velocity impacts into granular material in reduced gravity. *Monthly Notices of the Royal Astronomical Society*, 468(2):1259–1272, 2017.
- [79] K. A. Newhall and D. J. Durian. Projectile-shape dependence of impact craters in loose granular media. *Phys. Rev. E*, 68:060301, Dec 2003.
- [80] M Nishida, M Okumura, and K Tanaka. Effects of density ratio and diameter ratio on critical incident angles of projectiles impacting granular media. *Granular matter*, 12(4):337–344, 2010.
- [81] Naoko Ogawa, Fuyuto Terui, and Jun'ichiro Kawaguchi. Precise landing of space probe on asteroid using multiple markers. *IFAC Proceedings Volumes*, 43(15):172–177, 2010.
- [82] David Y Oh, Steve Collins, Dan Goebel, Bill Hart, Gregory Lantoine, Steve Snyder, Greg Whiffen, Linda Elkins-Tanton, Peter Lord, Zack Pirkel, et al. Development of the psyche mission for nasas discovery program. In 35th International Electric Propulsion Conference, No. IEPC-2017-153, Electric Rocket Propulsion Society, Fairview Park, OH, 2017.

- [83] Roberta JM Olson. and they saw stars: Renaissance representations of comets and pretelescopic astronomy. Art Journal, 44(3):216–224, 1984.
- [84] Stanley Osher and Ronald Fedkiw. Level set methods and dynamic implicit surfaces, volume 153. Springer Science & Business Media, 2006.
- [85] Steven J Ostro, Jean-Luc Margot, Lance AM Benner, Jon D Giorgini, Daniel J Scheeres, Eugene G Fahnestock, Stephen B Broschart, Julie Bellerose, Michael C Nolan, Christopher Magri, et al. Radar imaging of binary near-earth asteroid (66391) 1999 kw4. Science, 314(5803):1276–1280, 2006. DOI: [dx.doi.org/10.1126/science.1133622](https://doi.org/10.1126/science.1133622).
- [86] Steven J Ostro, Jean-Luc Margot, Lance AM Benner, Jon D Giorgini, Daniel J Scheeres, Eugene G Fahnestock, Stephen B Broschart, Julie Bellerose, Michael C Nolan, Christopher Magri, et al. Radar imaging of binary near-earth asteroid (66391) 1999 kw4. Science, 314(5803):1276–1280, 2006.
- [87] Steven J Ostro, Jean-Luc Margot, Lance AM Benner, Jon D Giorgini, Daniel J Scheeres, Eugene G Fahnestock, Stephen B Broschart, Julie Bellerose, Michael C Nolan, Christopher Magri, et al. Radar imaging of binary near-earth asteroid (66391) 1999 kw4. Science, 314(5803):1276–1280, 2006.
- [88] M Patzold, T Andert, Matthias Hahn, SW Asmar, J-P Barriot, MK Bird, B Hausler, Kerstin Peter, Silvia Tellmann, E Grun, et al. A homogeneous nucleus for comet 67p/churyumovgerasimenko from its gravity field. Nature, 530(7588):63–72, 2016.
- [89] Petr Pravec and Alan W Harris. Fast and slow rotation of asteroids. Icarus, 148(1):12–20, 2000.
- [90] Petr Pravec, Alan W Harris, P Scheirich, P Kušnirák, L Šarounová, Cark W Hergenrother, Stefano Mottola, MD Hicks, Gianluca Masi, Yu N Krugly, et al. Tumbling asteroids. Icarus, 173(1):108–131, 2005.
- [91] DA Rajon and WE Bolch. Marching cube algorithm: review and trilinear interpolation adaptation for image-based dosimetric models. Computerized Medical Imaging and Graphics, 27(5):411–435, 2003.
- [92] Derek C Richardson, Thomas Quinn, Joachim Stadel, and George Lake. Direct large-scale n-body simulations of planetesimal dynamics. Icarus, 143(1):45–59, 2000.
- [93] David Scott Richeson. Euler’s gem: the polyhedron formula and the birth of topology, volume 6. Springer, 2008.
- [94] Reinhard Roll and Lars Witte. Rosetta lander philae: touch-down reconstruction. Planetary and Space Science, 125:12–19, 2016.
- [95] Takanao Saiki, Hirotaka Sawada, Chisato Okamoto, Hajime Yano, Yasuhiko Takagi, Yasuhiro Akahoshi, and Makoto Yoshikawa. Small carry-on impactor of hayabusa2 mission. Acta Astronautica, 84:227–236, 2013.
- [96] Paul Sánchez and Daniel J Scheeres. Simulating asteroid rubble piles with a self-gravitating soft-sphere distinct element method model. The Astrophysical Journal, 727(2):120, 2011.

- [97] Paul Sánchez and Daniel J Scheeres. The strength of regolith and rubble pile asteroids. Meteoritics & Planetary Science, 49(5):788–811, 2014.
- [98] F Sansò. On the aliasing problem in the spherical harmonic analysis. Bulletin géodésique, 64(4):313–330, 1990. DOI: [dx.doi.org/10.1007/BF02538406](https://doi.org/10.1007/BF02538406).
- [99] John Satterly. The moments of inertia of some polyhedra. The Mathematical Gazette, 42(339):11–13, 1958.
- [100] Shujiro Sawai, Kawaguchi Junichiru, Daniel Scheeres, Naoki Yoshizawa, and Masahiro Ogasawara. Development of a target marker for landing on asteroids. Journal of Spacecraft and Rockets, 38(4):601–608, 2001.
- [101] Hanspeter Schaub and John L Junkins. Analytical mechanics of space systems. Aiaa, 2003.
- [102] Daniel J Scheeres. Orbital Motion in Strongly Perturbed Environments. 2012.
- [103] Daniel J Scheeres, Steven J Ostro, R Scott Hudson, Eric M DeJong, and Shigeru Suzuki. Dynamics of orbits close to asteroid 4179 toutatis. Icarus, 132(1):53–79, 1998.
- [104] DJ Scheeres. Landslides and mass shedding on spinning spheroidal asteroids. Icarus, 247:1–17, 2015.
- [105] DJ Scheeres, SG Hesar, S Tardivel, M Hirabayashi, D Farnocchia, JW McMahon, SR Chesley, O Barnouin, RP Binzel, WF Bottke, et al. The geophysical environment of bennu. Icarus, 276:116–140, 2016. DOI: [dx.doi.org/10.1016/j.icarus.2016.04.013](https://doi.org/10.1016/j.icarus.2016.04.013).
- [106] Li Shen and Moo K Chung. Large-scale modeling of parametric surfaces using spherical harmonics. In 3D Data Processing, Visualization, and Transmission, Third International Symposium on, pages 294–301. IEEE, 2006. DOI: [dx.doi.org/10.1109/3DPVT.2006.86](https://doi.org/10.1109/3DPVT.2006.86).
- [107] Christian Sigg, Ronald Peikert, and Markus Gross. Signed distance transform using graphics hardware. In Proceedings of the 14th IEEE Visualization 2003 (VIS'03), page 12. IEEE Computer Society, 2003.
- [108] Steve Squyres. Caesar: Project overview. In 18th Meeting of the NASA Small Bodies Assessment Group, 2018.
- [109] DO Staley. Man on an asteroid. Journal of Geophysical Research, 75(28):5571–5573, 1970.
- [110] William James Stronge. Impact mechanics. Cambridge university press, 2004.
- [111] Dirk J Struik. Lectures on classical differential geometry. Courier Corporation, 2012.
- [112] S Tardivel, E Canalias, M Deleuze, AT Klesh, and DJ Scheeres. Landing mascot on asteroid 1999 ju3: Solutions for deploying nanosats to small-body surfaces. In Lunar and Planetary Science Conference, volume 46, page 1182, 2015.
- [113] Simon Tardivel. The limits of the mascons approximation of the homogeneous polyhedron. In AIAA/AAS Astrodynamics Specialist Conference, page 5261, 2016.
- [114] Simon Tardivel. Optimization of the ballistic deployment to the secondary of a binary asteroid. Journal of Guidance, Control, and Dynamics, pages 2790–2798, 2016.

- [115] Simon Tardivel, Patrick Michel, and Daniel J Scheeres. Deployment of a lander on the binary asteroid (175706) 1996 fg3, potential target of the european marcopolo-r sample return mission. Acta Astronautica, 89:60–70, 2013.
- [116] Simon Tardivel and Daniel J Scheeres. Ballistic deployment of science packages on binary asteroids. Journal of Guidance, Control, and Dynamics, 36(3):700–709, 2013.
- [117] Simon Tardivel and Daniel J Scheeres. Accurate deployment of landers to dynamically challenging asteroids. In American Astronautical Society, AAS 15-424, 2015.
- [118] Simon Tardivel, Daniel J Scheeres, Patrick Michel, Stefaan Van wal, and Paul Sánchez. Contact motion on surface of asteroid. Journal of Spacecraft and Rockets, 51(6):1857–1871, 2014.
- [119] Simon Tardivel, Daniel J Scheeres, Patrick Michel, Stefaan Van wal, and Paul Sánchez. Contact motion on surface of asteroid. Journal of Spacecraft and Rockets, 51(6):1857–1871, 2014.
- [120] Simon Charles Vincent Tardivel. The deployment of scientific packages to asteroid surfaces. PhD thesis, University of Colorado at Boulder, 2014.
- [121] P Thomas, J Veverka, and S Dermott. Small satellites. In Satellites, pages 802–835, 1986.
- [122] Florian Thuillet, Patrick Michel, Clara Maurel, Ronald-Louis Ballouz, Yun Zhang, Derek C Richardson, Jens Biele, Eri Tatsumi, and Seiji Sugita. Numerical modeling of lander interaction with a low-gravity asteroid regolith surface. 2018.
- [123] Akira Tsuchiyama, Masayuki Uesugi, Takashi Matsushima, Tatsuhiro Michikami, Toshihiko Kadono, Tomoki Nakamura, Kentaro Uesugi, Tsukasa Nakano, Scott A Sandford, Ryo Noguchi, et al. Three-dimensional structure of hayabusa samples: origin and evolution of itokawa regolith. Science, 333(6046):1125–1128, 2011.
- [124] Yuichi Tsuda, Sei-ichiro Watanabe, Takanao Saiki, Makoto Yoshikawa, and Satoru Nakazawa. Cruise status of hayabusa2: Round trip mission to asteroid 162173 ryugu. Acta Astronautica, 136:176–181, 2017.
- [125] Yuichi Tsuda, Makoto Yoshikawa, Masanao Abe, Hiroyuki Minamino, and Satoru Nakazawa. System design of the hayabusa 2asteroid sample return mission to 1999 ju3. Acta Astronautica, 91:356–362, 2013.
- [126] R Pl Turco, OB Toon, C Park, RC Whitten, JB Pollack, and P Noerdlinger. An analysis of the physical, chemical, optical, and historical impacts of the 1908 tunguska meteor fall. Icarus, 50(1):1–52, 1982.
- [127] J. S. Uehara, M. A. Ambroso, R. P. Ojha, and D. J. Durian. Low-speed impact craters in loose granular media. Phys. Rev. Lett., 90:194301, May 2003.
- [128] David A Vallado. Fundamentals of astrodynamics and applications, volume 12. Springer Science & Business Media, 2001.
- [129] Stefaan Van wal, Robert Reid, and Daniel Scheeres. Parallellized small-body lander/hopper simulations with distributed contact and procedural noise. In 2017 AAS/AIAA Astrodynamics Specialist Conference. AAS/AIAA, 2017.

- [130] Stefaan Van wal, Simon Tardivel, Paul Sánchez, Darius Djafari-Rouhani, and Daniel Scheeres. Rolling resistance of a spherical pod on a granular bed. Granular Matter, 19(1):17, 2017.
- [131] Stefaan Van wal, Simon Tardivel, and Daniel Scheeres. High-fidelity small-body lander simulations. In 6th International Conference on Astrodynamics Tools and Techniques. ESA, 2016.
- [132] Stefaan Van wal, Simon Tardivel, and Daniel Scheeres. Parametric study of ballistic lander deployment. Journal of Spacecraft and Rockets, 2017.
- [133] Stefaan Van Wal, Simon Tardivel, and Daniel Scheeres. Parametric study of ballistic lander deployment to small bodies. Journal of Spacecraft and Rockets, pages 1–26, 2017/09/29 2017.
- [134] Stefaan Van wal, Simon Tardivel, and Daniel J Scheeres. Exploring Small Body Surfaces with Landed Pods. In 12th International Planetary Probe Workshop, Cologne, Germany, 2015.
- [135] J Veverka, B Farquhar, M Robinson, P Thomas, S Murchie, A Harch, PG Antreasian, SR Chesley, JK Miller, WM Owen Jr, et al. The landing of the near-shoemaker spacecraft on asteroid 433 eros. Nature, 413(6854):390, 2001.
- [136] J Veverka, B Farquhar, Mark Robinson, P Thomas, S Murchie, A Harch, PG Antreasian, SR Chesley, JK Miller, WM Owen, et al. The landing of the near-shoemaker spacecraft on asteroid 433 eros. Nature, 413(6854):390–393, 2001.
- [137] Tycho T Von Rosenvinge, John C Brandt, and Robert W Farquhar. The international cometary explorer mission to comet giacobini-zinner. Science, 232(4748):353–356, 1986.
- [138] Hakon Wadell. Volume, shape, and roundness of quartz particles. The Journal of Geology, 43(3):250–280, 1935.
- [139] Amanda M. Walsh, Kristi E. Holloway, Piotr Habdas, and John R. de Bruyn. Morphology and scaling of impact craters in granular media. Phys. Rev. Lett., 91:104301, Sep 2003.
- [140] Robert A Werner and Daniel J Scheeres. Exterior gravitation of a polyhedron derived and compared with harmonic and mascon gravitation representations of asteroid 4769 castalia. Celestial Mechanics and Dynamical Astronomy, 65(3):313–344, 1996.
- [141] Allen Wilkinson and Alfred DeGennaro. Digging and pushing lunar regolith: Classical soil mechanics and the forces needed for excavation and traction. Journal of Terramechanics, 44(2):133–152, 2007.
- [142] K Willner, J Oberst, H Hussmann, B Giese, H Hoffmann, K-D Matz, T Roatsch, and T Duxbury. Phobos control point network, rotation, and shape. Earth and Planetary Science Letters, 294(3-4):541–546, 2010.
- [143] K Willner, X Shi, and J Oberst. Phobos’ shape and topography models. Planetary and Space Science, 102:51–59, 2014.
- [144] H. Yabuta, N. Hirata, R. Honda, Y. Ishihara, K. Kitazato, and M. Komatsu. Hayabusa2 landing site selection (lss) training: Summary report of scientific evaluation. In Lunar and Planetary Science Conference, volume 49, 2018.

- [145] Hajime Yano, Takashi Kubota, Hiridy Miyamoto, Tatsuaki Okada, D Scheeres, Yasuhiko Takagi, Kazuya Yoshida, Masanao Abe, Shinsuke Abe, O Barnouin-Jha, et al. Touchdown of the hayabusa spacecraft at the muses sea on itokawa. Science, 312(5778):1350–1353, 2006.
- [146] DK Yeomans, PG Antreasian, J-P Barriot, SR Chesley, DW Dunham, RW Farquhar, JD Giorgini, CE Helfrich, AS Konopliv, JV McAdams, et al. Radio science results during the near-shoemaker spacecraft rendezvous with eros. Science, 289(5487):2085–2088, 2000.
- [147] Donald K. Yeomans and Tao Kiang. The long-term motion of comet halley. Monthly Notices of the Royal Astronomical Society, 197(3):633–646, 1981.
- [148] Tetsuo Yoshimitsu. Development of autonomous rover for asteroid surface exploration. In Robotics and Automation, 2004. Proceedings. ICRA'04. 2004 IEEE International Conference on, volume 3, pages 2529–2534. IEEE, 2004.
- [149] Tetsuo Yoshimitsu, Takashi Kubota, Tadashi Adachi, and Yoji Kuroda. Advanced robotic system of hopping rovers for small solar system bodies. In Proc. 11th Int. Symp. Artif. Intel., Robot. Autom. in Space, 2012.
- [150] Tetsuo Yoshimitsu, Takashi Kubota, and Ichiro Nakatani. Minerva rover which became a small artificial solar satellite. In Small satellite conference, 2006.
- [151] Tetsuo Yoshimitsu, Takashi Kubota, Ichiro Nakatani, Tadashi Adachi, and Hiroaki Saito. Hopping rover minerva for asteroid exploration. In Artificial Intelligence, Robotics and Automation in Space, volume 440, page 83, 1999.
- [152] Xu Zhen-Tao, FR Stephenson, and Jiang Yao-Tiao. Astronomy on oracle bone inscriptions. Quarterly Journal of the Royal Astronomical Society, 36:397, 1995.